Hodge Theory For Geometrically Frustrated Magnetism

by

Kristian Tyn Kai Chung

A thesis
presented to the University of Waterloo
in fulfillment of the
thesis requirement for the degree of
Doctor of Philosophy
in
Physics

Waterloo, Ontario, Canada, 2022

© Kristian Tyn Kai Chung 2022

Examining Committee Membership

The following served on the Examining Committee for this thesis. The decision of the Examining Committee is by majority vote.

External Examiner:	Peter P. Orth Associate Professor, Dept. of Physics and Astronomy, Iowa State University
Supervisor(s):	Michel J. P. Gingras Professor, Dept. of Physics and Astronomy, University of Waterloo
Internal Member:	Roger G. Melko Professor, Dept. of Physics and Astronomy, University of Waterloo
Internal-External Member:	Pierre-Nicholas Roy Professor, Dept. of Chemistry, University of Waterloo
Other Member(s):	David G. Hawthorn Professor, Dept. of Physics and Astronomy, University of Waterloo

Author's Declaration

I hereby declare that I am the sole author of this thesis. This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

I understand that my thesis may be made electronically available to the public.

Abstract

We present an analytical formalism based on the concept of discrete differential forms and Hodge theory as a framework for understanding geometrically frustrated magnetism. The primary insight is to treat spins as a 1-form field, for which we define the Helmholtz Hodge decomposition of spin configurations into rotational, irrotational, and harmonic components. As a physical example of its application, we demonstrate how nearest neighbor spin ice—the canonical model of a geometrically frustrated magnet—and its generalizations extended and dipolar spin ice fall neatly within this framework. Our framework clarifies analytically the dumbbell approximation of dipolar spin ice and its relation to projective equivalence, and allows us to write down the first order correction which captures the lifting of rotational mode degeneracy by the dipole interactions We demonstrate how polarized neutron diffraction cross sections can be interpreted within this framework, and in particular how it can be utilized to isolate correlations of the rotational modes. We then construct a new frustrated magnetic model which is complementary to nearest-neighbor spin ice, but whose low temperature spin liquid phase can be described as a 2-form gauge theory with 1-dimensional string-like excitations. Finally, we consider anisotropic Heisenberg models and show that Coulombic physics can be observed in the spin wave fluctuations of a frustrated magnet in a long range ordered phase.

Acknowledgements

First and foremost, I must acknowledge my advisor, Michel Gingras, for supporting me through my Ph.D., giving me the freedom to do the exploration that eventually led to the formalism developed in this work, and for instilling in me a deeper appreciation for the importance of experimental results in guiding theory. I must also express my thanks to my Master's advisor Jeff Greensite, who taught me most of what I know about the lattice gauge theories which in many ways inspired this work, and my undergraduate advisor Sayandeb Basu, for helping me develop a deep passion for physics through high-minded conversations about philosophy, quantum gravity, and the universe. I am indebted as a physicist to all of them.

I also must acknowledge my close friends who I had the privilege to work with over the course of the past five years. Daniel Lozano-Gómez has been a constant foil who I have bounced innumerable ideas off of and collaborated and learned with. Naman Gupta, for the many long walks late at night and discussions about physics and philosophy and the inseparable connection between the physics we study and the lives we live. Aritro Mukherjee, my dear friend and brother, a learning partner and fellow spiritual seeker. And Darren Pereira, who taught me to unify God and science. The brief period when we all worked together was a true joy.

And I must acknowledge the other fellow travellers who I have had the pleasure to work with. Wen Jin, Alex Hickey, Itamar Aharony, Lei Gioia Yang, Jeremy Goh, Cyrus Kulis, Jennifer Reid, and James Chow. Thank you all for your friendship and support.

Dedication

To my family

Table of Contents

List of Figures xiii

List of Tables xxvi

1 Introduction 1

1.1 From Electrons to Classical Spin Models . . . 3

1.1.1 Pyrochlore Rare-Earth Magnets . . . 4

1.1.2 Pseudo-Spin Hamiltonians . . . 5

1.1.3 Classical Spins . . . 6

1.2 Magnetic Order and Disorder: Field Theory . . . 7

1.2.1 Landau Paradigm: Symmetry Breaking and 0-Form Field Theory . . . 7

1.2.2 Frustration, Spin Liquids, Fractionalization and Gauge Fields . . . 8

1.3 Spin Ice and the Coulomb Phase . . . 10

1.3.1 Quantum Spin Ice and Emergent QED . . . 11

2 Methods: Theory of Classical Soft Spins 13

2.1 The Gaussian, Spherical, and Large-N Approximations . . . 13

2.1.1 The Self Consistent Gaussian Approximation . . . 14

2.1.2 The Spherical Approximation . . . 16

2.1.3 The Large-N Approximation . . . 17

2.2 Soft and Hard Spin Configuration Spaces . . . 20

2.2.1 The Interaction Matrix and SCGA Effective Hamiltonian 22
2.2.2 The Physical Significance of λ and the Spectral Decomposition of $\mathcal{G}$ . . . 24
2.2.3 The Low Temperature Behavior of λ_0 in the Thermodynamic Limit . . . 26
2.3 Spatial Symmetries and Representation Theory 29
2.3.1 The Fourier Basis: Irreducible Representations of Translations 29
2.3.2 Bands of the Interaction Matrix 32
2.3.3 Long Wavelength Modes and Point Group Symmetries 33
2.4 Frustration of Soft Spins 35
2.4.1 Unfrustrated Square Lattice 35
2.4.2 Fully Frustrated Clusters: Triangular and Simplex Motifs 37
2.4.3 Edge Sharing Frustrated Units: Triangular vs. J_1-J_2 Square Lattices . . . 40
2.5 Corner Sharing Frustrated Units: Ice Models and Line Graphs 44
2.5.1 Line Graph Lattices and Flat Bands 47
2.5.2 Long Wavelength: Emergent Field Theory and Gauss' Law 48
2.6 Correlations and Neutron Scattering: Pinch Points 55
2.6.1 Structure Factors 56
2.6.2 Neutron Scattering 57
2.6.3 Pinch Points in Spin Ice 59
2.7 Classical Monte Carlo 63

3 Hodge Theory for Geometrically Frustrated Magnetism 64
3.1 Hodge Theory: Discrete Differential Forms 65
3.1.1 Cellular Complexes 65
3.1.2 Orientations 66
3.1.3 Chains and Cochains 67
3.1.4 The Boundary of a Chain 68
3.1.5 Stokes' Theorem and the Differential 70
3.1.6 The de Rham Complex 70

3.1.7 Connection to Vector Calculus: Integration on Chains . . . 71
3.1.8 Homology . . . 73
3.1.9 Cohomology . . . 74
3.1.10 Inner Product of Cochains . . . 76
3.1.11 The Codifferential and Coboundary . . . 76
3.1.12 The Hodge Laplacian and Harmonic Forms . . . 78
3.1.13 The Helmholtz-Hodge Decomposition . . . 78
3.1.14 "Gauge Invariance" and the HHD . . . 79
3.1.15 Computing the HHD: Poisson's Equation . . . 80
3.1.16 Hodge Duality and the Dual Cell Complex . . . 81
3.2 The Laplacian and the Adjacency Matrix . . . 83
3.2.1 The Graph Laplacian . . . 84
3.2.2 Line Graphs and 1-Forms . . . 86
3.3 Spin Models as Discrete Field Theories . . . 88
3.3.1 Discrete 0-Form (Scalar) Field Theory: Ising Ferromagnet . . . 88
3.3.2 Ice Models as Discrete 1-Form Field Theories . . . 91
3.3.3 Helmholtz-Hodge Decomposition and Flat Bands: Rotational and Irrotational Modes . . . 93
3.4 Spin Ice: 1-Form Field Theory on the Diamond Cell Complex . . . 96
3.4.1 Eigenvectors of the Adjacency Matrix . . . 99
3.4.2 HHD Projectors . . . 103
3.4.3 Explicit Computation of the HHD for Spin Configurations . . . 109
3.4.4 Ground State Constraints and Projections . . . 111
3.4.5 Topology of the Constraint . . . 112
3.4.6 Correlations of Rotational and Irrotational Modes . . . 115
3.4.7 Further Neighbor Interactions: Extended Ice Models . . . 117
3.4.8 Dipolar Spin Ice: Dumbbell Model and Projective Equivalence . . . 121
3.4.9 Corrections to the Dumbbell Approximation . . . 125
3.4.10 Corrections to the Flat Bands: Vorticity . . . 129
3.5 Conclusion and Outlook . . . 133

4 Polarized Neutron Scattering as a Probe of Flat Band Physics in Spin Ice Systems 136
4.1 Polarized Neutron Scattering in Spin Ice: Spin Flip and Non Spin Flip Channels . 136
4.2 The Flat NSF . . . 137
4.3 NSF as Rotational Mode Correlator . . . 139
4.4 NSF as Chain-Chain Correlator . . . 141
4.5 The Non-Flat NSF of Dipolar Spin Ice . . . 145
4.6 Hard Spin Corrections to the Flat NSF . . . 147
4.7 Conclusion and Outlook . . . 149

5 The Vorticity Model 151
5.1 The Vorticity Model . . . 152
5.2 Band Structure, Long Wavelength Expansion, and Gaussian Correlations . . . 153
5.3 The 2-Coulomb Phase: Membranes and Vortex Strings . . . 157
5.3.1 Hodge Duality: Spins as 2-Forms . . . 157
5.3.2 The 2-Coulomb Phase: A Soup of Closed Membranes . . . 160
5.3.3 Open Membranes and Vortex Tubes . . . 161
5.4 Thermodynamics with Hard Spins . . . 162
5.4.1 Monte Carlo Cluster Updates: Membrane Moves . . . 163
5.4.2 Specific Heat: An Unexpected Phase Transition . . . 164
5.4.3 Structure Factors Above the Transition . . . 166
5.4.4 AIAO Order by Disorder: Symmetry Breaking Coexisting with a Spin Liquid . . . 168
5.5 Interpolation to the 3D Ising Ferromagnet: Phase Diagram . . . 172
5.5.1 Interpolated Bands . . . 175
5.5.2 Star Moves . . . 175
5.5.3 Phase Diagram . . . 176
5.5.4 Finite Size Scaling . . . 179
5.6 Conclusion and Outlook . . . 184

6 Coulombic Physics In Symmetry-Broken Long Range Ordered Phases **186**
6.1 Pinch Point Features in an Ordered Phase . . . 188
6.2 3-Component Spins on the Pyrochlore Lattice . . . 189
6.2.1 The Heisenberg Antiferromagnet Coulomb Phase . . . 190
6.2.2 Irreducible Representations Analysis . . . 191
6.3 Rank-2 Formalism, Irreps, and Fields . . . 194
6.4 Helmholtz-Hodge Decomposition for 3-Component Spins . . . 199
6.5 Coulomb Physics Below T_c . . . 201
6.5.1 The Easy Axis Model . . . 201
6.5.2 Low Temperature Expansion . . . 202
6.5.3 HHD Analysis of the Hessian Band Structure . . . 203
6.5.4 Correlations and Pinch Points . . . 205
6.5.5 Rank-2 Tensor Field Theory . . . 207
6.5.6 Dynamics, HHD, and Inelastic Neutron Scattering . . . 209
6.6 Conclusion . . . 211

7 Conclusion and Synthesis **214**

References **217**

APPENDICES **234**

A The Pyrochlore Lattice and Coordinate Conventions **235**

B Adjacency Matrices for the Pyrochlore Lattice and Projective Equivalence **239**
B.1 Adjacency Matrices up to Third Neighbor . . . 240
B.2 Adjacency Matrices Beyond Third Neighbor . . . 241
B.3 Corrections to the Projective Equivalence Model of the Dipole-Dipole Interaction Matrix . . . 241

C The 1/*N* Expansion **250**

C.1 Correlation Functions 250

C.2 Saddle Point Expansion 251

C.3 Diagrammatic Corrections to the Correlator 255

D Derivations for *N*=3 Spins **257**

D.1 Classical Low Temperature Expansion 257

D.1.1 Classical Low Temperature Expansion for the Easy Axis Model 259

D.2 Classical Spin Waves 260

D.3 HHD Norm 263

List of Figures

1.1 (a) two tetrahedra of the pyrochlore lattice, left: a spin ice ground state obeying the 2-in/2-out ice rule, right: a spin ice excited state created by flipping a single spin looks like a pair of charges. (b) In the Coulomb phase spins align into closed loops without ends. 11

2.1 The configuration space Ω_{spin} for (a) $N_{\text{spins}} = 2$ and (b) $N_{\text{spins}} = 3$ soft Ising ($N = 1$) spins. The hard spin constraint defines a subset of points in this space $\mathcal{M}_{\text{hard}}$ given by the corners of a cube (or a Cartesian product of N-spheres for arbitrary N). The spherical approximation replaces the hard spin manifold with the sphere in which this cube is inscribed. 21

2.2 The SCGA effective potential for two Ising spins interacting antiferromagnetically as a function of temperature. At high temperature the potential is isotropic and all states are equally likely, at low temperature the potential prefers the lower-energy states. 23

2.3 (a) Temperature dependence of λ for the two-spin antiferromagnet without shifting the diagonal of the interaction matrix, which diverges as $T \to 0$ as ϵ_0/T. (b) Temperature dependence of $\lambda_0 \equiv \lambda + \beta\epsilon_0$, equivalent to the value of λ if we shift the minimum eigenvalue of the interaction matrix to zero, which is finite at all temperature. 24

2.4 The temperature evolution of λ_0 for toy models with extensive (a,b,f) and subextensive (c,d,e) zero-energy density of states. 27

2.5 (a) The bands of the nearest-neighbor model on the square lattice with ferromagnetic (top) or antiferromagnetic (bottom) interactions. When interactions are ferromagnetic, the the band minimum is at zone center, and when the interactions are antiferromagnetic the bands are minimum at the zone boundary. (b,c) In the antiferromagnetic case, the minimum is at the zone corner M point, which corresponds to a plane wave with two-sublattice repetition. This corresponds to the true ground state of the square lattice antiferromagnetic Ising model, i.e. the Néel state. 36

2.6 The SCGA ground state manifold for a single frustrated unit of n spins is the $n-1$-dimensional space orthogonal to $\sum_{i=1}^{n}|i\rangle$. This is depicted for $N=1$ spins with (a) $n=2$, (b) $n=3$, and (c) $n=4$. The 1-dimensional subspace spanned by $\sum_{i=1}^{n}|i\rangle$ is indicated in red and the orthogonal ground state subspace is depicted in blue. For $N=1$ spins with n odd, such as in (b), the soft spin ground state space (blue disk) is not intersected by the hard spin manifold (cube corners), and the hard spin ground states violate the constraint $\sum_{i=1}^{n}S_i=0$. For n even, such as (a) and (c), some of the hard spin states do lie in the soft spin ground state subspace, and those are the hard spin ground states of the cluster. 39

2.7 (a) Bands of the triangular lattice antiferromagnet, which has a minimum at the K-points despite being highly frustrated. (b) The square lattice with second neighbor interactions is unfrustrated despite being constructed of frustrated units. (c) At the special point $J_2/J_1=1/2$ the square lattice model becomes frustrated again, indicated by the completely flat dispersion along the zone boundary. 40

2.8 (a) Ground states of the triangular lattice with hard Ising spins cannot be formed from the minimum eigenvalue wavevector of the interaction matrix, which is a planewave with three sublattice repetition. (b) Standing wave with three symmetry-equivalent ground state wavevectors. (c) Plane wave with wavevector corresponding to the X point. The band is not minimum at this wavevector and yet this is a ground state. 42

2.9 (a) when $J_1=J_2$ the ground state of the system a plane wave corresponding to the X point. The system is unfrustrated despite each crossed square individually is a frustrated unit. (b) The Néel state is not a ground state when $J_1=J_2$ despite satisfying $\sum_{i\in c}S_i=0$ on each cluster. It is a ground state when $J_2=J_1/2$, because at that point the Hamiltonian splits into a sum over frustrated units. (c) An example of a frustrated ground state satisfying the zero-net-spin condition on each crossed square. 43

2.10 Line graph construction of the (a) checkerboard, (b) kagome, and (c) pyrochlore lattices. 46

2.11 Bands of the adjacency matrices for the (a) checkerboard and (b) kagome lattices, demonstrating flat bands at the bottom of their spectrum, with a band touching point at zone center. 47

2.12 (a) The first Brillouin zone (FBZ) of the face centered cubic (FCC) lattice with high-symmetry points labeled, along with a table of their positions as fractions of the cubic reciprocal lattice vectors, with the $(hk0)$ (tan) and (hhl) (purple) high-symmetry reciprocal planes indicated. See Appendix A for detailed figures showing the structure of the pyrochlore lattice and its relation to the FCC lattice. (b) The bands of the pyrochlore lattice in the $(hk0)$ plane. (c) The bands of the pyrochlore lattice in the (hhl) plane. There are two flat bands at the bottom of the spectrum which are flat in the entire zone. 49

2.13 Points group symmetries of the pyrochlore lattice. (a) The site symmetry group D_{3d} of the pyrochlore lattice includes a single C_3 easy axis (black), three mirror planes (one shown in light blue), three C_2 axes (dark blue) each orthogonal to one mirror plane, and inversion through the site. (b) The site symmetry group of the diamond lattice T_d contains four C_3 axes (white), three C_2 axes (black), three fourfold roto-inversion axes (black, coinciding with the C_2 axes), and six mirror planes (light blue). (c) The character table for T_d; columns are equivalence classes and rows are irreducible representations. 50

2.14 A single FCC unit cell of the pyrochlore lattice, which contains two tetrahedra (two diamond lattice sites, white spheres) and four pyrochlore sites, colored black, red, blue, and yellow. One site sits in the middle (black). The remaining sites sit in the center of a face of the cell and are each shared between two FCC cells, thus yielding only four sites per cell. Additional detailed figures for the pyrochlore lattice can be found in Appendix A. 52

2.15 When performing polarization analysis, the neutron beam is polarized along a fixed direction $\hat{\boldsymbol{z}}_{\text{sc}}$. For each scattering wavevector $\boldsymbol{q}$ in the plane orthogonal to $\hat{\boldsymbol{z}}_{\text{sc}}$ (here spanned by $(hh0)$ and $(00l)$), there is a natural orthonormal basis spanned by $\hat{\boldsymbol{z}}_{\text{sc}}$, $\hat{\boldsymbol{x}}_{\text{sc}} \equiv \hat{\boldsymbol{q}}$, and $\hat{\boldsymbol{y}}_{\text{sc}} \equiv \hat{\boldsymbol{z}}_{\text{sc}} \times \hat{\boldsymbol{x}}_{\text{sc}}$. The neutron is insensitive to the components of the magnetic moments along $\boldsymbol{q}$. Components of the magnetic moments along $\hat{\boldsymbol{y}}_{\text{sc}}$, orthogonal to the neutron polarization, can induce spin flipping, which components along $\hat{\boldsymbol{z}}_{\text{sc}}$ do not induce spin flipping. 58

2.16 Structure factors for nearest-neighbor spin ice at high and low temperature, showing the development of singular pinch points features reflecting singularities in the spin-spin correlation matrix. 61

2.17 In attempting to visual three dimensional correlations with two dimensional cuts it is important to not lose track of the three dimensional structure. (a) The $S_Q(\boldsymbol{q})$ structure factor (first column of Fig. 2.16) showing the intersection of the $(hk0)$ and (hhl) planes and their relation to the FBZ. (b) The pinch point in $S_Q(\boldsymbol{q})$ centered at (220), showing that the intensity is pinched along the $(00l)$ (vertical) direction and remains along the two orthogonal directions. 62

3.1 A foam of soap bubbles forms a natural 3-dimensional cellular complex. Photograph by Paul VanDerWerf, distributed under a CC-BY 2.0 license. 65

3.2 An example of the boundary operator acting on a chain. Top: the boundary of an oriented triangular 2-cell is the sum of its three 1-cell edges with orientations (indicated by arrows) determined by the orientation of the "bulk". Bottom: the boundary of an oriented 1-cell is the sum of the two endpoints with opposite orientation, where the terminating endpoint is oriented "in" (filled black circle) and the starting endpoint is oriented "out" (empty circle). Equivalently, it is the *difference* of the 0-cells with the *same* orientation, here both oriented "in". We take the convention that all 0-cells are always oriented "in" in our formulas. . . . 69

3.3 Depiction of homologous 1-dimensional closed contours (cycles) on the 2-dimensional torus. Two 1-cycles are homologous if their difference is the boundary of an oriented 2-dimensional region (red). The two cycles in (a) are both homologous to 0, i.e. the trivial cycle, because they are both independently boundaries of a 2-dimensional region on the surface of the torus. The cycles in (b) and (c) are not homologous to the trivial cycle because each cycle by itself does not bound a 2-dimensional region, but the two cycles in (b) and homologous to each other, and the two cycles in (c) are homologous to each other. 74

3.4 A depiction of the de Rham complex, which is an exact sequence. The kernel of d_k is a subspace of Ω^k, and the image of d_{k-1} is a subspace of the kernel of d_k, which follows from $\mathrm{d}_k \mathrm{d}_{k-1} = 0$. Elements of $\operatorname{im} \mathrm{d}_{k-1}$ are *exact* k-forms and elements of $\ker \mathrm{d}_k$ are *closed* k-forms. Every exact k-form is closed, but not every closed k-form need be exact. Cohomology measures the failure of a closed form to be exact. Elements of H^k are equivalence classes of closed forms which differ by an exact form. 75

3.5 Top: coboundary of a 0-cell in a two dimensional cell complex, which is the sum of all six oriented edges ending terminating at that cell. Bottom: The coboundary of a 1-cell is the sum of all 2-cells containing that cell in their boundary. 77

3.6 (a) The triangular lattice cell complex and its dual cell complex, the honeycomb lattice. (b) To construct the dual complex we start by placing a dual vertex (open circle) inside every top-dimensional cell (triangle) of the original complex. We then insert the dual 1-cells (dashed line) by gluing the boundary to two vertices whose original 2-cells form the coboundary of an original 1-cell. (c) We then insert the dual 2-cells (hexagon) by gluing the boundary along the dual 1-cells (dashed lines) whose original 1-cells form the coboundary of an original 0-cell. In this way the boundary relation of the dual complex reflects the coboundary relation of the original complex. 82

3.7 Depiction of (a) the differential d_0, (b) the codifferential δ_1, and (c) the 0-form Laplacian $\Delta_0 = \delta_1 d_0$. 85

3.8 Demonstration of Eq. (3.71) when edges can be oriented consistently "in" and "out" at every vertex. Triangles indicate edges of the line graph lattice of the underlying parent lattice. 88

3.9 (a) A configuration for discrete 0-form field on a 2D cellular complex. (b) A single fiber (internal space) of the 0-form field attached to a vertex i of the cell complex. 89

3.10 As a basis for the space of rotational modes we can take 1-forms $\delta_2 \alpha_p$, where $\alpha_p(p) = 1$ if p is oriented as shown by the orange arrow and 0 on all other plaquettes. Then $\delta_2 \alpha_p(e)$ on the boundary edges of p is ± 1 depending on the orientation taken for e and 0 on all other edges. 95

3.11 The bipartite diamond lattice is self-dual. Here we show the diamond lattice (red, orange spheres, blue edges) and its dual (gray, white spheres, blue edges). The diamond and its dual together form a BCC lattice. On the left we show a single BCC unit cell, and on the right we show four BCC unit cells which either surround a single diamond site or fill the diamond "cage." These four BCC cells together form the Wigner Seitz cell for the FCC lattice. 97

3.12 (a) The FCC wigner Seitz unit cell can be subdivided into two symmetric halves separated by a buckled hexagon (blue and red faces). (b) The Wigner Seitz cell cell splits into two equivalent cells (red, blue), cf. right hand column of Fig. 3.11. (c) The buckled hexagon forms a 2-cell of the diamond cell complex, shown here along with the 1-skeleton (blue edges, red and orange vertices) of the complex in a single cubic unit cell. (d) A single 3-cell of the diamond lattice cell complex, with four buckled hexagonal 2-cells on its boundary (two of which are not visible). (e) Two 3-cells of the diamond cellular complex formed by the separation shown in (a,b), which are inverted copies of each other. These two cells together for a unit cell for the FCC lattice, containing the same volume as the original Wigner Seitz cell. 98

3.13 A single diamond A sublattice site (red) and its four neighboring B sublattice sites (orange). The four displacement vectors $\boldsymbol{\delta}_\mu$ are shown in gray. The four pyrochlore sites sitting at the midpoints of the diamond bonds are shown as black spheres, forming a tetrahedron surrounding the central diamond site. The vectors $(\boldsymbol{\delta}_{\mu_1} \pm \boldsymbol{\delta}_{\nu_1})/2$ which appear in the matrix elements of the irrotational projector Eq. (3.146) are shown, with the positive sign corresponding to the blue vectors pointing to the cube face centers and the negative sign corresponding to the pyrochlore separation vectors along the edges of the tetrahedron, with one example vector shown in green. 106

3.14 Helmholtz-Hodge decomposition for a 1-form on the diamond cell complex which is supported on a single oriented edge on which it has unit magnitude. The subextensive harmonic portion has been multiplied by L^3, in this example we took $L = 6$. Only the 1-skeleton of the complex is shown. We can visually observe that the irrotational part has non-zero divergence, and the rotational part has zero divergence and forms closed loops. 110

3.15 The divergence-divergence correlator structure factor defined in Eq. (3.183) computed for the nearest-neighbor spin ice model in the SCGA. 116

3.16 The bands of the interaction matrix of the ESI model, with $J_1 \equiv J > 0$ and $J_2 = J_{3a} = J'$. Upper row: $J' < 0$, tuning J' tunes the dispersion of the top band (yellow), until at a critical value of $J'/J = -1/2$ the upper band becomes degenerate with the lower dispersive band, after which it drops below the flat bands, leading to a condensation of the divergence at every site at sufficiently low temperature. Bottom row: $J' > 0$, tuning J' changes the dispersion of both the upper and lower irrotational bands, until at a critical value of $J'/J = 1/6$ the lower band goes soft, after which it disperses quadratically downward. 119

3.17 The matrix elements of $\Delta_{ee'}$ computed from Eq. (3.219) for $L = 60$ up to twelve times the nearest-neighbor pyrochlore separation or 100'th nearest neighbor, which decay as $1/r^5$. 127

3.18 (a) The bands of the nearest-neighbor interaction matrix $\mathrm{d}_0\delta_1$. (b) The bands of dipole-dipole interaction matrix, normalized such that the nearest-neighbor coupling is 1, computed by truncating the correction terms Δ_{n,s_n} to the projective equivalence interaction matrix at 100'th nearest-neighbor. (c) The two rotational bands of the dipole-dipole interaction matrix, showing a very weak dispersion and a slight lifting of the degeneracy of the two bands. 128

3.19 The minimal hexagonal loop motif of the pyrochlore lattice, with the diamond cell complex shown inside the tetrahedra. second and third-type-b neighbors are indicated. Six spins (red arrows) are illustrated forming a closed loop around the hexagon, maximizing the total circulation. 130

3.20 The band spectrum of $\mathcal{J}^{(\omega)}/|J_\omega|$ defined in Eq. (3.223), or equivalently the operator $\delta_2\mathrm{d}_1$. There are two irrotational bands and two dispersive rotational bands which are completely degenerate. 132

3.21 Bands of the vorticity Hamiltonian with negative coefficient (green) compared to the bands of the dipole-dipole interaction matrix computed by truncating corrections at 100'th neighbor (transparent, yellow). The dipole interaction matrix has been normalized so that the nearest-neighbor coupling is unity. The coefficient J_ω was chosen to match the bandwidths. Since the dipole interaction matrix splits the degeneracy of the two rotational bands, we show separate comparisons in (a) and (b) for the two bands, in both $(hk0)$ and (hhl). 134

4.1 The SF, NSF, and unpolarized neutron scattering cross sections in the (hhl) plane for the nearest-neighbor spin ice model computed in the SCGA. The SF channel shows clear pinch points while the NSF is completely flat. 138

4.2 (a) A single tetrahedron of the pyrochlore lattice, with the direct-space $[hhl]$ plane shown and sublattices labeled. The local $\hat{\boldsymbol{z}}_\mu$ easy axis vectors are shown in red, where the central diamond lattice site (in the center of the tetrahedron) is an A site. (b) Depiction of the alpha chains of the pyrochlore lattice, which correspond to 1-dimensional cycles of the diamond lattice cell complex wrapping around the periodic boundaries. When projected to the direct-space $[hhl]$ plane they form a triangular lattice. 142

4.3 The spin flip, non-spin-flip, and unpolarized neutron scattering cross sections computed for the dipolar spin ice model in the SCGA. The non-flat NSF reflects the corrections to the flat bands which lift the rotational mode degeneracy discussed in Section 3.4.10. Notice also the presence of pinch points in the SF even at high temperatures: in the high-temperature series expansion at lowest order one would already obtain pinch points from the fact that the dipolar interaction matrix itself is singular. 146

4.4 (a) The (hhl) NSF of NNSI measured in Monte Carlo at $T/J = 0.01$ for $L = 16$ with 10^{10} samples. (b) $1/N$ corrected (hhl) NSF at $T/J = 0.01$ computed for $L = 12$. 147

4.5 Eigenvalues of the spin-spin correlation matrix measured in Monte Carlo for NNSI at at zero temperature with $L = 16$ averaged over 10^{10} samples, compared with the Large-N, $1/N$ corrected eigenvalues at $T/J = 0.01$ with the same system size. Eigenvalues at each $\boldsymbol{q}$ are sorted by magnitude. The first and second column show weak deviations from the SCGA expectation of flat bands with eigenvalue $\lambda_0^{-1} \to 2$ at zero temperature, with qualitative agreement in the second column. . 148

5.1 Structure factors for the vorticity model computed in the SCGA, which should be contrasted with those of the nearest neighbor spin ice model in Fig. 2.16. 155

5.2 The (220) pinch point in $S_Q(\boldsymbol{q})$ for the vorticity model at $T/J_\omega = 0.01$, demonstrating that the pinch is along two directions, corresponding to the suppression of the two rotational modes, and intensity remains along the longitudinal direction. Compare to Fig. 2.17(b) 156

5.3 (a) A "star" of the diamond lattice cell complex, the arrows represent the 1-form $\mathrm{d}_0 |v\rangle$, where each edge carries unit magnitude. (b) Every vertex (0-cell) sits at the center of a volume (3-cell) of the dual cell complex, and each edge (1-cell) passes through the center of plaquette (2-cell) of the dual complex. 158

5.4 (a) The minimal excitations of the vorticity models are vortex tube. The vorticity 2-form $\omega = \mathrm{d}_1 S$ measures the circulation around each plaquette of the direct lattice. Its Hodge dual $\tilde{\omega}$ is a 1-form on the dual lattice. Flipping a single spin (center, red arrow) in the 2-form description creates a flux through a single plaquette of the dual complex (blue hexagon), and $\tilde{\omega} = \tilde{\delta}_2 \tilde{S}$ encodes the corresponding circulation around this dual plaquette (red arrows). (b) A closed membrane consisting of twelve dual plaquettes, where all twelve spins have the same orientation relative to the membrane surface. Flipping these spins will result in a new configuration with the same energy. (c) An open oriented membrane. Flipping this membrane in the ground state results in an closed string excitation localized along the boundary with energy proportional to its length. In the direct lattice this excitation is a vortex tube. 161

5.5 Specific heat per spin of the vorticity Ising model as a function of temperature in units of J_ω, measured in Monte Carlo simulations at various systems sizes varying from $L = 6$ to $L = 32$. (a) Broad temperature range with 10^5 sampled spin configurations per temperature, showing a sharp peak which grows with system size, indicating a phase transition. (b) Detailed structure of the peak at various system sizes with 10^7 samples per temperature. 165

5.6 Entropy per spin computed from the Monte Carlo specific heat data, using $L = 32$ data for temperatures near the specific heat peak and $L = 20$ data away from the peak. The red line shows $S(\infty) = \ln(2)$. For $T/J_\omega > 100$ the entropy is estimated by fitting the high-temperature tail of the specific heat, yielding a very small contribution of $S(\infty) - S(100J_\omega) \approx (2.7 \pm 0.1) \times 10^{-4}$. The green line shows Pauling's estimate of $(1/2)\ln(3/2) \approx 0.202733$, while we obtain a slightly larger value of $S(0.1J_\omega) \approx 0.2061 \pm 0.0001$. 165

5.7 Structure factors computed at three temperatures above $T_c/J_\omega \sim 1.15$ in Monte Carlo for $L = 16$ with 10^5 samples, and at $T/J_\omega = 0.01$ in the Large-N with the $1/N$ correction. 167

5.8 Top: Structure factors for the vorticity model at $T/J_\omega = 1.0$, just below the transition at $T_c/J_\omega \sim 1.15$, for $L = 16$ with 10^5 samples. Bottom: the corresponding structure factors for the fully ordered AIAO state, with intensity divided by L^3. Bragg peaks for the vorticity model below T_c appear at the positions and with the relative intensities expected for AIAO order, but their intensity contains only a fraction of the total spectral weight, and we can still see the spin liquid correlations in the background. Note that we have scaled the color bar intensity logarithmically to highlight the low-intensity background in the top panels. 169

5.9 All-in-all-out order parameter as a function of temperature for system sizes varying from $L = 6$ to 32. (a) Broad temperature range with 10^5 sampled spin configurations per temperature, showing that the order parameter saturates , indicating a phase transition. (b) Detailed structure of the peak at various system sizes with 10^7 samples per temperature. 170

5.10 Monte Carlo structure factors at $T = 0$ obtained by starting from one AIAO configuration and only applying star moves randomly to flippable stars. Top: We observe the same Bragg peaks as seen in Fig. 5.8, which serves as a partial validation for the theory that the symmetry breaking is due to a subextensive bottleneck in the ground state manifold. Bottom: Adjusting the color scale to highlight the low-intensity background, beneath the Bragg peaks we observe the correlations of the 2-Coulomb phase, which contains the large majority of the spectral weight. Notice that the pinch points intensity zone centers is zero, because we are always in the zero-magnetization sector and do not have harmonic mode fluctuations. 173

5.11 The band structure of the interpolated interaction matrix when tuning from the vorticity model at the left end to the nearest-neighbor Ising ferromagnet at the right end. Top row: (hhl), bottom row: $(hk0)$. Due to the choice of Hamiltonian the rotational (green) and irrotational (red,orange) bands remain distinct, since we tune the interaction in the sectors separately. At the left end there are two flat irrotational bands and two dispersive rotational bands, while at the right end there are two flat rotational bands and two dispersive irrotational bands Note that the degeneracy of the threefold band touching point is preserved throughout, and furthermore the two vorticity bands remain degenerate throughout. 174

5.12 (a) Specific heat for the Hamiltonian Eq. (5.20) as a function of temperature T/J and the interpolation parameter x, measured in Monte Carlo for $L = 16$ with 10^5 sampled configurations per temperature and x. At the left end ($x = 0$) is the vorticity model and at the right end ($x = 1$) is the AIAO ferromagnet. Left: linear-linear scale. Right: log-log scale. (b) The AIAO order parameter as a function of T/J and x. The partially ordered 2-Coulomb phase is visible as the light blue region. 177

5.13 Cuts at various x of the specific heat and AIAO order parameter with $L = 16$ and 10^5 samples per temperature. 178

5.14 The variance of the AIAO order parameter as a function of temperature and x for $L = 16$ with 10^5 samples per data point. For $x > x^*$ this quantity exhibits a strong peak, indicating diverging fluctuations of the AIAO order parameter in the thermodynamic limit. For $x < x^*$ however, it exhibits no divergence whatsoever. . 180

5.15 System size dependence of (a) the AIAO order parameter, (b) the specific heat, (c) the AIAO variance. 182

5.16 (a) The AIAO Binder cumulant defined in Eq. (5.25) at various system sizes and for various values of x. At large x the Binder cumulants for different system sizes cross near a single temperature, even when $L = 4$. (b) The Binder cumulant plotted vs $L^{1/\nu}(T - T_c)/T_c$, where T_c/J is estimated from the specific heat peak and $\nu = 0.63$ corresponding to the critical exponents of the 3D Ising universality class. 183

6.1 Unpolarized neutron scattering cross section in the (hhl) plane measured in Monte Carlo for the easy axis model. At high temperatures $T \gg T_c$ the system is paramagnetic with weak antiferromagnetic correlations. Just above the critical temperature $T \gtrsim T_c$ the system has begun forming pinch points indicative of the correlations of a Coulomb phase, but this spin liquid is co-opted by a phase transition to long range order. Deep in the ordered phase $T \ll T_c$ sharp Bragg peaks have formed (green points), but remnant pinch-point-like features remain in the correlations of fluctuations about the ordered configuration (note the intensity scale). 189

6.2 Symmetry broken ground states of the nearest-neighbor anisotropic Heisenberg model labeled by their corresponding irreps. The E irrep has a continuously degenerate ground state manifold, where all four spins can simultaneously rotate about their local [111] axes, indicated by circles, from which the high-symmetry ψ_2 and ψ_3 are selected by order by disorder. The T_2 irrep ground states (also called Palmer Chalker states) form a spiral when viewed along a Cartesian axis (black vertical line). 192

6.3 Illustration of the T_1 ground states as the canting angle is varied. At the maximum canting angle $\theta = \cos^{-1}(1/\sqrt{3})$, $T_{1\theta}$ corresponds to the ice configurations with spins along the local [111] axes, and $T_{1\bar{\theta}}$ has spins in the local xy planes. At $\theta = \pi/2$ the $T_{1\parallel}$ and $T_{1\perp}$ irrep swap. 199

6.4 (a) The four bands of the pyrochlore adjacency matrix $\mathcal{A}$, or equivalently the twelve bands of $\mathcal{J}_{\text{Heis}}$, in the $(hk0)$ and $(hh\ell)$ planes, restricted for the first Brillouin zone (FBZ) of the fcc lattice. (b) Line cuts of the bands, with minimum eigenvalue shifted to zero. The two degenerate flat bands (blue) correspond to rotational modes, while the two dispersive bands (red) correspond to irrotational modes. . . 200

6.5 (a) The eight bands of the Hessian for the easy axis model with $D = 0.25J$ in the FBZ in the $(hk0)$ and $(hh\ell)$ planes, colored by the irrotational projection of the corresponding eigenvectors, $\mathbb{P}(\phi_{\boldsymbol{q},n})$. (b) Corresponding line cuts of the Hessian bands along a high-symmetry path, the same as Fig. 6.4, again colored by HHD projection. There are two degenerate fully rotational (blue) flat bands at the bottom of the spectrum with $\epsilon_{\text{flat}}(\boldsymbol{q}) = D$. The upper bands, with energy of order J, are dominant irrotational (red). 205

6.6 Spin structure factors in the $(hk0)$ (left) and (hhl) (right) planes for (a) the Heisenberg AFM Coulomb phase; (b) the easy axis model below the critical temperature, computed from the Hessian using Eq. (6.44); (c) the easy axis model restricting the sum in Eq. (6.42) to only the upper six dispersive bands; (d) the same but for only the bottom two flat bands. Adding the contributions of (c) and (d) yields the total structure factor in (b), with the complementary fourfold pinch points combining to create non-singular, broadened pinch point features. 206

6.7 (a,b) Magnon frequency bands colored by irrotational projection of their eigenvectors in (a) the $(hk0)$ plane and the $(hh\ell)$ plane, (b) along with line cuts along a high symmetry path in reciprocal space (cf. Fig. 2.12). The two lower bands are strongly rotational (blue), while the upper two bands have mixed character, with rotational and irrotational modes meeting at a threefold band touching point. (c) Inelastic neutron scattering intensity from the bottom three bands (upper band removed for clarity) in the $(hk0)$ plane centered at (020) and the $(hh\ell)$ plane centered at (002), showing singularities at the threefold band touching point. (d) integrated scattering intensity of the upper two (top) and lower two (bottom) bands in the $(hk0)$ (left) and (hhl) (right) planes, demonstrating that singular pinch points are observable via energy integration of inelastic neutron scattering data. Intensity units are arbitrary. 212

A.1 Pyrochlore lattice cubic conventional cell. Cubic lattice vectors are shown in red, FCC lattice vectors are shown in blue. Sites of the pyrochlore lattice are shown as colored spheres, with four sublattices per FCC unit cell colored accordingly. . . 236

A.2 Demonstration of the oxygen positions (white spheres) in the $A_2B_2O_7$ pyrochlore compounds. Red spheres are *A* sites, blue spheres are *B* sites, black spheres are oxygen vacancies. 238

List of Tables

1.1 A summary of the differential forms appearing in Maxwell's equations and their physical interpretation. See [1] for a pedagogical introduction. 2

B.1 Further neighbor symmetry classes up to 100'th nearest neighbor. 242

B.2 Further neighbor symmetry classes up to 100'th nearest neighbor. 243

B.3 Further neighbor symmetry classes up to 100'th nearest neighbor. 244

B.4 Further neighbor symmetry classes up to 100'th nearest neighbor. 245

B.5 Further neighbor symmetry classes up to 100'th nearest neighbor. 246

B.6 The further-neighbor couplings appearing in the difference Δ between the magnetic dipolar interaction matrix and the projector, up to 100'th neighbor. 247

B.7 The further-neighbor couplings appearing in the difference Δ between the magnetic dipolar interaction matrix and the projector, up to 100'th neighbor. 248

B.8 The further-neighbor couplings appearing in the difference Δ between the magnetic dipolar interaction matrix and the projector, up to 100'th neighbor. 249

Chapter 1

Introduction

At the heart of physics, lies geometry. Perhaps the most fundamental physical observation about the world we live in is its geometry: the universe is, at least locally to us humans, a three dimensional Euclidean space. This basic fact forms the basis for Newtonian mechanics and all of the classical mechanics that follows from it. Einstein took this geometric picture of physics to the extreme, and showed that the universe is in fact described as a curved four dimensional pseudo-Riemannian manifold with the theory of General Relativity. In doing so physicists were confronted with the mathematical description of smooth geometric spaces, the language of differential geometry, whose objects of study are smooth manifolds and their geometric structure. But this language of smooth geometry is by no means confined to the study of General Relativity, and indeed suffuses many areas of physics, especially when it comes to describing field theories, which describe fluctuating collective degrees of freedom filling space. Field theory is perhaps the most powerful tool physicists have for describing the universe, encompassing everything from classical electrodynamics and fluid flow, to quantum field theory and the Standard Model of particle physics.

The primary insight of interest to us from differential geometry is the use of differential forms to describe physical quantities. A differential form is a natural generalization of a function on a manifold which properly encodes the physicist's notion of a density. Although differential forms are largely excluded from undergraduate physics curriculum in favor of vector calculus, they are a more natural mathematical object to use even from a physical perspective. Take for example classical electromagnetism. The electric field E and the magnetic field B in the vector calculus language are both treated as vector fields, but these are quite different quantities. The electric field can be integrated along a 1-dimensional contour to obtain the net work required to displace a unit charge (a 0-dimensional object) along that contour, thus E should be understood as a so-called 1-form. On the other hand, integrating B along a contour is less sensical, instead B should be integrated over a 2-dimensional surface to obtain the total flux through that surface, and thus

Symbol	Name	SI unit	Degree	Differential Form Description
E	Electric Field	V/m	1-form	Work per unit charge per unit length
H	Magnetic Field	A/m	1-form	Work per unit charge per unit length
D	Electric Induction	C/m^2	2-form	Electric flux per unit area
B	Magnetic Induction	Wb/m^2	2-form	Magnetic flux per unit area
j	Current Density	A/m^2	2-form	Current per unit area
ρ	Charge Density	C/m^3	3-form	Electric charge per unit volume

Table 1.1: A summary of the differential forms appearing in Maxwell's equations and their physical interpretation. See [1] for a pedagogical introduction.

B is a 2-form. In Table 1.1, we summarize the primary quantities playing a role in Maxwell electromagnetism and their differential form interpretation. Notice that even the SI units reflect their meaning as forms!

In the study of condensed matter systems, however, the use of a continuum field description is often somewhat ill-defined, because the underlying systems we are interesting in describing are fundamentally discrete collections of atoms in a material. Field theories in condensed matter are often justified through a coarse-graining prescription or long-wavelength approximation, and such approaches are hugely powerful through the application of renormalization techniques, but it is often desirable to have an intrinsic description of our system at the microscopic level without losing any information in the process of coarse-graining or integrating out short wavelength degrees of freedom.

Mathematicians, on the other hand, generally have no issues working with discrete spaces and systems. Indeed discrete spaces can often be more powerful to the mathematician because they avoid the numerous issues when trying to deal directly with the continuum. This is most clear in the field of combinatorial topology, which seeks to define topological invariants to classify topological spaces often based on discrete structures. In general, the topological properties of a space may be independent of its smooth structure (i.e. the "differential" part of "differential geometry"), and can be computed more easily by considering a triangulation of the space. Thus, a triangle (without interior) is topologically a circle and a tetrahedron (without interior) is topologically a sphere. The most famous such topological invariant is the Euler characteristic, $\chi = \sum_{k=0}^{d}(-1)^k n_k$, where n_k is the number of k-dimensional "faces" (0-dimensional vertices, 1-dimensional edges, 2-dimensional faces, etc.) of a discrete topological space and d is its dimension. The next most basic topological invariants are the homology groups (originally introduced by Emmy Noether) which classify the k-dimensional holes in a topological space and are most straightforwardly computed from a

triangulation of a space [2]. The next step in this topological hierarchy is *co*homology, which in its most basic conception is defined using differential forms, where it is referred to as de Rham cohomology and is a central object of study Hodge theory [3]. Hodge theory studies the topology of manifolds by studying the kernel of the Laplace operator, and many of these ideas were brought over to mathematical physics by Michael Atiyah to study Yang Mills theory [4] and later topological quantum field theories [5, 6] and more [7]. Much of Chapter 3 will be devoted to laying out the basics of discrete Hodge theory using discrete versions of smooth differential forms, which could be called "discrete forms" but are called "cochains" by mathematicians [3].

Returning now to physics, the study of discrete versions of field theories has a long history stemming from the desire to study non-perturbative effects in quantum chromodynamics (QCD) and non-Abelian Yang Mills theory dating back to the work of Wilson [8, 9] and Creutz [10, 11]. Furthermore, the study of topology in condensed matter has exploded in the past two to three decades [12–21], and a number of recent works have attempted to construct topologically ordered phases using discrete building blocks [22–24], which has a similar flavor to what we will present here.

In this thesis, we will motivate the use of a discrete differential forms based on Hodge theory to describe the physics of geometrically frustrated magnets, with the primary insight being that in a certain class of models called ice models, the spin degrees of freedom should be treated as a discrete 1-form field. The basis for this perspective already exists in the literature, starting from the construction of a lattice gauge field description of *quantum* spin ice [25]. We will take a step back and demonstrate that the underlying *classical* systems are already described very naturally in the language of discrete differential forms. This will allow us to unify a number of results in the field in a single framework based on discrete Hodge theory, which will be described in Chapter 3. We will then, in Chapter 4, show how this perspective allows us to understand neutron scattering cross sections, and motivate the use of neutron polarization analysis as a tool to study certain aspects of these systems. In Chapter 5, motivated by the insights gleaned from the discrete Hodge perspective, we will introduce a new frustrated magnetic model which has novel properties. Lastly, in Chapter 6, we will study the physics of frustrated magnets in an ordered phase using the same ideas to expose non-trivial behavior in the low-temperature excitations around a classical ground state.

1.1 From Electrons to Classical Spin Models

If there is one slogan for condensed matter physics, it is "emergence." Down at the Planck scale, we know that our standard understanding of the universe in terms of continuum fields and smooth space must break down, but we can't say what the physics at that scale really looks like. On the

scales that we can probe, however, particle physics has classified essentially all of the building blocks of matters, described in the Standard model by quarks, electrons, neutrinos, gauge bosons, and the Higgs boson. We know that quarks bind into hadrons, and that hadrons then form atomic nuclei, which bind with electrons to form atoms, which bind together to make the world around us. The fundamental degrees of freedom and forces once we have arrived at the scale of a condensed matter system (say, an angstrom) are essential just nuclei, electrons, and photons. Quarks and gluons are confined, weak force carriers are heavy with a lengthscale less than the nuclear distance, neutrinos are anti-social, and all the heavier generations are too unstable to matter at the scale of condensed matter physics, so the nuclei, electrons, and photons are all we have. And yet with such basic building blocks an endless host and variety of complex phases and phenomena emerge through strong correlations and collective behaviors. Apparently nuclei, electrons, and photons are more than enough to obtain emergent phases of matter more complex than particle physicists ever dreamed possible. It seems likely, at least to the author, that the frontiers and great unknowns of the Planck scale and quantum gravity will inevitably be pushed forward by the study of emergent phenomena in condensed matter systems [26].

In this thesis, we will discuss a few of the phenomena of frustrated magnetism. In doing so, we will encounter emergent quantum electrodynamics, which motivates us to develop a theory based on treating spins as 1-forms, and we will eventually even see an emergent 2-form gauge theory, which has no analogues in the Standard Model. But before we dive into frustrated magnetism we will quickly sketch the path from electrons to the spin models we will be studying.

1.1.1 Pyrochlore Rare-Earth Magnets

Electrons in condensed matter come in two forms: localized or delocalized (insulating or conducting). We will only be studying localized degrees of freedom, but it is worth mentioning that some of our results may apply to delocalized systems as well. In the delocalized case one can either start from a perspective of nearly-free electrons moving in the potential induced by the nuclei (usually forming a crystalline lattice), or start from a simple tight-binding approximation where electrons hop from atom to atom.[1] Almost all of the emergent phenomena of condensed matter physics can then be understood as collective behaviors of a sea of electrons interacting via Coulomb interactions, both with the lattice and with each other. The emergent behaviors are strong correlation effects induced by a combination of interactions and quantum entanglement effects in the many body wavefunction. In principle, a wide range of such systems can be understood via the

[1] Electrons in fully filled shells of the atoms generally do not contribute to the many body physics. This warrants a slight restatement that condensed matter physics is built upon atomic nuclei, electrons, and photons. A more precise list would be ions with closed shells, electrons in unfilled shells, and photons. Indeed we don't usually deal with the photons directly, instead they simply mediate the practically instantaneous Coulomb potential.

Hubbard interaction, the on-site Coulomb repulsion of two electrons with opposite spin (same spin being forbidden by Pauli). In practice, the ground state of such a many body system can be studied through a combination of density functional theory plus various methods for treating interactions such as dynamical mean field theory.

We will not be studying such systems, instead we start from a fully localized picture. Even more specifically, our interest is magnetic rare-earth insulators, where the unpaired electrons sit in a partially filled f-shell, which can hold up to 14 electrons [27]. Starting from the single ion, the electronic ground state is deduced phenomenologically according to Hund's rules. Noting that spin orbit coupling is strong for f electrons, the free ion ground state is characterized by the total angular momentum quantum number j and has degeneracy $2j + 1$ due to rotational symmetry. This ion then sits inside a lattice of other ions, which in the simplest crystal field approximation are treated as point charges [28, 29], acting as an external potential for our rare-earth ion. This potential reduces the symmetry and thus splits the $2j + 1$ single-ion ground state degeneracy into a hierarchy of energy levels. In the case of the pyrochlore lattice, which is our primary interest, the local symmetry group is D_{3d}. For most of the Lanthanide ($4f$) ions in their 3+ ionization state we are left with a well-isolated degenerate ground state doublet, in particular for Dy^{3+}, Ho^{3+}, Er^{3+}, Yb^{3+} [30], Pr^{3+} [31], Nd^{3+} [32], and Ce^{3+} [33]. The wavefunctions of this doublet is a linear combination of the free ion ground states,

$$|\pm\rangle = \sum_{m=-j}^{j} c_m^{(\pm)} |j, m\rangle \,, \tag{1.1}$$

where the gap to the first excited crystal field level is of the order 50-100K [30]. For temperatures well below this, the excited crystal field levels become irrelevant to the low temperature physics. In the case of Kramers ions (with an odd number of electrons) the ground state doublet degeneracy is protected by time reversal symmetry, while for non-Kramers ions it can be lifted by effects which lower the local symmetry further, primarily disorder or an external field.

1.1.2 Pseudo-Spin Hamiltonians

Once we have our ground state doublets, we are left with what amounts to an array of qubits. The Hilbert space of any 2-state system is equivalent to the Hilbert space of a spin-1/2, so it is convention to refer to the resulting system as a pseudo-spin. Since the pyrochlore lattice has a single high-symmetry axis for each magnetic ion it is natural to choose this axis as the quantization axis, and denote the operators on the Hilbert space by S^z and $S^\pm$. While these effective degrees of freedom are not spins, they are related to the physical magnetic moment $\boldsymbol{m}$ of the ion through the

g-tensor [28],

$$m^\alpha = \mu_B \sum_\beta g^{\alpha\beta} S^\beta, \qquad g = g_\parallel \hat{z}\hat{z}^T + g_\perp (\hat{x}\hat{x}^T + \hat{y}\hat{y}^T), \tag{1.2}$$

where μ_B is the Bohr magneton, $\hat{z}$ is the local high-symmetry axis, and $\hat{x}$ and $\hat{y}$ are the local orthogonal directions. The values of $g_\parallel$ and $g_\perp$ are generally determined experimentally.

We can now construct a Hamiltonian to describe the interactions of these qubits. The physical interactions between the magnetic ions in pyrochlore insulators predominantly arise from either super-exchange mediated by an intermediate non-magnetic oxygen ion or via long-range dipole-dipole interactions. The various interaction pathways can be considered explicitly [34, 35], but for our purposes and from a modelling perspective, it is sufficient to consider the most general nearest-neighbor bilinear spin Hamiltonian which is consistent with the lattice symmetries. There are three possibilities depending on which coefficients $c_m^{(\pm)}$ are non-zero in Eq. (1.1), because the $|j, m\rangle$ transform differently under the lattice symmetries for different m. They are referred to as effective spin-1/2, dipolar-octupolar, and non-Kramers [36].

1.1.3 Classical Spins

In this thesis, we will almost entirely focus on classical approximations of the underlying quantum pseudo-spin-1/2 Hamiltonians, for reasons which we will discuss rather extensively in Chapter 2. There are essentially two approximations we work with: classical Ising models in Chapters 2 to 5 and classical anisotropic Heisenberg models in Chapter 6. The two correspond to very different classical approximations. A classical Ising model can be viewed as a spin-1/2 system in which there are no off-diagonal (non-commuting) terms in the Hamiltonian, and thus the eigenstates of the Hamiltonian can be specified in the S^z basis. These systems can thus be described by classical Ising variables $S_i = \pm 1$, where i indexes the spins. Classical Heisenberg models, on the other hand, can be understood as the $S \to \infty$ limit of a quantum spin, which can be represented by a unit vector $\boldsymbol{S}_i$ which can rotate continuously on the surface of a sphere. We will simply refer to the pseudo-spins as spins from here on and primarily think of them as classical spins, because in the cases of interest to us the classical approximation is highly accurate and a thorough understanding of the classical models allows for a more detailed understanding of where the quantum effects truly come into play.[2]

[2]To be specific, the classical spin ice compounds which we will discuss have large-j ground states, with the the dominant terms in the doublet Eq. (1.1) being the maximal $m = j$ terms. These are strongly Ising-like doublets, and the off-diagonal $S_i^\pm$ pseudo-spin operators which would generate quantum fluctuations would have to originate from operators in the full crystal field Hamiltonian connecting $|j, j\rangle$ with $|j, -j\rangle$ originating from high-multipole interactions which are expected to be very weak [35].

1.2 Magnetic Order and Disorder: Field Theory

At high temperatures relative to the interaction energy scale between neighboring spins J, which we will assume is much smaller than the first excited crystal field level, a magnet is trivially disordered. This is the paramagnetic phase with weak correlations between the orientations of the local spins. The paramagnetic phase is analogous to a gas. The value of J can generally be estimated by fitting the high-temperature magnetic susceptibility to a Curie law (because the system is highly polarizable in the paramagnetic phase due to weak inter-ion correlations), from which we obtain the Curie-Weiss temperature Θ_{CW}, whose sign provides a naive estimate of whether the interactions are predominantly ferromagnetic (positive) or antiferromagnetic (negative).[3] In almost all cases, however, the system will order at some critical temperature generally on the order of the Curie-Weiss temperature. There are a plethora of possibilities for magnetic ground states, ferromagnets, anti-ferromagnets, spin spirals, and so on. This is where field theory enters the picture.

1.2.1 Landau Paradigm: Symmetry Breaking and 0-Form Field Theory

Many phase transitions, and in particular many magnetic ones, can be understood in the framework of Landau symmetry breaking [37, 38], where the system spontaneously breaks a symmetry described by a group G such as global spin rotation (e.g. $G = SU(2)$ or $SO(3)$) symmetry by spontaneously ordering ("condensing") into a phase where the symmetry has been reduced to a subgroup $N \subset G$. The classic example is a Heisenberg ferromagnet, which classically has $G = O(3)$ rotation symmetry and spontaneously orders into a collinear state where the symmetry is reduced to $N = O(2)$ (rotation about the collinear axis). The ground state manifold is then given by the quotient space $G/N = O(3)/O(2)$, which is topologically a sphere [39]. The general prescription for the description of the phase transition is (i) identify an order parameter, which we will call ϕ, which is zero in the disordered phase and non-zero in the ordered phase and has the symmetries of the Hamiltonian, (ii) use this order parameter as an expansion parameter for the free energy in the vicinity of the critical temperature. Now, a physical system does not generally spontaneously order everywhere at once, so we wish to introduce some spatial dependence to the order parameter. Roughly speaking what this means is that we treat our thermodynamically large system as a collection of smaller patches, in each of which the order parameter is approximately constant. The free energy of the system can then be expanded as

$$F \sim \int d^d \boldsymbol{r} \left[(\nabla \phi)^2 + m^2 \phi^2 + g \phi^4 + \cdots \right], \tag{1.3}$$

[3] This statement must be taken with a heavy grain of salt. It is generally true for isotropic systems, but under a variety of anisotropic conditions it may be highly misleading.

where the integral serves as the continuum limit of a sum over patches, and the gradient term penalizes rapid spatial variations. Thus what we wind up with is a field theory describing the fluctuations of the order parameter near the critical temperature.[4] A key point here is that this is a 0-form field theory, meaning that the field $\phi(\boldsymbol{r})$ takes values at each point in space like a function, and is not meaningfully integrated along lines or surfaces. In other words, $\phi(\boldsymbol{r})$ is what a particle physicist would call a matter field, it represents the density of the order parameter at a point (or in a volume around that point).

The other possibility to obtain field theories in condensed matter systems is deep in the ordered phase as a low temperature expansion about the ground state of the system. We will see an example of this in Chapter 6 in the form of spin wave fluctuations of an ordered magnetic phase. The field no longer describes an order parameter, but rather the density of deviations from the ground state of the system. Once again, in most cases such a field is a 0-form field.

1.2.2 Frustration, Spin Liquids, Fractionalization and Gauge Fields

Frustration is the primary mechanism by which one can avoid the symmetry breaking scenario, or at least suppress it to very low temperatures, to obtain different types of physics which are not described by some 0-form field theory. Frustration has many manifestations, but the general idea can be stated as an obstruction to a system being able to satisfy the energetic requirements of all of its parts simultaneously [40, 41]. What this means for spin models is generally that the pairwise interactions of spins cannot all be minimized. The classic example is to try to color the corners of a triangle black and white without two neighbor being the same color (or place three spins on the corners antiferromagnetically aligned). The best that can be done is to satisfy two of the edges while the third is unsatisfied, which can be done in six possible ways. Thus frustration tends to increase the low energy density of states, and in the most extreme cases it can lead to an extensive[5] ground state degeneracy, which we will discuss in much more detail in Chapter 2.

There are generally two ways to achieve frustration: via random disordered interactions, or via exploiting geometry such as in the example of the antiferromagnetic triangle. The former tends to lead to glassy behavior [40], where (very roughly speaking) the system has a very large number of low energy "pockets" of metastable configurations, separated by energy barriers, and becomes trapped at a critical freezing temperature in one of these pockets. [42]. The latter is referred to as geometric frustration [43], and is the topic of the remainder of this thesis.

[4] This action is only schematic and serves as a cartoon, there could be a cubic term, the order parameter could have multiple components, etc.

[5] Extensive is meant here in the thermodynamic sense of proportional to the size of the system.

In the context of magnetism, geometric frustration tends to lead not to glassy behavior but rather *spin liquid* behavior [44–46]. Heuristically, if a paramagnet is analogous to a gas and a long range ordered magnetic phase (e.g. a ferromagnet) is akin to a solid, then a spin liquid is like its namesake, a liquid. Spin liquids can generally be characterized as low temperature disordered phases, often adiabatically connected to a paramagnetic phase, with non-trivial strong correlations. They are essentially always associated with some sort of *fractionalization*, where their elementary excitations carry fractional quantum numbers due to being collective excitations of the liquid [46, 47]. For example, in a spin-1/2 ferromagnet, flipping a single spin creates a single magnon excitation carrying spin-1, while in a spin-1/2 spin liquid a single spin flip will produce two spinons, each carrying spin-1/2 [46].

Unlike the symmetry broken phases we described previously, spin liquids are generally not described by 0-form field theories, but rather by *gauge field theories*, such as those describing the fundamental forces in the Standard Model [48]. There are many reasons for this, none of which is entirely satisfactory, but we will offer just one. In the example of the Landau order parameter field describing the ordering density, or the low temperature expansion field describing the density of fluctuations away from a fully ordered ground state, we know precisely what field configuration we are expanding about. We always define the field to be *zero* in either the disordered phase for the order parameter example or in the fully ordered ground state for the low temperature expansion example. In a spin liquid there is generally no such unique zero state of the field to expand about. In a classical spin liquid there are an extensive number of ground states which the system can transition between, and in a quantum spin liquid the system can be in a superposition of an extensive number of classical configurations. Gauge fields are essentially just the ingredient necessary for such a situation. The primary reason for this is because gauge fields are essentially *frame fields* (here frame can be interchanged with "basis"), and gauge transformations amount to an inhomogeneous change of basis. Without delving any further into this, we note that this perspective on gauge fields can be formalized in the language of connections on a principal bundle [2], a perspective first introduced to physics by Ref. [49], which we will very briefly touch upon in Section 3.4.5. For our purposes, the important insight is that gauge fields are generally described by 1-form fields,[6] and this is the core conceptual point we will exploit in developing a

[6]To be more specific: the physical degrees of freedom of a gauge field are described by the field strength 2-form F. This 2-form satisfies the closure property $\mathrm{d}F = 0$, where d is the exterior derivative of differential forms. As such it can be represented as $F = \mathrm{d}A$ where A is a 1-form field. More precisely, A is a Lie-algebra-valued 1-form, which defines a G-equivariant section of a principal-G bundle, where G is the gauge group, and A takes its values in the Lie algebra of G. If we have a "matter" field described by sections of a vector bundle with fiber V, then it has an associated principal-GL(V) bundle, where GL(V) is the general linear group of invertible linear maps $V \to V$, which is a Lie group. The space of frames for V is isomorphic to GL(V) since, given a reference basis, the columns of an invertible linear map are linearly independent and therefore form another basis. Thus the principal-GL(V) bundle P is the frame bundle of the vector bundle, and the $\mathfrak{gl}(V)$-valued 1-form A on P defines a frame field for the vector bundle,

theory of geometrically frustrated classical spin liquids in Chapter 3.

1.3 Spin Ice and the Coulomb Phase

Now on to the main topic of the bulk of this thesis, spin ice. We will discuss ice models in some detail in Chapter 2, but here we summarize the qualitative picture of the physics of spin ice. We will not review the history of the topic, which can be found in a number of sources [50–54]. Spin ice is one of the most important examples of a spin liquid, because (i) its classical versions have experimental realizations and are relatively well understood theoretically, and (ii) it offers a basic jumping off point to understanding quantum spin liquids through the minimal quantum spin ice models which realize emergent quantum electrodynamics.

Classical spin ice is realized in the compounds $Ho_2Ti_2O_7$ ($j = 8$) and $Dy_2Ti_2O_7$ ($j = 15/2$), in which the magnetic ions form a pyrochlore lattice of corner sharing tetrahedra with the ions sitting at the corners of the tetrahedra. Due to their large angular momentum, the dominant coefficients in Eq. (1.1) are the maximal $m = \pm j$ terms, and so these compounds are well approximated by classical Ising models [35]. The large j also implies that they have large magnetic moments, meaning that long-range dipole-dipole interactions are strong [35]. Despite this, the basic physics of spin ice is very well approximated in a nearest-neighbor approximation, for reasons which we will discuss in detail in Chapter 3.

In terms of the magnetic moments $\boldsymbol{m}_i$, the dominant nearest-neighbor interactions in spin ice are *ferromagnetic*. It is conventional to express the magnetic moment in terms of the pseudo-spin as $\boldsymbol{m}_i = g_{\parallel} S_i^z \hat{\boldsymbol{z}}_i$, where $\hat{\boldsymbol{z}}_i$ is the local easy axis direction. If for a single tetrahedron we take all four of the $\hat{\boldsymbol{z}}_i$ to point "in" (or all "out"), then the resulting nearest-neighbor Ising model has *anti-ferromagnetic* interactions between the S_i^z.

There are $2^4 = 16$ Ising spin configurations for a single tetrahedron, of which six are ground states, corresponding to states where two moments point "in" and two point "out" of the tetrahedron. This 2-in 2-out characterization of the ground state is called the "ice rule", due to a direct analogy with the proton disorder in water ice, from whence spin ice gets its name. In the left hand image in Fig. 1.1(a) we show two tetrahedra of the pyrochlore lattice, both of which satisfy the ice rule. If we think of the red arrows as fluxes in and out of the tetrahedron, then we can consider the ice rule to be a zero-divergence constraint on the ground states. This zero-divergence constraint implies that in the ground states of spin ice have spins aligned head-to-tail forming closed loops, as depicted in Fig. 1.1(b). All such ice rule obeying configurations are degenerate, and as was originally shown by Pauling, there are an exponentially large number of such configurations,

which defines a connection on the vector bundle and therefore a covariant derivative.

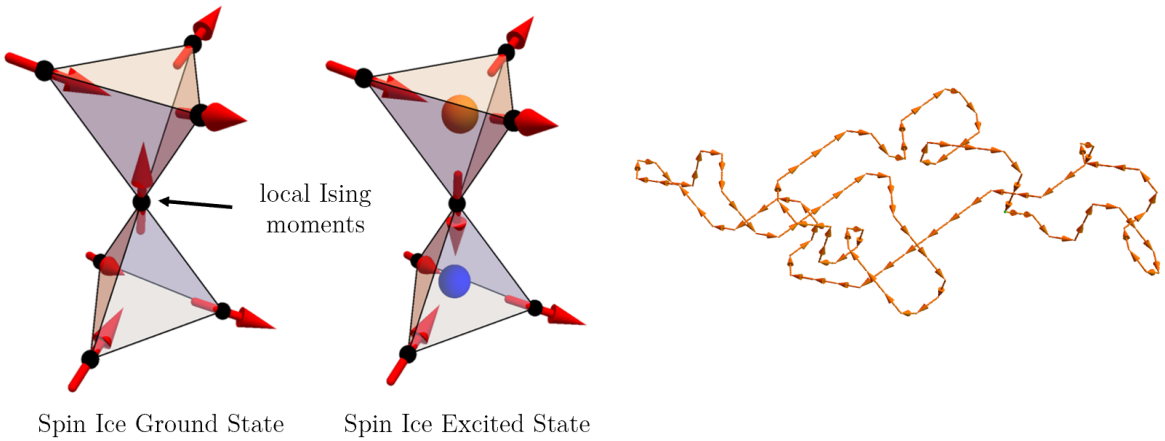


Figure 1.1: (a) two tetrahedra of the pyrochlore lattice, left: a spin ice ground state obeying the 2-in/2-out ice rule, right: a spin ice excited state created by flipping a single spin looks like a pair of charges. (b) In the Coulomb phase spins align into closed loops without ends.

resulting in a macroscopic ground state degeneracy and non-zero entropy at low temperature. This phase, in which the spins collectively organize themselves into closed strings with zero divergence, is called the Coulomb phase [55].

Excitations out of the Coulomb phase can be created by flipping a single spin, as illustrated in the second image in Fig. 1.1(a). This creates a non-zero divergence at each end of the flipped spin, which act like sources and sinks of the fluxes and thus as oppositely charged particles. These are precisely the classical realizations of the fractionalized spinons we mentioned earlier, where a single excitation of the spins created a pair of excitations. As we will discuss further later, these excitations are deconfined, meaning that they can be separated infinitely far from each other by flipping a longer and longer open string of spins at a *finite* energy cost.

1.3.1 Quantum Spin Ice and Emergent QED

What we described previously sounds qualitatively similar to classical electromagnetism: field lines forming closed loops, and charges acting as sources and sinks of field lines. It is quite fashionable to describe the classical spin ice model as having an "emergent gauge structure" [56], its excitations as "gauge charges" [57], and also sometimes to describe the Coulomb phase as having a form of "topological order" [58] or the excitations as "topological charges" [59]. These

phrases as used in the literature tend to be somewhat vague, not being used with the their precise meaning in the theory of gauge fields or the mathematics of topology. Much of Chapter 3 will be devoted to trying to provide a language and framework by which these statements can be made more precise in the classical model.

Before we continue, however, it is important to have a basic understanding of one model for which these statements can be made precise, and that is the ring exchange model of quantum spin ice. This model was studied in great detail in Ref. [25], where it was derived perturbatively from a model of spin ice with weak transverse exchange terms $S_i^+ S_j^-$ between nearest-neighbors which serves to flip two neighboring spins. At third order in perturbation theory, this term yields the ring exchange term, which flips six spins aligned head-to-tail around the boundary of the minimal plaquette of the pyrochlore lattice, which is a hexagon (cf. Fig. 3.19). The authors of Ref. [25] used this to map their quantum spin ice model to a lattice U(1) 1-form gauge theory, in which the S^z operators are treated as an electric field and the $S^\pm \sim e^{\pm iA}$ are treated as U(1) link variables [11]. The result is that the ring exchange term produces a photon, and the resulting theory is to a good approximation equivalent to quantum electrodynamics [60, 61]. Thus quantum spin ice is well described as a gauge theory. Much of our work on classical spin ice in Chapter 3 amounts to the development of a lattice field theory for the classical model in which the spins are taken to collectively describe a 1-form field. We will also attempt to make some small clarification on the meaning of topology in ice models, a topic discussed by [62, 63], because "topology" has become a major buzzword in condensed matter that often gets used in very imprecise ways.

Chapter 2

Methods: Theory of Classical Soft Spins

In this chapter we will describe in detail the classical soft spin approximation and its usefulness for studying spin liquid physics. The soft spin approximation will be important to our development of a discrete differential form theory for ice models in Chapter 3 and will be a working approximation throughout this thesis, allowing us to approximate for example neutron scattering cross sections. Various mathematical tools and notation will also be introduced in the present chapter which will be used throughout the remaining of the thesis, one of the most important of which is the analysis of the band structure of the interaction matrix and its relation to frustration. We will also give the general prescription for performing a long wavelength expansion which we will use multiple times, as well as discuss spin correlations and neutron scattering.

2.1 The Gaussian, Spherical, and Large-N Approximations

The most broadly studied classical spin Hamiltonians are those involving bilinear spin interactions,[1] and can generically be written in the form

$$H = \frac{1}{2} \sum_{ij} \sum_{\alpha\beta} S_i^\alpha \mathcal{J}_{ij}^{\alpha\beta} S_j^\beta \tag{2.1}$$

[1]The most prominent example of non-bilinear spin interactions is the four-spin terms appearing in the toric code model of Kitaev [64], and its variety of generalizations and extension, such as the eight- and twelve-spin terms appearing in the X-cube model [65]. These models are interesting because they admit exact solutions, and because they cleanly demonstrate interesting physics such as topological order [66] and fractons [67]. They are somewhat unrealistic from a condensed matter perspective, where spin interactions arise generally from electronic exchange between ions and bilinear interactions will generically dominate over higher multi-spin interactions. Such multi-spin interactions are not directly amenable to the analytical approximations which we will describe here.

where i, j index sites on a lattice and $\alpha, \beta = 1, \ldots N$ index N spin vector components. The case $N = 1$ corresponds to Ising models, $N = 2$ to XY models, and $N = 3$ to Heisenberg models. The matrix $\mathcal{J}_{ij}^{\alpha\beta}$ is the interaction matrix, which can in general be anisotropic (i.e. $\mathcal{J}_{ij}^{\alpha\beta} \not\propto \delta^{\alpha\beta}$) but we assume is consistent with the spatial symmetries of the lattice. We will say more about the anisotropic case in Chapter 6.

Our primary interest is to compute thermodynamic quantities, both bulk quantities such as specific heat and magnetization, and thermally-averaged spatially-resolved or momentum-resolved quantities, particularly spin-spin correlation functions. To do so, we need to compute the partition function, $Z = \mathrm{Tr}\exp\{-\beta H\}$ where β is the inverse temperature. For a classical spin model, the partition function is given by

$$Z = \int \mathcal{D}\boldsymbol{S} \exp\left(-\frac{1}{2}\sum_{ij}\sum_{\alpha\beta} S_i^\alpha \beta \mathcal{J}_{ij}^{\alpha\beta} S_j^\beta\right), \tag{2.2}$$

where the integration measure is uniform over the surface of an N-dimensional sphere. A convenient way to express this is to let the variables S_i^α vary from $-\infty$ to ∞ and enforce the spin-length constraint with a delta function, i.e.

$$\int \mathcal{D}\boldsymbol{S} \equiv \prod_i \int_{-\infty}^{\infty} \mathrm{d}S_i^1 \cdots \mathrm{d}S_i^N \, \delta(|\boldsymbol{S}_i|^2 - S^2). \tag{2.3}$$

In this form it is clear that the delta function constraint makes the partition function intractable to solve exactly. To do so we will make some approximations to "soften" the constraint. There are three nearly-equivalent methods to take to do this, which all lead to the same result. As we will see, they each provide a slightly different perspective.

2.1.1 The Self Consistent Gaussian Approximation

The first is the so-called self-consistent Gaussian approximation (SCGA), where the delta functions are replaced by Gaussians,

$$\delta(|\boldsymbol{S}_i|^2 - S^2) \to \exp\left[-(|\boldsymbol{S}_i|^2 - S^2)/2\lambda_i^{-1}\right], \tag{2.4}$$

where λ_i^{-1} is the variance of the Gaussian. In this approximation the spin length is not fixed, but is enforced on average via the self-consistency conditions[2]

$$\langle|\boldsymbol{S}_i|^2\rangle_{\{\lambda\},T} = S^2 \quad \forall\, i, \tag{2.5}$$

[2]The symbol $\forall$ is to be read "for all".

where the thermodynamic average depends on the values of the λ_i and temperature T. These equations yield a λ_i as a function of T. We refer to the spin variables in this approximation as "soft spins."

Substituting Eq. (2.4) into the partition function yields

$$Z_{\mathrm{SCGA}}[\{\lambda\}, T] = \prod_i e^{\lambda_i S^2} \int_{-\infty}^{\infty} \mathrm{d}S_i^1 \cdots \mathrm{d}S_i^N \; \exp\left\{-\frac{1}{2}\sum_{ij}\sum_{\alpha\beta} S_i^\alpha (\lambda_i \delta_{ij}\delta^{\alpha\beta} + \beta \mathcal{J}_{ij}^{\alpha\beta}) S_j^\beta\right\}. \tag{2.6}$$

In this form, we can see that the λ_i serve as Lagrange multiplier for the spin length constraint, which is now enforced energetically. The term leading the integral $\exp\left(S^2 \sum_i \lambda_i\right)$ is a constant which cancels in the computation of the normalized thermodynamic quantities we are interested in and can thus be dropped entirely. The resulting integral is a simple multi-dimensional Gaussian which can be solved exactly. We define the matrix $\mathcal{G}$ whose inverse is the covariance matrix of the above Gaussian,

$$[\mathcal{G}^{-1}]_{ij}^{\alpha\beta} \equiv \lambda_i \delta_{ij}\delta^{\alpha\beta} + \beta \mathcal{J}_{ij}^{\alpha\beta}. \tag{2.7}$$

Then by standard Gaussian integration, we directly obtain the spin-spin correlation functions,

$$\langle S_i^\alpha S_j^\beta \rangle = \mathcal{G}_{ij}^{\alpha\beta}. \tag{2.8}$$

The self-consistency conditions Eq. (2.5) are then a constraint on the diagonal elements of $\mathcal{G}$,

$$\sum_\alpha \mathcal{G}_{ii}^{\alpha\alpha} = S^2 \quad \forall\, i, \tag{2.9}$$

where the left hand side depends on the λ_i and temperature, $T = \beta^{-1}$.

In the case that the interaction matrix $\mathcal{J}_{ij}$ has certain symmetries, e.g. is translationally invariant, then $\mathcal{G}_{ij}$ will also have those symmetries and the solution of the equations Eq. (2.9) will reflect them. In particular, we are generally interested in spin systems on a crystalline lattice, and then it is usually the case that *every* lattice site is equivalent under the crystal symmetries, in which case the solutions to the self-consistency will have all λ_i equal, $\lambda_i \equiv \lambda$.[3] In that case we only need solve the equations for a single site i, or equivalently we can add them all up to obtain

$$\sum_{i,\alpha} \mathcal{G}_{ii}^{\alpha\alpha} \equiv \mathrm{Tr}\, \mathcal{G} = N_{\mathrm{spins}} S^2 \quad \text{(every site equivalent by symmetry)}. \tag{2.10}$$

[3] More generally, the λ_i are equal for all sites which are related by symmetry. For example, in a thin film which has layers which are translationally invariant in two dimensions, the λ_i are a equal within a layer but become a function of depth [68].

This is the form of the equation we will usually solve, since computing the trace of $\mathcal{G}$ is usually a straightforward task.

The SCGA obviously does not capture the exact spin length constraint, but it is generally a powerful tool for approximating correlations in paramagnetic phases, particularly at high temperatures. Note that at infinite temperature, from Eq. (2.7) with $\beta = 0$, $\mathcal{G}_{ij}^{\alpha\beta} = \lambda_i^{-1}\delta_{ij}\delta^{\alpha\beta}$, and the self-consistency equations Eq. (2.9) read simply

$$\frac{N}{\lambda_i} = S^2 \;\Rightarrow\; \lambda_i = \frac{N}{S^2}, \tag{2.11}$$

which is clearly independent of i, with N being the number of spin components. It is generally possible to solve the self-consistency equation at temperatures much larger than J without issue. For a finite N_{spins} it is in fact always possible to find a solution at any temperature, though in the thermodynamic limit there may be a critical temperature at which the equations cease to have solutions, predicting a phase transition. The SCGA is not capable of describing the low temperature symmetry broken phase below the critical temperature. We will discuss all of this in more detail in the rest of this chapter.

2.1.2 The Spherical Approximation

The fact that the SCGA yields a translationally invariant λ means that we can write the self-consistency conditions as a single equation,

$$\sum_i \langle |\boldsymbol{S}_i|^2 \rangle = \operatorname{Tr} \mathcal{G} = N_{\text{spins}} S^2, \tag{2.12}$$

which says that the *total* spin length of all spins is constant. This suggests considering a model with precisely this constraint, i.e. where we replace all of the spin-length constraining delta functions with a single global constraint,

$$\prod_i \delta(|\boldsymbol{S}_i|^2 - S^2) \to \delta\left(\sum_i |\boldsymbol{S}_i|^2 - N_{\text{spins}} S^2\right). \tag{2.13}$$

This is the precise logic that led Mark Kac in 1947 to first consider this approximation, which he later solved with the help of Theodore Berlin, and which they dubbed the "spherical model" [69, 70]. Conceptually, one can imagine that we have replaced the many "small" hard spins with fixed length S by a single "big" hard spin with fixed length $N_{\text{spins}}^{1/2} S$.

For our purposes we need not go through the exact solution, except to state that the resulting correlation functions are exactly those of the SCGA in the paramagnetic phase. The primary difference between the two methods is that the spherical model has a well-defined phase transition and symmetry broken phase, because the single "big spin" still has a hard spin length constraint.

We will not utilize the spherical approximation in this thesis, primarily because our interest is in paramagnetic spin liquid phases where the SCGA is sufficient. In Chapter 6, we will discuss symmetry broken ground states, but there we already know the exact ground state with the "small" spin constraints and the spherical approximation is not needed. However, it is quite useful to keep the spherical "one big spin" approximation in mind as a conceptual tool.

2.1.3 The Large-N Approximation

Another approximation which we can make is to take the number of spin components N to infinity [71]. Essentially by the central limit theorem, we will see that in this limit the correlations reduce to those of the SCGA [72]. Thus in a sense, the large-N approximation can be taken as a justification for using the SCGA, but more importantly it serves as a basis to develop a controlled $1/N$ perturbative expansion away from the $N = \infty$ limit, and thus a method to compute corrections to the SCGA.

To compute the large-N partition function, we start from the N-component partition function, and replace the spin length S with S_N, where $S_N^2 \equiv S^2 N/N_{\mathrm{t}}$, where N_{t} is the "target" number of components (i.e. in the end we wish to evaluate the correlation functions of Ising or Heisenberg spins, so $N_{\mathrm{t}} = 1$ or 3, respectively).[4],

$$Z_{\text{Large-}N} = \int_{-\infty}^{\infty} \prod_i \mathrm{d}S_i^1 \cdots \mathrm{d}S_i^N \, \delta(|\boldsymbol{S}_i|^2 - S_N^2) \, \exp\left(-\frac{1}{2}\sum_{ij} \beta \mathcal{J}_{ij} \boldsymbol{S}_i \cdot \boldsymbol{S}_j\right). \tag{2.14}$$

Note that we did not write the Hamiltonian with anisotropic interactions. This is because the specific anisotropic terms which can appear in the Hamiltonian depend on N and are generally not well-defined for $N > 3$. In particular, the allowed anisotropic couplings for $N = 3$ Heisenberg spins will be described in detail in Chapter 6 and depend on how the spins transform under the three dimensional spatial symmetries of rotations and reflections—in general they should transform as angular momentum pseudo-vectors, and there is no unique way to define such transformations for vectors with $N > 3$ components. We will return to this point at the end of this section.

[4]This convention of defining the N-component spin length comes from [73]. Another common convention is to simply replace S^2 with N.

The large-N derivation relies on the representation of the delta functions via the formula

$$\delta(x) \propto \int_{\lambda-i\infty}^{\lambda+i\infty} \mathrm{d}\mu \, e^{-\mu x/2}, \tag{2.15}$$

where λ is a real number. The proportionality is irrelevant since it will appear as an overall constant in the partition function, and the factor of two is a convenient scaling. Substituting this identity, we obtain

$$Z_{\text{Large-}N} = \int_{\lambda-i\infty}^{\lambda+i\infty} \prod_i \mathrm{d}\mu_i \, e^{\mu_i S_N^2/2} \int_{-\infty}^{\infty} \prod_i \mathrm{d}S_i^1 \cdots \mathrm{d}S_i^N \, \exp\left(-\frac{1}{2}\sum_{ij}(\mu_i\delta_{ij} + \beta\mathcal{J}_{ij})\boldsymbol{S}_i \cdot \boldsymbol{S}_j\right). \tag{2.16}$$

We now note that the integrals over spin components fully decouple, so that we can write the partition function as

$$Z_{\text{Large-}N} = \int_{\lambda-i\infty}^{\lambda+i\infty} \prod_i \mathrm{d}\mu_i \, e^{\mu_i S_N^2/2} \prod_{\alpha=1}^{N} \left[\int_{-\infty}^{\infty} \prod_i \mathrm{d}S_i^\alpha \, \exp\left(-\frac{1}{2}\sum_{ij}(\mu_i\delta_{ij} + \beta\mathcal{J}_{ij})S_i^\alpha S_j^\alpha\right)\right]. \tag{2.17}$$

Each integral for fixed α is identical since the Hamiltonian is isotropic. As before define the covariance matrix

$$[\mathcal{G}(\mu)^{-1}]_{ij} = \mu_i\delta_{ij} + \beta\mathcal{J}_{ij}, \tag{2.18}$$

then we have (up to overall constants which we ignore)

$$Z_{\text{Large-}N} = \int_{\lambda-i\infty}^{\lambda+i\infty} \prod_i \mathrm{d}\mu_i \, e^{\mu_i S_N^2/2} \, \det[\mathcal{G}(\mu)]^{N/2}, \tag{2.19}$$

$$= \int_{\lambda-i\infty}^{\lambda+i\infty} \prod_i \mathrm{d}\mu_i \, \exp\left(-\frac{N}{2}\left[-\sum_i \mu_i \frac{S^2}{N_{\mathrm{t}}} + \mathrm{Tr}\log[\mathcal{G}(\mu)^{-1}]\right]\right). \tag{2.20}$$

In the final form, it is clear that this integral can be solved via the method of steepest descent in the limit $N \to \infty$. In order to apply the method of steepest descent, the exponent must be expanded about the saddle point where its derivative with respect to the μ_i vanishes. The value of λ is then determined by the requirement that the integration contour run over the saddle point [74, 75]. The

saddle point condition reads

$$0 = \frac{\partial}{\partial \mu_k}\left[-\sum_i \mu_i \frac{S^2}{N_\mathrm{t}} + \mathrm{Tr}\log[\mathcal{G}^{-1}]\right] \tag{2.21}$$

$$= -\frac{S^2}{N_\mathrm{t}} + \mathrm{Tr}\left[\mathcal{G}\frac{\partial \mathcal{G}^{-1}}{\partial \mu_k}\right] \tag{2.22}$$

$$= -\frac{S^2}{N_\mathrm{t}} + \sum_i \sum_j \mathcal{G}_{ij}\frac{\partial}{\partial \mu_k}(\mu_i \delta_{ji} + \beta \mathcal{J}_{ji}) \tag{2.23}$$

$$= -\frac{S^2}{N_\mathrm{t}} + \sum_i \sum_j \mathcal{G}_{ij}\delta_{ij}\delta_{ik} \tag{2.24}$$

$$= -\frac{S^2}{N_\mathrm{t}} + \mathcal{G}_{kk} \tag{2.25}$$

for every k. In other words,

$$N_\mathrm{t}\mathcal{G}(\mu_\mathrm{sp})_{ii} = S^2 \quad \forall\, i \qquad \text{(saddle point condition)}. \tag{2.26}$$

Noting that $\mathcal{G}(\mu)$ is the spin-spin correlation function for a single spin component with fixed $\{\mu\}$, i.e.

$$\langle |S_i^\alpha|^2 \rangle_{\{\mu\}} = \mathcal{G}(\mu)_{ii}, \tag{2.27}$$

we can identify $N_\mathrm{t}\mathcal{G}(\mu)_{ii} \equiv \langle |\boldsymbol{S}_i|^2 \rangle_{\{\mu\}}$ since we have assumed an isotropic model, in which case the saddle point condition can be interpreted as an *average* spin length condition

$$\langle |\boldsymbol{S}_i|^2 \rangle_{\{\mu_\mathrm{sp}\}} = S^2 \quad \forall\, i \qquad \text{(saddle point condition)}, \tag{2.28}$$

The solution to this equation generically occurs at a real and positive value of μ_i which is translationally invariant, and is generally denoted $(\mu_\mathrm{sp})_i \equiv \lambda$. It determines the offset λ in Eq. (2.15) for the integration contour in the complex μ plane, which must be chosen to run over the saddle point [74].

The saddle point condition Eq. (2.28) is precisely the self consistency condition in the Gaussian approximation, Eq. (2.9). The large-N approximation can thus be taken as a justification for the SCGA, though as we stated previously the large-N limit is not a per se natural approximation for anisotropic spin Hamiltonians (because the N integrals in Eq. (2.17) only decouple if the Hamiltonian is isotropic), while the Gaussian approximation is perfectly valid for anisotropic Hamiltonians. The upside to the large-N formulation is that it allows for a controlled perturbative expansion away from $N = \infty$ in powers of $1/N$. A derivation of the $1/N$ series expansion is given in Appendix C. In the remainder of this thesis we will work almost exclusively with the SCGA for our analytical calculations, though we will be interested in the $1/N$ correction to the correlation functions in Sections 4.6 and 5.4.3.

2.2 Soft and Hard Spin Configuration Spaces

In the classical soft spin approximations discussed above, the spin variables S_i^α are treated as real numbers which vary continuously. There are $\mathcal{N} = N_{\text{spins}} \times N$ real variables S_i^α, and therefore the space of soft spin configurations is an $\mathcal{N}$-dimensional real vector space, which we denote Ω_{spin}. Each vector in this place represents a soft spin configuration of the entire system. This is one of the primary advantages for working with soft spins: we can apply standard methods of linear algebra to describe and study the physics of classical spins models.

We will use bra-ket notation to work with vectors in this space because it will allow us to write convenient and transparent coordinate-free formulas. We denote a spin configuration by $|S\rangle \in \Omega_{\text{spin}}$. There is a canonical basis for this vector space indexed by i and α, $\{|i,\alpha\rangle\}$, and a canonical Hermitian inner product (the standard L^2 inner product), such that the spin component variables are given by

$$S_i^\alpha = \langle i,\alpha|S\rangle\,, \tag{2.29}$$

and a spin configuration can be written as

$$|S\rangle = \sum_{i,\alpha} |i,\alpha\rangle\,\langle i,\alpha|S\rangle \equiv \sum_{i,\alpha} S_i^\alpha\,|i,\alpha\rangle\,. \tag{2.30}$$

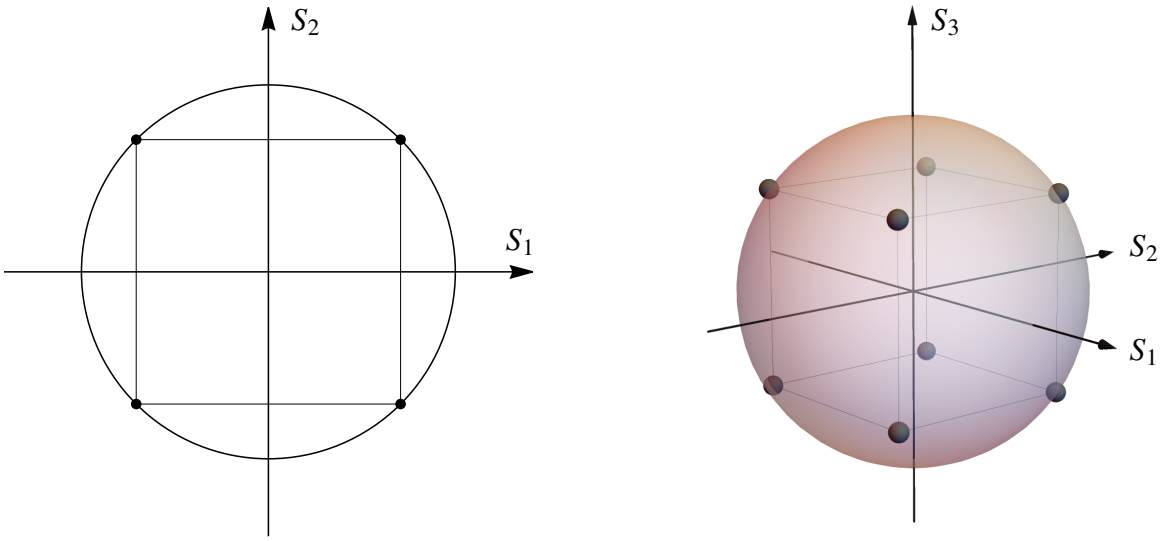
The N_{spins} hard spin constraints $|\boldsymbol{S}_i|^2 = S^2$ are quadratic polynomials defining a constraint manifold in Ω_{spin} which is isomorphic to the Cartesian product of N_{spins} copies of the $N-1$-sphere.[5] We denote this manifold by $\mathcal{M}_{\text{hard}}$. The spherical approximation replaces this hard spin constraint surface with the sphere on which

$$\langle S|S\rangle = \sum_{i,\alpha} |S_i^\alpha|^2 = N_{\text{spins}} S^2. \tag{2.31}$$

Note that $\mathcal{M}_{\text{hard}}$ is always embedded in this sphere. The SCGA removes the constraint surface entirely, allowing the length of $|S\rangle$ to fluctuate with Gaussian distribution whose variance is energetically fixed to $N_{\text{spins}}S^2$ by the Lagrange multiplier λ.

This is depicted in Fig. 2.1 for $N_{\text{spins}} = 2, 3$ with $N = 1$. The hard spin configuration space $\mathcal{M}_{\text{hard}}$ satisfying $|S_i|^2 = S^2$ or $S_i = \pm S$ is given by the corners of N_{spins}-dimensional cube of edge length $2S$ centered at the origin. The corners sit at a distance $\sqrt{N_{\text{spins}}}\,S$ from the origin. In the spherical approximation the hard spin constraint manifold is replaced with the sphere of radius $\sqrt{N_{\text{spins}}}\,S$ in which the cube is inscribed. In the Gaussian approximation the soft spins can explore the entire space but with the total spin length obeying a Gaussian distribution with fixed variance, i.e. the root-mean-square total spin length is fixed to $\sqrt{N_{\text{spins}}}\,S$.

[5] For example, for $N = 1$ Ising spins, the constraint surface is a product of 0-spheres, where a 0-sphere consists of two points, equivalent to a hypercube. For $N = 3$ Heisenberg spins, the constraint surface is a product of 2-spheres.



(a) (b)

Figure 2.1: The configuration space Ω_{spin} for (a) $N_{\text{spins}} = 2$ and (b) $N_{\text{spins}} = 3$ soft Ising ($N = 1$) spins. The hard spin constraint defines a subset of points in this space $\mathcal{M}_{\text{hard}}$ given by the corners of a cube (or a Cartesian product of N-spheres for arbitrary N). The spherical approximation replaces the hard spin manifold with the sphere in which this cube is inscribed.

2.2.1 The Interaction Matrix and SCGA Effective Hamiltonian

A fully general bilinear spin Hamiltonian has the form

$$H = \frac{1}{2} \sum_{ij} \sum_{\alpha\beta} S_i^\alpha \mathcal{J}_{ij}^{\alpha\beta} S_j^\beta, \tag{2.32}$$

where $\mathcal{J}$ is the interaction matrix. For Ising ($N = 1$) spins there are no spin component indices. The Hamiltonian can be considered to be a quadratic form on Ω_{spin}, meaning we can rewrite Eq. (2.32) in a coordinate free form as

$$H = \frac{1}{2} \langle S | \mathcal{J} | S \rangle , \tag{2.33}$$

which can be expanded by inserting two resolutions of the identity $\mathbb{1} = \sum_{i,\alpha} |i, \alpha\rangle \langle i, \alpha|$,

$$\frac{1}{2} \sum_{ij} \sum_{\alpha\beta} \langle S | i, \alpha \rangle \langle i, \alpha | \mathcal{J} | j, \beta \rangle \langle j, \beta | S \rangle , \tag{2.34}$$

We identify the interaction matrix components,

$$\mathcal{J}_{ij}^{\alpha\beta} \equiv \langle i, \alpha | \mathcal{J} | j, \beta \rangle . \tag{2.35}$$

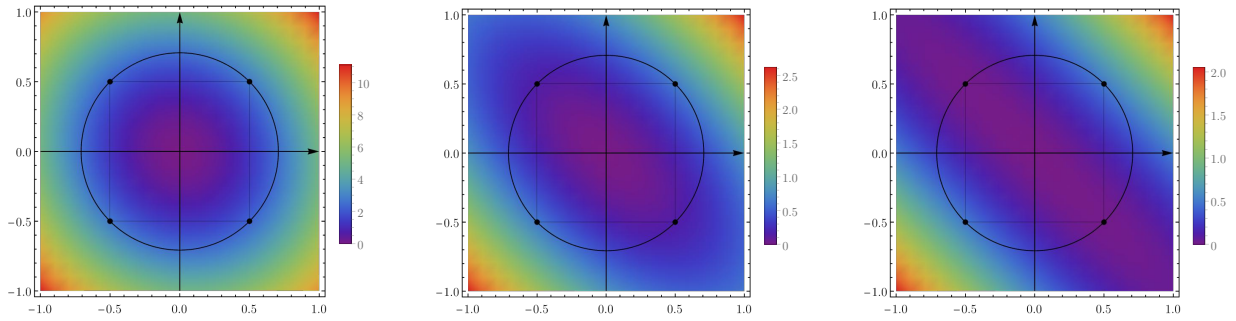
From the SCGA partition function Eq. (2.6) we can read off an *effective* Hamiltonian for the soft spins,

$$H_{\text{SCGA}} = \frac{1}{2} \sum_{ij} \sum_{\alpha\beta} S_i^\alpha [\lambda T \delta_{ij} \delta^{\alpha\beta} + \mathcal{J}_{ij}^{\alpha\beta}] S_j^\beta . \tag{2.36}$$

This allows us to define an effective interaction matrix for the soft spins,

$$\mathcal{J}_{\text{SCGA}} \equiv \lambda T \mathbb{1} + \mathcal{J}, \tag{2.37}$$

which differs from the original interaction only by the addition of diagonal terms proportional to λ. For hard spins with fixed spin length, adding diagonal elements to the interaction matrix produces terms in the Hamiltonian $\mathcal{J}_{ii} |\boldsymbol{S}_i|^2/2$, which is a constant shift to the energy that has no influence on the physics. In the SCGA on the other hand, the diagonal term λT in the effective interaction matrix serves to control the fluctuating spin length. The value of λ is fixed at each temperature by the self-consistency condition Eq. (2.12) that the thermally averaged total spin length squared is constant.



(a) $T/JS^2 = 10$ (b) $T/JS^2 = 1$ (c) $T/JS^2 = 0.1$

Figure 2.2: The SCGA effective potential for two Ising spins interacting antiferromagnetically as a function of temperature. At high temperature the potential is isotropic and all states are equally likely, at low temperature the potential prefers the lower-energy states.

The SCGA Hamiltonian can be considered as a potential energy on the space of soft spin configurations.[6] Consider the simplest example of two $N = 1$ spins interacting via the Hamiltonian

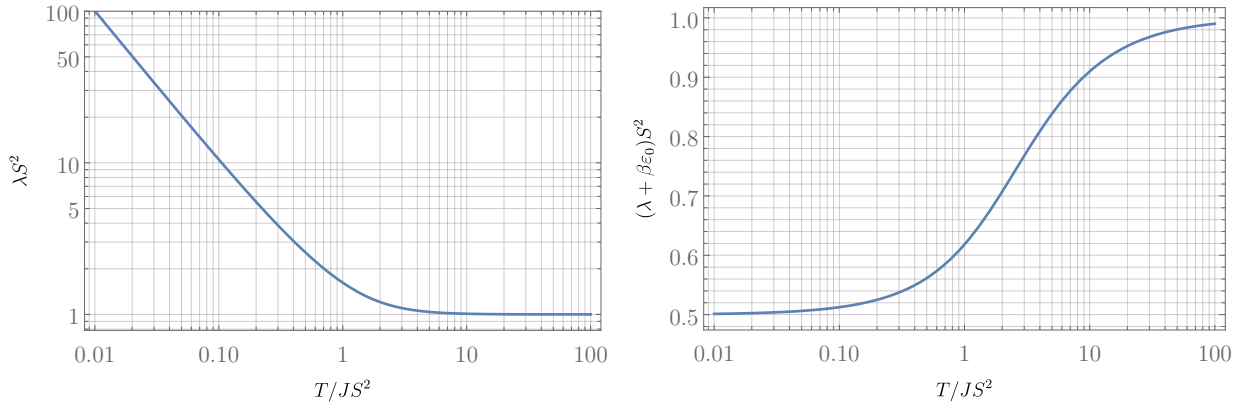
$$H_{\text{SCGA}} = \frac{1}{2}\lambda T(S_1^2 + S_2)^2 + JS_1S_2 = \frac{1}{2}\begin{pmatrix} S_1 & S_2 \end{pmatrix}\begin{pmatrix} \lambda T & J \\ J & \lambda T \end{pmatrix}\begin{pmatrix} S_1 \\ S_2 \end{pmatrix}. \tag{2.38}$$

The eigenvalues of the interaction matrix are $\pm J$, and the corresponding eigenvectors are given by $|\varphi_\pm\rangle = [|1\rangle \pm |2\rangle]/\sqrt{2}$. The covariance matrix is $\mathcal{G} = (\beta\mathcal{J}_{\text{SCGA}})^{-1}$, and the self-consistency condition can be written as

$$\operatorname{Tr}\mathcal{G} = \frac{1}{\lambda + \beta J} + \frac{1}{\lambda - \beta J} = 2S^2. \tag{2.39}$$

The effective SCGA potential energy landscape for this simple system is shown in Fig. 2.2 at three different temperatures. At high temperatures $T \gg JS^2$ the potential is isotropic. As temperature decreases the value of λ evolves. The potential increases quadratically along the direction of the higher-energy eigenvector of $\mathcal{J}$, suppressing the ferromagnetic states with $S_1 = S_2$, while the two anti-ferromagnetic states with $S_1 = -S_2$ sit in the minimum of the potential. We will discuss the physical meaning of the temperature evolution of λ in the next section.

[6] We will address spin kinetics at the end of this chapter when we discuss neutron scattering, and further in Chapter 6. For now, we are only interested in the thermal physics of these systems.



(a) Minimum eigenvalue of $\mathcal{J}$ negative.

(b) Minimum eigenvalue of $\mathcal{J}$ shifted to zero.

Figure 2.3: (a) Temperature dependence of λ for the two-spin antiferromagnet without shifting the diagonal of the interaction matrix, which diverges as $T \to 0$ as ϵ_0/T. (b) Temperature dependence of $\lambda_0 \equiv \lambda + \beta\epsilon_0$, equivalent to the value of λ if we shift the minimum eigenvalue of the interaction matrix to zero, which is finite at all temperature.

2.2.2 The Physical Significance of λ and the Spectral Decomposition of $\mathcal{G}$

The solution to Eq. (2.39) is

$$\lambda S^2 = \frac{1}{2} + \sqrt{\frac{1}{4} + (\beta J S^2)^2}. \tag{2.40}$$

At low temperature $\beta J S^2 \gg 1$ this diverges as T^{-1}, as shown in Fig. 2.3a. This divergence had to be the case to prevent the denominator of the second term in Eq. (2.39) from ever going to zero, at which point there would be no solution of the self-consistency condition. This occurs due to the fact that $\mathcal{J}$ has negative eigenvalues. Let $\epsilon_0 < 0$ be the lowest eigenvalue of $\mathcal{J}$, then at low temperatures generically λ will diverge as $\beta|\epsilon_0|$ to keep the denominator positive. This is not a problem of course—the eigenvalues of the interaction matrix correspond to the energies of its eigenstates and we can always shift the zero point energy. Consider adding and subtracting $\epsilon_0 \mathbb{1}$ to the original Hamiltonian, which achieves precisely this shift,

$$H = \frac{1}{2} \langle S | \left(\mathcal{J} - \epsilon_0 \mathbb{1} + \epsilon_0 \mathbb{1} \right) | S \rangle \tag{2.41}$$

$$= \underbrace{\frac{1}{2} \epsilon_0 \langle S | S \rangle}_{E_0 \equiv \frac{1}{2}\epsilon_0 N_{\text{spins}} S^2} + \frac{1}{2} \langle S | \left(\mathcal{J} - \epsilon_0 \mathbb{1} \right) | S \rangle . \tag{2.42}$$

The first term is the energy of the spherical model ground state,[7] while in the second term $\mathcal{J} - \epsilon_0 \mathbb{1}$ is a positive semi-definite matrix with no negative eigenvalues, and measures the excitation energy above the ground state. Consider now doing the same for the SCGA effective interaction matrix, Eq. (2.37),

$$\mathcal{J}^{-1}_{\text{SCGA}} = (\lambda T \mathbb{1} + \epsilon_0 \mathbb{1}) + \mathcal{J} - \epsilon_0 \mathbb{1}) = (\lambda + \beta\epsilon_0) T \mathbb{1} + (\mathcal{J} - \epsilon_0 \mathbb{1}). \tag{2.43}$$

The second term is again positive definite, while in the first term the Lagrange multiplier λ is shifted by $\beta\epsilon_0$. The quantity

$$\lambda_0 \equiv \lambda + \beta\epsilon_0 \tag{2.44}$$

should be finite, since we removed the divergence caused by the negative eigenvalues. Indeed, setting $\epsilon_0 = -J$ for the 2-spin antiferromagnet, we observe in Fig. 2.3 that λ_0 is always finite.

The quantity λ_0 carries the important physics information, which can be seen by performing a spectral decomposition of the correlation matrix $\mathcal{G}$. Denote the eigenvalues of $\mathcal{J}$ by ϵ_n with corresponding eigenvectors $|\varphi_n\rangle$. Then the spectral decomposition of $\mathcal{G}$ is given by

$$\mathcal{G} = \sum_{k=1}^{N_{\text{spins}} N} \frac{1}{\lambda_0 + \beta(\epsilon_k - \epsilon_0)} |\varphi_k\rangle \langle\varphi_k| \,. \tag{2.45}$$

The eigenvectors of $\mathcal{J}$ define the normal modes of the system, which we denote $S_n \equiv \langle\varphi_n|S\rangle$. The eigenvalues of $\mathcal{G}$ are the variances of the normal modes,

$$\left\langle |S_k|^2 \right\rangle = \frac{1}{\lambda_0 + \beta(\epsilon_k - \epsilon_0)}, \tag{2.46}$$

where the brackets indicate thermal average. Notice that the left hand side of Eq. (2.46) is non-negative, so the eigenvalues of $\mathcal{G}$ are also necessarily non-negative. We will refer to these eigenvalues as the generalized susceptibilities of the normal modes S_k. They can also be interpreted as a generalized occupation number for these modes. The SCGA self-consistency condition reads

$$\operatorname{Tr} \mathcal{G} = \sum_k \left\langle |S_k|^2 \right\rangle = \sum_k \frac{1}{\lambda_0 + \beta(\epsilon_n - \epsilon_0)} = N_{\text{spins}} S^2, \tag{2.47}$$

which we can read to say that the total generalized occupation of all modes is a constant, independent of temperature. At infinite temperature $\beta = 0$, all modes are occupied with equal probability. As temperature decreases, the occupations of higher energy modes decreases, and the occupations of

[7] That is, it is the energy of the lowest-energy state $|S\rangle$ on the spherical model constraint surface. We will see an example shortly of a frustrated system where this is not the hard spin ground state and thus the true ground state energy of the system is larger than this naive ground state energy estimate.

lower energy modes increases. This implies that λ_0 is a monotonically decreasing function of temperature, as in Fig. 2.3. In particular, the lowest energy modes with $\epsilon_n = \epsilon_0$ have a generalized susceptibility/occupation of exactly

$$\langle |S_0|^2 \rangle = \frac{1}{\lambda_0}. \tag{2.48}$$

Therefore, the physical meaning of λ_0^{-1} is the generalized susceptibility of the ground state modes.[8]

2.2.3 The Low Temperature Behavior of λ_0 in the Thermodynamic Limit

For finite systems such as the 2-spin antiferromagnet of the previous section, λ_0 will remain nonzero at all temperatures. In a thermodynamically large system ($N_{\text{spins}} \to \infty$) however, it is possible that λ vanishes at a critical temperature T_c. To see that this occurs, let us briefly consider some toy models. In the thermodynamic limit, the energy levels of the system generally form a continuum, and we can write the self-consistency condition in terms of the density of states $g(E)$ as

$$\frac{\operatorname{Tr} \mathcal{G}}{N_{\text{spins}} N} = \frac{\int_0^\infty dE\, g(E)\, \frac{1}{\lambda_0 + \beta E}}{\int_0^\infty dE\, g(E)} = S^2, \tag{2.49}$$

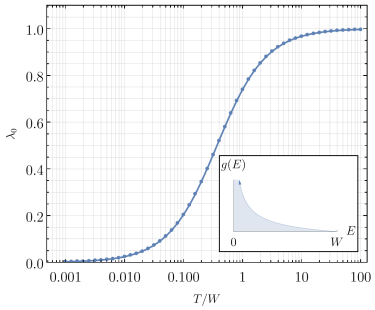
where we have implicitly set the ground state energy to zero. We will from here on set $S = 1$. The low temperature behavior of λ_0 is controlled by the low-energy portion of $g(E)$. We therefore consider simple toy densities given by

$$g(E) \propto \begin{cases} E^p & 0 \le E \le W \\ 0 & \text{otherwise} \end{cases} \tag{2.50}$$

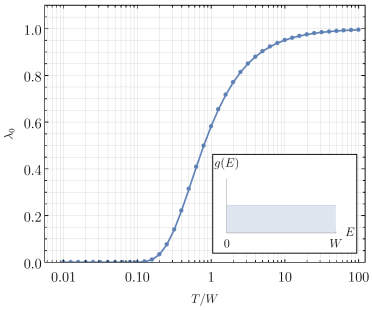
The proportionality constant in $g(E)$ cancels in the ratio in Eq. (2.49). If $p > 0$ then $g(0) = 0$, the ground states are subextensive and the system should order at sufficiently low temperature. In the case $p < 0$, the density of states is extensive and large at zero energy, implying an extensively degenerate ground state or an extensive number of states just above the ground state, where the system may avoid ordering.

The integrals can be solved analytically and the resulting equation solved numerically, with the results for $\lambda_0(T)$ given in Fig. 2.4, for $p = -1/2$, 0, $1/2$, 1, and 2. When $p \le 0$ the system has

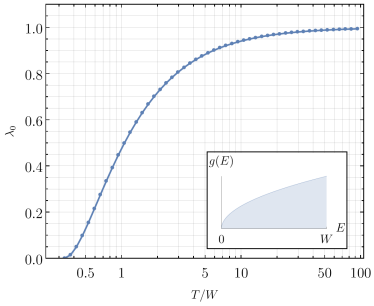
[8] More generally, one could have chosen ϵ_0 not to be the lowest eigenvalue of $\mathcal{J}$ but rather the eigenvalue of some excited mode of interest. In Ref. [76], the convention was chosen to pick ϵ_0 to correspond to the bulk magnetization mode, which we will discuss more later in this chapter. In that case, $\lambda_0 = \lambda + \beta\epsilon_0$ can diverge at low temperature if this mode is not the true ground state (e.g. if the ground state breaks translation symmetry), and λ_0^{-1} goes to zero at low temperature, meaning that the fluctuations of the bulk magnetization are suppressed and the bulk susceptibility goes to zero. We will only ever take ϵ_0 to be the lowest eigenvalue of $\mathcal{J}$.



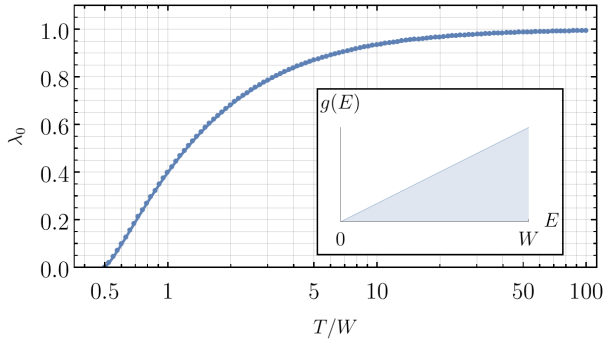
(a) $p = -1/2$
$T_c/W = 0$
$\lambda_0 \sim T$



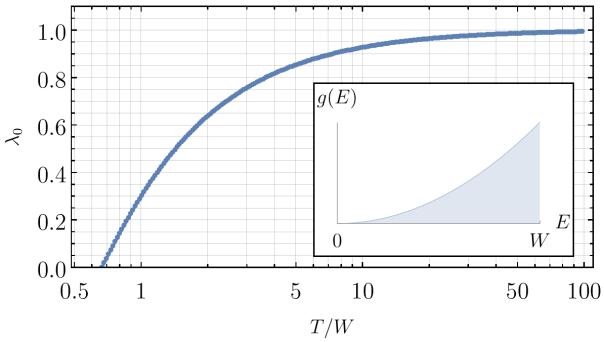
(b) $p = 0$
$T_c/W = 0$
$\lambda_0 = (\exp(-W/T)-1)^{-1}(W/T)$



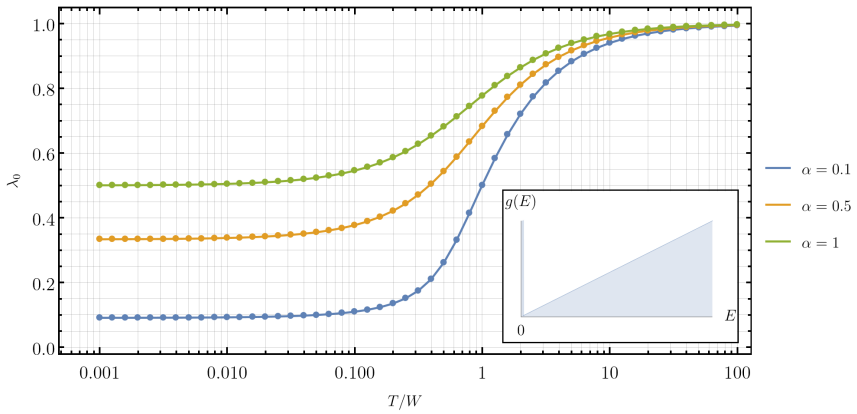
(c) $p = 1/2$
$T_c/W = 1/3$
$\lambda_0 \sim (T - T_c)^2$



(d) $p = 1, T_c/W = 1/2, \lambda_0(T - T_c)$ non-polynomial



(e) $p = 2, T_c/W = 2/3, \lambda_0 \sim (T - T_c)$



(f) Flat bands at zero energy, no phase transition, λ_0 saturates to a constant.

Figure 2.4: The temperature evolution of λ_0 for toy models with extensive (a,b,f) and subextensive (c,d,e) zero-energy density of states.

an extensive number of low energy modes and λ_0 monotonically decreases to zero as temperature is lowered. For $p > 0$, the system has a subextensive number of low energy modes and λ_0 reaches zero at a finite temperature critical temperature T_c. The exact critical temperatures can be computed by expanding Eq. (2.49) about $\lambda_0 = 0$, as well as the functional dependence of λ_0 on $T - T_c$ which are summarized in the captions of Fig. 2.4. As the excited states depopulate, the conserved total spin length enforced by the self consistency condition causes the occupation of the lower energy modes to increase, where the population of a mode with energy E is $1/(\lambda_0 + \beta E)$. At the critical temperature, the fluctuations of the ground state mode with $E = 0$ diverges as $\langle |S_0|^2 \rangle = \lambda_0^{-1}$ or, put another way, this mode condenses and obtains a macroscopic occupation number. When the density of states near zero energy is large, the critical temperature is suppressed, until at $p = 0$ it reaches zero. With a constant density of states we obtain an exact expression for λ_0, and find that it approach zero exponentially fast at low temperature. For $p = -1/2$ λ_0 goes to zero but only linearly with T.

Frustration tends to increase the low energy density of states of a system, and in the most extreme cases lead to a macroscopic number of ground states. We can model this extreme case with a density of states given by

$$g(E) \propto E^p + \alpha \frac{W^{p+1}}{p+1} \delta(E), \tag{2.51}$$

with $p > 0$. The density of states decreases at low energy but there are a macroscopic number of states at $E = 0$ corresponding to the presence of a perfectly flat band at zero energy, which we will see is the case for ice models. The dimensionless constant α is the ratio of the number of states in the flat band to the number of states with $E > 0$. The resulting temperature evolution of λ_0 is shown in Fig. 2.4(f) for $p = 1$. The important qualitative result is that λ_0 saturates to a constant and does not go to zero. Thus the SCGA predicts no phase transition, even at zero temperature. The system remains in the paramagnetic phase at all temperatures, with the constraint at zero temperature that the excited modes are completely unoccupied and the configuration space is reduced to a constrained but extensive subspace. Since the total occupation of all modes is conserved by the self consistency condition, at zero temperature each flat band mode has an occupation equal to the total number of states divided by the number of states in the flat band (assuming that the occupations are equally distributed among states in the flat band). The occupation of the zero energy modes is precisely λ_0^{-1}, and so λ_0 saturates to $\alpha/(1 + \alpha)$, which is the fraction of the total number of states in the flat band.

2.3 Spatial Symmetries and Representation Theory

So far we have described spin models in a general way, without exploiting any symmetries in our analysis. In this section, we will discuss how symmetries can be systematically folded into our analysis of soft and hard spin models. The primary symmetries we will be interested in are spatial ones.

The symmetries (isometries) of d-dimensional Euclidean space are described by the Euclidean group $E(d)$, which contains rotations, translations, and inversions. The Euclidean group can be expressed as a semi-direct product $O(d) \ltimes T(d)$, where $O(3)$ is the orthogonal group of rotations and inversion[9] in d dimensions and $T(d)$ is the group of translations in d dimensions. The product is semidirect because rotations and translations do not commute.[10]

Crystalline lattices break the smooth Euclidean symmetries down to a discrete subgroup called a space group. There are 17 space groups in two dimensions and 230 space groups in three dimensions. In every space group the translations are broken to a discrete subgroup isomorphic to $\mathbb{Z}^d$, splitting space into discrete copies of a single unit cell which tiles space. The quotient of the space group by the translations describe the symmetries of a single unit cell, which is a subgroup of $O(d)$ called the point group.

We will make extensive use of these symmetries in the analysis of spin models. In particular, the soft spin configuration space Ω_{spin} is a vector space equipped with an action of the crystalline space group, meaning that we can make use of representation theory. A representation ρ of a group G is a linear action of G on a vector space V, meaning that $\rho : G \to GL(V)$ is a group homomorphism from G to the general linear group of V (the group of all invertible linear maps $V \to V$). Under the action of the full space group, Ω_{spin} will decompose into a direct sum of subspaces which are the orbits of the representation, meaning that under the action of any space group symmetry these subspaces are mapped to themselves. The induced representations on these subspaces are called irreducible representations. In the next section we will introduce the Fourier basis, which is the irreducible basis for translations. After that, we will discuss how the irreducible representations of the point group control the long-wavelength physics.

2.3.1 The Fourier Basis: Irreducible Representations of Translations

We assume that the spins are located on the sites of a d-dimensional crystalline lattice, which can be described by a spatially repeating unit cell and a sublattice basis of crystal sites within a

[9] Alternatively, rotations and reflections.

[10] More precisely, a semidirect product of groups $A \ltimes B$ is a group where B acts on itself in the usual way and A acts on B as an inner automorphism. We will not be concerned with these details.

cell. Let $\boldsymbol{a}_i$, $i = 1, \ldots, d$ be the primitive lattice translation vectors, and let $\boldsymbol{c}_\mu$ be the sublattice positions within a cell. A lattice site is then indexed by three integers n_i and a sublattice index μ, and is located at position $\boldsymbol{R}$ given by

$$\boldsymbol{R} = \sum_{i=1}^{d} n_i \boldsymbol{a}_i + \boldsymbol{c}_\mu. \tag{2.52}$$

We will always consider a finite size lattice where $n_i \in \{0, \ldots, L-1\}$ so that there are exactly L^d unit cells, with periodic boundaries (i.e. $n_i \in \mathbb{Z}_L$), containing $L^d N_{\text{sl}}$ lattice sites, where N_{sl} is the number of sublattices. It is helpful to consider the system with periodic boundaries to be a super cell with $L^d N_{\text{sl}}$ sublattices, which infinitely tiles space in a repeating pattern. Alternatively, we can consider our system to be on a torus.

In Section 2.2 we indexed our basis vectors by site index i and spin component α. We will now break the site index i into a unit cell index $\boldsymbol{r} = \sum_i n_i \boldsymbol{a}_i$ and sublattice index μ, so the vectors previously indexed $|i, \alpha\rangle$ can be specified as $|\boldsymbol{r}, \mu, \alpha\rangle$. Our interest now is in the action of the translation subgroup of the space group on the space Ω_{spin}. A primitive translation shifts the entire system by a primitive vector $\boldsymbol{a}_i$. These translations act on Ω_{spin} by the linear maps T_i defined in the position basis as

$$T_i \, |\boldsymbol{r}, \mu, \alpha\rangle = |\boldsymbol{r} + \boldsymbol{a}_i, \mu, \alpha\rangle \, . \tag{2.53}$$

A general translation is then represented by $\prod_i T_i^{n_i}$, and the periodic boundaries imply that $T_i^L = \mathbb{1}$. The eigenvalues of T_i must therefore be the L'th roots of unity and the eigenvectors are plane waves,

$$|\boldsymbol{q}, \mu, \alpha\rangle \equiv \sum_{\boldsymbol{r}} \frac{1}{\sqrt{L^3}} e^{-i\boldsymbol{q}\cdot(\boldsymbol{r}+\boldsymbol{c}_\mu)} \, |\boldsymbol{r}, \mu, \alpha\rangle \, , \tag{2.54}$$

which satisfy the eigenvalue equation for the translation operator T_i with eigenvalues $e^{i\boldsymbol{q}\cdot\boldsymbol{a}_i}$,

$$\begin{aligned} T_i \, |\boldsymbol{q}, \mu, \alpha\rangle &= \sum_{\boldsymbol{r}} \frac{1}{\sqrt{L^3}} e^{-i\boldsymbol{q}\cdot(\boldsymbol{r}+\boldsymbol{c}_\mu)} \, |\boldsymbol{r} + \boldsymbol{a}_i, \mu, \alpha\rangle & (2.55) \\ &= \sum_{\boldsymbol{r}} \frac{1}{\sqrt{L^3}} e^{-i\boldsymbol{q}\cdot(\boldsymbol{r}-\boldsymbol{a}_i+\boldsymbol{c}_\mu)} \, |\boldsymbol{r}, \mu, \alpha\rangle & (2.56) \\ &= e^{i\boldsymbol{q}\cdot\boldsymbol{a}_i} \sum_{\boldsymbol{r}} \frac{1}{\sqrt{L^3}} e^{-i\boldsymbol{q}\cdot(\boldsymbol{r}+\boldsymbol{c}_\mu)} \, |\boldsymbol{r}, \mu, \alpha\rangle & (2.57) \\ &= e^{i\boldsymbol{q}\cdot\boldsymbol{a}_i} \, |\boldsymbol{q}, \mu, \alpha\rangle \, . & (2.58) \end{aligned}$$

The eigenvalues are indexed by $\boldsymbol{q}$ but are only unique up to addition of a reciprocal lattice vector $\boldsymbol{b}_i$

satisfying $\boldsymbol{b}_i \cdot \boldsymbol{a}_j = 2\pi\delta_{ij}$. The $\boldsymbol{q}$ are unique within the first Brillouin zone (FBZ) torus, defined by

$$\boldsymbol{q} = \sum_{i=1}^{d} \lambda_i \boldsymbol{b}_i \quad \lambda_i \in [0, 1). \tag{2.59}$$

The periodic boundary condition $T_i^L = 1$ implies that the eigenvalues $\exp(i\boldsymbol{q} \cdot \boldsymbol{a}_i)$ are L'th roots of unity, which means that the wavevectors $\boldsymbol{q}$ can only take discrete values. Specifically, note that $\boldsymbol{b}_i$ is orthogonal to every primitive lattice vector except $\boldsymbol{a}_i$, meaning that it is orthogonal to the plane spanned by the $\boldsymbol{a}_{j\neq i}$. The planes of constant phase therefore lie in the crystal planes of the lattice. With periodic boundaries, only certain wavevectors are commensurate with the periodicity of the super cell—the allowed plane waves must be constant on the planes of the super lattice, $\exp(i\boldsymbol{q} \cdot \boldsymbol{r}) = \exp(i\boldsymbol{q} \cdot (\boldsymbol{r} + L\boldsymbol{a}_i))$. The super lattice has primitive super-translations $L\boldsymbol{a}_i$, and corresponding reciprocal super-lattice vectors $\boldsymbol{b}_i/L$, which is the longest wavelength wavevector commensurate with the super cell.[11] Allowed wavevectors must be multiples of this "fundamental harmonic," so the set of unique wavevectors is given by

$$\boldsymbol{q} = \sum_{i=1}^{d} \frac{n_i}{L} \boldsymbol{b}_i \quad n_i \in \mathbb{Z}_L. \tag{2.60}$$

Therefore there are exactly L^d allowed wavevectors in the FBZ which are commensurate with the periodic boundaries, one for each unit cell. This must be the case in order for the Fourier and position basis to have the same number of basis vectors.

In the language of representation theory, the representation of the translation subgroup acting on Ω_{spin} splits into a sum of irreducible representations. The vector spaces spanned by the Fourier planes waves with fixed $\boldsymbol{q}$, which we denote

$$\Omega_{\text{spin}}^{(\boldsymbol{q})} = \text{span}\left\{|\boldsymbol{q}, \mu, \alpha\rangle | \forall\, \mu, \forall\, \alpha\right\}, \tag{2.61}$$

do not mix under the action of the T_i, because when restricted to these spaces the T_i act as $\exp(i\boldsymbol{q} \cdot \boldsymbol{a}_i)$ times the identity. Therefore Ω_{spin} splits into the direct sum of L^3 invariant subspaces,

$$\Omega_{\text{spin}} = \bigoplus_{\boldsymbol{q}} \Omega_{\text{spin}}^{(\boldsymbol{q})}. \tag{2.62}$$

[11] In precisely the same sense as $\boldsymbol{b}_i$ is the longest wavelength wavevector commensurate with the primitive unit cell. Just as $\exp(i(x\boldsymbol{a}_i) \cdot \boldsymbol{b}_i)$ completes one full cycle as x varies from 0 to 1 (crossing a single primitive cell), $\exp(i(x\boldsymbol{a}_i) \cdot (\boldsymbol{b}_i/L))$ completes a full cycle as x varies from 0 to L (crossing a single super cell).

This is essentially a statement of Bloch's theorem. The dimension of each subspace $\Omega_{\text{spin}}^{(\boldsymbol{q})}$ is $N_{\text{sl}} \times N$.

Note that the phase factor $\exp\left(-i\boldsymbol{q}\cdot\boldsymbol{c}_\mu\right)$ in Eq. (2.54) is a gauge freedom in the definition of the basis vectors. We choose the phase factor for each sublattice according to the sublattice position because this will be the proper basis for computing neutron scattering cross sections, since the neutron interacts with ions at their physical location in the lattice. It is sometimes also useful to consider another basis without the phase factors, which we denote

$$|\boldsymbol{q},\tilde{\mu},\alpha\rangle \equiv \sum_{\boldsymbol{r}} \frac{1}{\sqrt{L^3}} e^{-i\boldsymbol{q}\cdot\boldsymbol{r}} \,|\boldsymbol{r},\mu,\alpha\rangle\,. \tag{2.63}$$

The two bases are related by a unitary operator,

$$\Phi\,|\boldsymbol{q},\tilde{\mu},\alpha\rangle = e^{-i\boldsymbol{q}\cdot\boldsymbol{c}_\mu}\,|\boldsymbol{q},\mu,\alpha\rangle\,. \tag{2.64}$$

2.3.2 Bands of the Interaction Matrix

The interaction matrix $\mathcal{J}$ is assumed to be invariant under the action of the space group, meaning that it commutes with all generators of the representation of the space group action on Ω_{spin}. This means that $\mathcal{J}$ does not mix the irreducible subspaces and is block-diagonal in the irreducible basis. In particular in the Fourier basis, $\mathcal{J}$ will decomposes into L^d blocks of dimension $N_{\text{sl}}N \times N_{\text{sl}}N$, indexed by $\boldsymbol{q}$,

$$\langle\boldsymbol{q},\mu,\alpha|\,\mathcal{J}\,|\boldsymbol{q}',\nu,\beta\rangle = \delta_{\boldsymbol{q},\boldsymbol{q}'}\,\langle\boldsymbol{q},\mu,\alpha|\,\mathcal{J}\,|\boldsymbol{q},\nu,\beta\rangle \equiv \delta_{\boldsymbol{q},\boldsymbol{q}'}\mathcal{J}^{\alpha\beta}_{\mu\nu}(\boldsymbol{q}). \tag{2.65}$$

One can easily show by inserting resolutions of the identity in the position basis that

$$\mathcal{J}^{\alpha\beta}_{\mu\nu}(\boldsymbol{q}) = \sum_{j\in\nu} \mathcal{J}^{\alpha\beta}_{ij} e^{i\boldsymbol{q}\cdot(\boldsymbol{r}_j-\boldsymbol{r}_i)} \quad i\in\mu, \tag{2.66}$$

where the sum is over all j in sublattice ν with i fixed on sublattice μ. We will denote the matrix $\mathcal{J}$ restricted to the subspace $\Omega_{\text{spin}}^{(\boldsymbol{q})}$ by $\mathcal{J}(\boldsymbol{q})$, which will be a primary object of study throughout this thesis. The Hamiltonian in the Fourier basis reads

$$H = \frac{1}{2}\sum_{\boldsymbol{q}}\sum_{\mu\nu}\sum_{\alpha\beta} S^\alpha_\mu(\boldsymbol{q})^* \mathcal{J}^{\alpha\beta}_{\mu\nu}(\boldsymbol{q}) S^\beta_\nu(\boldsymbol{q}), \tag{2.67}$$

where $*$ denotes complex conjugation and

$$S^\alpha_\mu(\boldsymbol{q}) \equiv \langle\boldsymbol{q},\mu,\alpha|S\rangle = \sum_{\boldsymbol{r}} \frac{1}{\sqrt{L^d}} e^{-i\boldsymbol{q}\cdot(\boldsymbol{r}+\boldsymbol{c}_\mu)} S^\alpha_\mu(\boldsymbol{r}). \tag{2.68}$$

Denote the eigenvalues of the matrix $\mathcal{J}(\boldsymbol{q})$ by $\epsilon_n(\boldsymbol{q})$ and the corresponding normalized eigenvectors by $\left|\varphi_{\boldsymbol{q},n}\right\rangle$. We define the normal modes of the system to be the linear combinations

$$S_{\boldsymbol{q},n} \equiv \left\langle \varphi_{\boldsymbol{q},n} \middle| S \right\rangle . \tag{2.69}$$

The Hamiltonian Eq. (2.67) is diagonalized in this basis into a sum of quadratic normal modes,

$$H = \frac{1}{2} \sum_{\boldsymbol{q}} \sum_{n=1}^{N_{\mathrm{sl}} N} \epsilon_n(\boldsymbol{q}) \left| S_{\boldsymbol{q},n} \right|^2 . \tag{2.70}$$

The thermal equilibrium physics of this Hamiltonian in the soft spin approximation can be deduced from the eigenvalues $\epsilon_n(\boldsymbol{q})$ and their corresponding eigenvectors. In general, one expects the ground state of the system to be in the subspace spanned by the $\left|\varphi_{\boldsymbol{q},n}\right\rangle$ with minimum $\epsilon_n(\boldsymbol{q})$. It may happen, particularly in frustrated systems, that this subspace does not intersect the hard spin constraint surface $\mathcal{M}_{\mathrm{hard}}$, meaning that the hard spin ground states appear as excited states relative to the naive soft spin ground state. We will see an example of this shortly in the triangular lattice Ising antiferromagnet.

2.3.3 Long Wavelength Modes and Point Group Symmetries

It is often the case that we are interested in the macroscopic, rather than microscopic behavior of the system, in which case the most important modes are those at long wavelength. In particular, this is the case when the bottom of the spectrum of $\mathcal{J}$ is at $\boldsymbol{q} = \boldsymbol{0}$, then at low temperature the high-energy modes are depopulated and the physics is dominated by the (usually quadratically dispersing) long-wavelength modes. The long-wavelength physics of many systems can in principle be described by a coarse-grained field theory. The idea is to define a continuum field $\phi(\boldsymbol{r})$ which represents an appropriate average of the degrees of freedom within some mesoscopic volume surrounding the point $\boldsymbol{r}$ of the underlying discrete system. The coarse-grained picture is attractive because it is intuitive and allows us to bring to bear methods of field theory to discrete systems. It is therefore desirable to have a systematic method for developing long-wavelength field theories, and we will demonstrate how this can be done via a long-wavelength expansion of the interaction matrix.

The general procedure to perform a long wavelength expansion of the Hamiltonian is to begin by analyzing the interaction matrix at infinite wavelength or zero momentum, $\boldsymbol{q} = \boldsymbol{0}$. The corresponding normal modes by definition are translationally invariant, corresponding to spin configurations which are exactly the same in every unit cell of the lattice, and therefore can be analyzed by simply considering the degrees of freedom in a single unit cell. The

symmetries of a single unit cell are described by the point group, which contains only rotations and inversions/reflections. The space of $\boldsymbol{q} = \mathbf{0}$ modes will decompose further into irreducible representations of the action of the point group indexed by I,

$$\Omega_{\text{spin}}^{(\mathbf{0})} = \bigoplus_I \Omega_{\text{spin}}^{(\mathbf{0},I)}. \tag{2.71}$$

Equivalently, the eigenvectors of $\mathcal{J}(\mathbf{0})$ can be indexed by the irreducible representations (irreps) of the point group, $|\mathbf{0}, I, i_I\rangle$, where I labels the irreps and $i_I = 1, \ldots, d_I$ where d_I is the dimension of irrep I. Within each irrep subspace, the eigenvalue of $\mathcal{J}(\mathbf{0})$ will be d_I-fold degenerate, meaning that in the spectrum of $\mathcal{J}$ there is a d_I-fold band touching point at $\boldsymbol{q} = \mathbf{0}$. Our interest is in the band touching point with the lowest energy which, at low temperature, controls the long-wavelength physics as higher bands depopulate. It is also possible for multiple irreps to be degenerate, but these degeneracies are not protected by symmetries and generally require fine tuning.

The $\boldsymbol{q} = \mathbf{0}$ normal modes are

$$S_{\mathbf{0},I,i_I} \equiv \langle \mathbf{0}, I, i_I | S\rangle = \sum_{\mu,\alpha} \underbrace{\langle \mathbf{0}, I, i_I | \mathbf{0}, \mu, \alpha\rangle}_{\mathcal{T}_{(I,i_I)}{}^{(\mu,\alpha)}} \langle \mathbf{0}, \mu, \alpha | S\rangle\,, \tag{2.72}$$

where the matrix $\mathcal{T}$ is the diagonalizing unitary matrix in the $\boldsymbol{q} = \mathbf{0}$ subspace. The basis $|\mathbf{0}, I, i_I\rangle$ is the eigenbasis at $\boldsymbol{q} = \mathbf{0}$, but in general we do not know the exact eigenvectors of $\mathcal{J}$ for $\boldsymbol{q} \neq \mathbf{0}$. We use the $\boldsymbol{q} = \mathbf{0}$ unitary matrix $\mathcal{T}$ to construct an ansatz basis for the long wavelength expansion,

$$|\boldsymbol{q}, I, i_I\rangle \equiv \sum_{\mu,\alpha} |\boldsymbol{q}, \mu, \alpha\rangle \langle \mathbf{0}, \mu, \alpha | \mathbf{0}, I, i_I\rangle\,. \tag{2.73}$$

This is *not* the eigenbasis of $\mathcal{J}$ for $\boldsymbol{q} \neq \mathbf{0}$, but the claim is that it is close to the correct eigenbasis, with an error which goes to zero in the limit $\boldsymbol{q} \to \mathbf{0}$, and is the correct basis to begin the long wavelength expansion. We define the approximate normal modes for $\boldsymbol{q} \neq \mathbf{0}$ as

$$S_{\boldsymbol{q},I,i_I} \equiv \langle \boldsymbol{q}, I, i_I | S\rangle = \sum_{\mu,\alpha} \mathcal{T}_{(I,i_I)}{}^{(\mu,\alpha)} \langle \boldsymbol{q}, \mu, \alpha | S\rangle\,. \tag{2.74}$$

We obtain the matrix elements of $\mathcal{J}$ in this basis, $\mathcal{J}_{IJ}^{i_I j_J}(\boldsymbol{q}) \equiv \langle \boldsymbol{q}, I, i_I | \mathcal{J} | \boldsymbol{q}, J, j_J\rangle$, which we then expand to quadratic order in $\boldsymbol{q}$. The resulting matrix will not in general be diagonal. Assuming that one irrep (or set of degenerate irreps) has minimum eigenvalue and we are interested in energy scales smaller than the gap to the first excited irrep, the high-energy irrep modes can be integrated out. In practice this means that their matrix elements are entirely ignored, because integrating them out produces new couplings among the ground state modes at quartic and higher order in

$\boldsymbol{q}$. Limiting only to the matrix elements among the ground state irreps, the resulting (generally small) matrix can be diagonalized to obtain the long-wavelength normal modes and their quadratic dispersions. The resulting effective long-wavelength Hamiltonian can then be Fourier transformed back to direct space to obtain a field theory for the long-wavelength modes. We will see an explicit example of this procedure for spin ice in Section 2.5.2.

2.4 Frustration of Soft Spins

In this section, we will apply the soft spin analysis developed so far to a few basic models in order to demonstrate how it can be used to study frustration. First, we consider for juxtaposition the unfrustrated square lattice. We then turn to discussing the basic frustrated unit, the triangular motif, and briefly examine the triangular lattice antiferromagnet and the square lattice with second neighbor interactions. We then turn to ice models, briefly discussing kagome and square ice, then moving to three dimensions and discussing spin ice. We will only consider $N = 1$ Ising models here. The cases $N = 2, 3$ will be discussed in Chapter 6. For Ising spins with a single component $S_i^z \equiv \boldsymbol{S}_i \cdot \hat{\boldsymbol{z}}_i$, where $\hat{\boldsymbol{z}}_i$ is the quantization axis (generally the local high-symmetry easy axis), we will omit the superscript z and denote them by classical variables $S_i = \pm 1$.

2.4.1 Unfrustrated Square Lattice

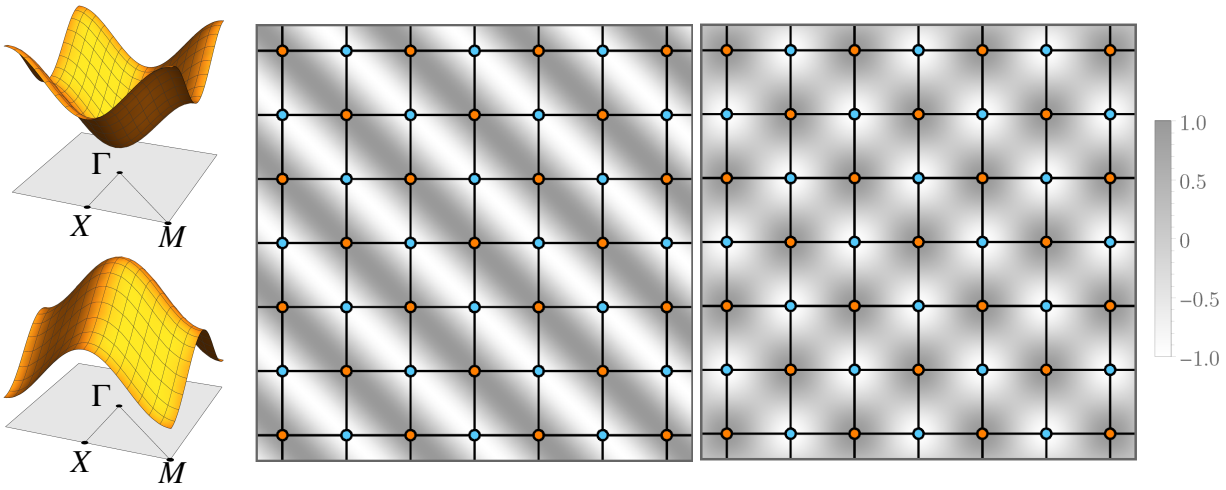
We quickly apply the above methodology to the Ising ferromagnet and anti-ferromagnet on a 2-dimensional square lattice for illustration, before turning to frustrated spin models. The square lattice is a Bravais lattice, meaning $N_{\text{sl}} = 1$ and Ising spins have $N = 1$ component. Ω_{spin} is therefore an L^d dimensional vector space. With only nearest-neighbor interactions, the Hamiltonian is

$$H = J \sum_{\langle ij \rangle} S_i S_j \tag{2.75}$$

where J is an energy scale, $J < 0$ for the ferromagnetic case. We take the primitive vectors to be $\boldsymbol{a}_1 = a\hat{\boldsymbol{x}}$ and $\boldsymbol{a}_2 = a\hat{\boldsymbol{y}}$ where a is the lattice spacing. The Fourier transform of $\mathcal{J}$ can be easily computed from Eq. (2.66). Since the square lattice has no sublattices $\mathcal{J}$ is diagonalized in the Fourier basis, and the Hamiltonian Eq. (2.67) reads

$$H = J \sum_{\boldsymbol{q}} [\cos(\boldsymbol{q} \cdot \boldsymbol{a}_1) + \cos(\boldsymbol{q} \cdot \boldsymbol{a}_2)] |S(\boldsymbol{q})|^2. \tag{2.76}$$

The band is plotted in Fig. 2.5(a) for $J < 0$ (ferromagnetic) and $J > 0$ (antiferromagnetic). In the ferromagnetic case, the band minimum is at $\boldsymbol{q} = \boldsymbol{0}$ and the band disperses quadratically upward.



(a) (b) (c)

Figure 2.5: (a) The bands of the nearest-neighbor model on the square lattice with ferromagnetic (top) or antiferromagnetic (bottom) interactions. When interactions are ferromagnetic, the the band minimum is at zone center, and when the interactions are antiferromagnetic the bands are minimum at the zone boundary. (b,c) In the antiferromagnetic case, the minimum is at the zone corner M point, which corresponds to a plane wave with two-sublattice repetition. This corresponds to the true ground state of the square lattice antiferromagnetic Ising model, i.e. the Néel state.

The 1-dimensional subspace of Ω_{spin} spanned by $|\mathbf{0}\rangle$ intersects the hard-spin constraint surface $\mathcal{M}_{\text{hard}}$, the corners of an N_{spins}-dimensional hypercube, at exactly two points, corresponding to all spins up or all spins down. In the antiferromagnetic case, the band minimum is at the zone corner and the ground state is the Néel state which has maximum crystal momentum, shown in Fig. 2.5(b). Notice that the wave is perfectly commensurate with the lattice, with lattice sites sitting precisely at the maxima and minima. Since the crystal momenta are periodic, the four zone corners are actually all equivalent, and we can represent the ground state in terms of a sum over the four symmetrically-equivalent wavevectors to obtain a standing wave pattern shown in Fig. 2.5(c). This is the prototypical example of an unfrustrated antiferromagnet, following from the fact the the lattice is bipartite and so supports a 2-sublattice symmetry broken ground state.

2.4.2 Fully Frustrated Clusters: Triangular and Simplex Motifs

Before moving on to lattice models for frustrated magnetism, we first consider the minimal frustrated unit: the antiferromagnetic triangle, and the general concept of a frustrated cluster. This will allow us to distinguish two types of frustration effects.

As we discussed earlier, the simplest antiferromagnet involves just two spins. The configuration space is shown in Fig. 2.6(a). The excited state is the ferromagnetic $S_1 = S_2$ state, which sits in the subspace spanned by $|\psi_+\rangle \propto |1\rangle + |2\rangle$, indicated by the red line. The ground state is the antiferromagnetic $S_1 = -S_2$ which sits in the orthogonal subspace spanned by the ground state eigenvector $|\psi_-\rangle \propto |1\rangle - |2\rangle$. Note that the hard spin manifold (the corners of the square) is completely contained in these two eigenspaces of the interaction matrix. The SCGA ground state is the subspace indicated by the blue line, and the spherical model ground state is given by the intersection of the circle with the blue line, which is the true hard spin ground state. The correlation matrix has a spectral expansion

$$\mathcal{G} = \frac{1}{\lambda_0 + 2\beta J} |\psi_+\rangle \langle\psi_+| + \frac{1}{\lambda_0} |\psi_-\rangle \langle\psi_-| . \tag{2.77}$$

In the limit of zero temperature $\beta \to \infty$, the excited state is fully suppressed and gives no contribution to the correlations.

Now consider the same problem but for three spins on triangle, which is the minimal frustrated unit. The Hamiltonian can be written as

$$H = J(S_1S_2 + S_2S_3 + S_3S_1) = \frac{1}{2} \begin{pmatrix} S_1 & S_2 & S_3 \end{pmatrix} \begin{pmatrix} 0 & J & J \\ J & 0 & J \\ J & J & 0 \end{pmatrix} \begin{pmatrix} S_1 \\ S_2 \\ S_3 \end{pmatrix} \tag{2.78}$$

The interaction matrix has an eigenvector $|\psi_+\rangle \propto |1\rangle + |2\rangle + |3\rangle$ with eigenvalue $+2J$, and the remaining two eigenvectors are degenerate with eigenvalue $-J$. Adding and subtracting the minimum eigenvalue on the diagonal as per Eq. (2.42) is equivalent to writing the Hamiltonian as

$$H = -\frac{J}{2}(S_1^2 + S_2^2 + S_3^2) + \frac{J}{2}(S_1 + S_2 + S_3)^2. \tag{2.79}$$

The first term yields the spherical model ground state energy, and the second term is exactly $(J/2)|\langle\psi_+|S\rangle|^2$. The configuration space is shown in Fig. 2.6(b), with the subspace spanned by $|\psi_+\rangle$ indicated in red. The two ferromagnetic excited states (all-up and all-down, or all-in and all-out) lie in this subspace. The two-dimensional subspace orthogonal to $|\psi_+\rangle$ spanned by the remaining two degenerate eigenvectors is shown by the blue disk. Its intersection with the

spherical constraint surface is indicated by the blue circle. Notice that *this subspace does not include any of the hard spin states*. There are six hard spin ground states for this system with the three spins satisfying $S_i = S_j = -S_k$, indicated by the blue corners of the cube. These states are displaced from the soft spin ground state subspace along $|\psi_+\rangle$, and therefore have higher energy than the soft spin ground states lying on the spherical constraint surface indicated by the blue circle, because there are no hard spin states which can fully minimize $S_1 + S_2 + S_3$. Thus the ground state energy of this system is larger than that of the spherical model.

This inability to reach the ground state subspace of $\mathcal{J}$ on the hard spin constraint surface is a manifestation of frustration, but it is not a necessary criteria of frustration. To see why, consider a fully connected frustrated unit with four spins (e.g. a tetrahedron) interacting antiferromagnetically. Writing down the obvious analogues of Eqs. (2.78) and (2.79) we again find that the interaction matrix has one excited eigenvector $|\psi_+\rangle \propto |1\rangle + |2\rangle + |3\rangle + |4\rangle$, and the remaining three eigenvectors are degenerate. The hard spin configuration space is given by the sixteen corners of a four-dimensional hypercube, which is illustrated in Fig. 2.6c as a tesseract. The ground state subspace orthogonal to $|\psi_+\rangle$ satisfies $S_1 + S_2 + S_3 + S_4 = 0$ and is three dimensional, and in fact there are six hard spin states which lie in this subspace, colored blue. This model is certainly frustrated, in that ground states do not minimize every bond energy and some bonds remain ferromagnetically aligned, but the hard spin states are contained in the SCGA ground state subspace.

We can generalize these observations to a fully-connected cluster of n spins (an n-simplex), where each bond is antiferromagnetic. Every trio of spins in the cluster forms an anti-ferromagnetic triangle, so such clusters are maximally frustrated units. The Hamiltonian for the cluster can be written as

$$H_c = J \sum_{\langle ij \rangle \in c} S_i S_j = -\frac{J}{2} \sum_{i \in c} S_i^2 + \frac{J}{2} \left(\sum_{i \in c} S_i \right)^2 . \tag{2.80}$$

The interaction matrix always has one excited eigenvector $|\psi_+\rangle \propto \sum_{i \in c} |i\rangle$ and the soft spin ground state subspace is the $n-1$-dimensional subspace orthogonal to this vector. In the spherical approximation the ground state subspace is an $n-1$-sphere. For n even, the hard spin manifold (corners of an n-cube) intersects these ground states, corresponding to states for which $\sum_{i \in c} S_i = 0$, while for n odd the sum can at best be ± 1 and the hard spin ground states do not intersect the SCGA ground state manifold.

This means we should distinguish two mechanisms of frustration. The first is the usual one: not every pairwise interaction can be simultaneously satisfied. This mechanism is present in all frustrated models, it is a necessary but not sufficient criteria for frustration, and we will see an example in the next section where it is satisfied despite the system having a unique ground state. The second is an additional type of frustration which occurs when frustrated clusters have an odd number of *Ising* spins ($N = 1$): the constraint $\sum_{i \in c} S_i = 0$ cannot be satisfied exactly—the best

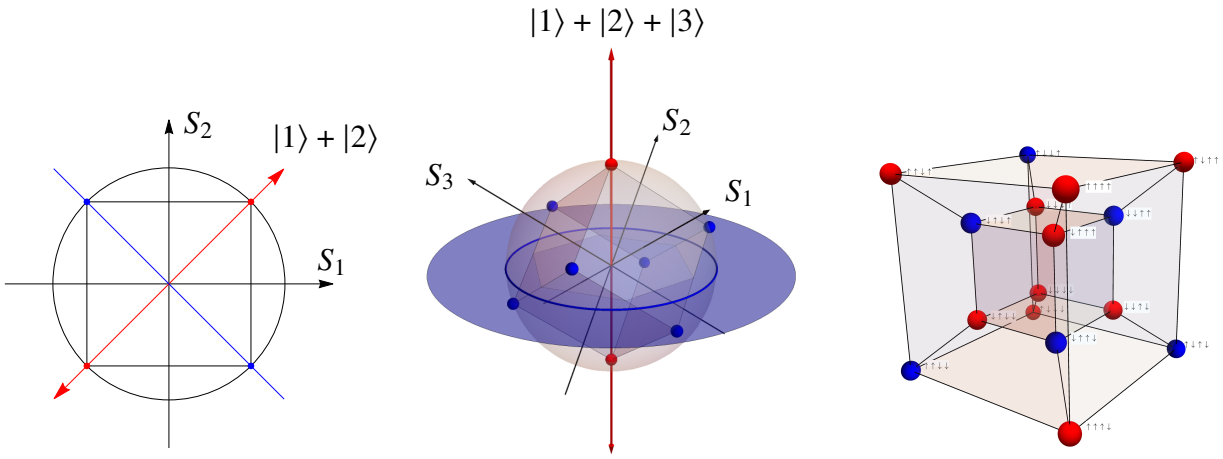
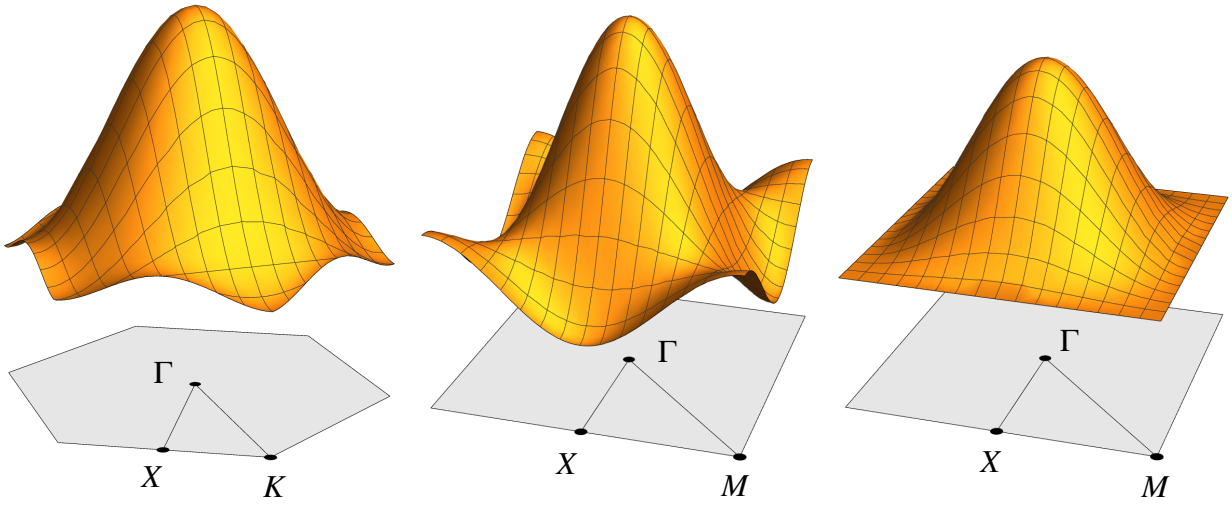


Figure 2.6: The SCGA ground state manifold for a single frustrated unit of n spins is the $n-1$-dimensional space orthogonal to $\sum_{i=1}^{n} |i\rangle$. This is depicted for $N = 1$ spins with (a) $n = 2$, (b) $n = 3$, and (c) $n = 4$. The 1-dimensional subspace spanned by $\sum_{i=1}^{n} |i\rangle$ is indicated in red and the orthogonal ground state subspace is depicted in blue. For $N = 1$ spins with n odd, such as in (b), the soft spin ground state space (blue disk) is not intersected by the hard spin manifold (cube corners), and the hard spin ground states violate the constraint $\sum_{i=1}^{n} S_i = 0$. For n even, such as (a) and (c), some of the hard spin states do lie in the soft spin ground state subspace, and those are the hard spin ground states of the cluster.



(a) Triangular lattice AFM

(b) $J_2/J_1 = 1$ square lattice

(c) $J_2/J_1 = 1/2$ square lattice

Figure 2.7: (a) Bands of the triangular lattice antiferromagnet, which has a minimum at the K-points despite being highly frustrated. (b) The square lattice with second neighbor interactions is unfrustrated despite being constructed of frustrated units. (c) At the special point $J_2/J_1 = 1/2$ the square lattice model becomes frustrated again, indicated by the completely flat dispersion along the zone boundary.

possibility is $\sum_{i\in c} S_i = \pm 1$—and the ground states do not correspond to the SCGA or spherical model ground states. This second type of frustration is relieved when $N > 1$, where the constraint becomes $\sum_{i\in c} \boldsymbol{S}_i = \mathbf{0}$. In the next section we will quickly consider two lattice models: the triangular lattice and the square lattice with second neighbor interactions, which are both built from edge-sharing frustrated units. The triangular lattice antiferromagnet is frustrated by the fact that the hard spin constraint surface does not intersect the SCGA ground state subspace, but is not frustrated for $N > 1$ spins.

2.4.3 Edge Sharing Frustrated Units: Triangular vs. J_1-J_2 Square Lattices

We now consider what happens when we build lattices out of frustrated units. We consider two juxtaposed examples: the triangular lattice and the square lattice with crossings, both with antiferromagnetic interactions and Ising spins. The triangular case is frustrated and has an extensive ground state degeneracy while the square lattice with crossings has a subextensive ground state despite being built of frustrated clusters. Our discussion will only be qualitative,

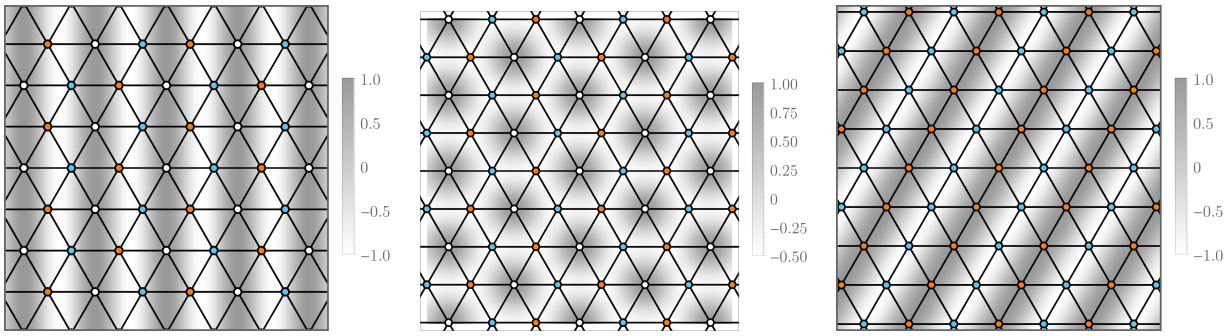
based on the intuition we have built up so far.

First we consider the triangular lattice with nearest-neighbor antiferromagnetic interactions. The bands are shown in Fig. 2.7(c). The bands have a minimum at the zone corner K points, just like the square lattice antiferromagnet. Unlike the square lattice the triangular lattice is not bipartite and cannot host the two-sublattice Néel state. Rather, the triangular lattice is tripartite. In Fig. 2.8(a), we show the wave corresponding to a band minimum K point wavevector, which has three-sublattice repetition. It is not hard to see why this is the preferred mode based on our discussion in the last section. The Hamiltonian for the triangular lattice can be written as a *sum over frustrated units*,

$$H = \frac{1}{2}\sum_t \left(J \sum_{\langle ij \rangle \in t} S_i S_j \right) = -\frac{1}{4} J \sum_i S_i^2 + \frac{J}{4} \sum_t \left(\sum_{i \in t} S_i \right)^2 , \tag{2.81}$$

where the $1/2$ in front of the first sum accounts for double counting of triangle edges when summing over all triangles. The naive ground state should satisfy $\sum_{i\in t} S_i = 0$ on every triangle, and the SCGA ground state manifold is orthogonal to every vector of the form $\sum_{i\in t} |i\rangle$. There are is one constraint per spins (one per triangle), so we would naively expect the ground state to be subextensive, which is consistent with the fact that the bands have minima at two wavevectors (the two K points). For soft spins, we could satisfy the constraint by, say, setting one spin to $S_i = 1$ and the other two $S_i = -1/2$. This is precisely satisfied by the band minimum K point eigenvector. The standing wave for this wavevector is shown in Fig. 2.8(b), where every triangle has value $+1$ on one corner and $-1/2$ on the remaining two.

But this is not possible for hard Ising spins, the best possible is to make two spins equal one each triangle. In Fig. 2.8(c), we show a ground state where on each triangle two spins are equal, but this is a plane wave with wavevector corresponding to the X point, which is not the band minimum! This model was solved by Wannier [77], who showed that there are an extensive number of ground states which can be constructed by separating the system into three sublattices, setting one to have $S_i = +1$, another to have $S_i = -1$, and the remaining spins on the third sublattice can be completely arbitrary, depicted by the orange, blue, and white sublattices in Fig. 2.8(a,b). These ground states are linear combinations of eigenvectors with multiple $\boldsymbol{q}$ vectors, they clearly carry some weight from the eigenvectors at the K point, and much also contain weight from eigenvectors with energy higher than that of the X point. Thus the frustration of the triangular lattice is driven by the fact that each frustrated unit has an *odd* number of spins, and the hard spin manifold does not lie in the soft spin ground state subspace. The triangular lattice is an interesting outlier among frustrated lattices with edge sharing frustrated units, where the frustration is driven by the inability to satisfy $\sum_{i\in c} S_i = 0$ (indeed for spins with $N > 1$ spin components, the triangular lattice is rather famously not frustrated, exhibiting 120° three-sublattice order, see [78] and references therein). This is a



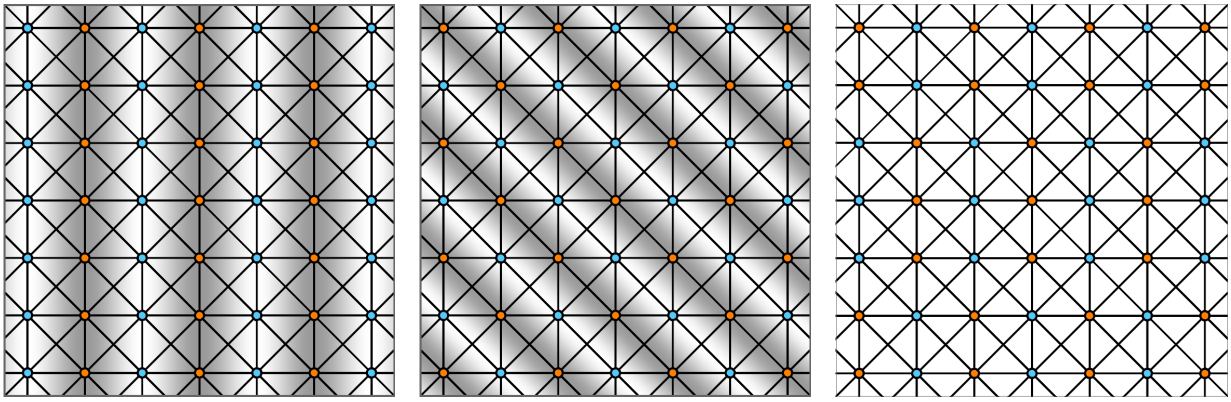
(a) (b) (c)

Figure 2.8: (a) Ground states of the triangular lattice with hard Ising spins cannot be formed from the minimum eigenvalue wavevector of the interaction matrix, which is a planewave with three sublattice repetition. (b) Standing wave with three symmetry-equivalent ground state wavevectors. (c) Plane wave with wavevector corresponding to the X point. The band is not minimum at this wavevector and yet this is a ground state.

distinct form of frustration from the other frustrated models we will consider in the remainder of this thesis.

We will quickly illustrate with another edge sharing example: the J_1-J_2 square lattice antiferromagnet with second neighbor interactions, shown in Fig. 2.9. When $J_1 = J_2$ the lattice is made of edge-sharing frustrated units—each crossed square is a fully connected cluster, and naively should satisfy $\sum_{i \in c} S_i = 0$. This naive expectation is incorrect, however, which can be seen by inspecting the band structure of the interaction matrix. The bands are shown in Fig. 2.7(b) for $J_1 = J_2$, and we can see that despite being made of frustrated units, there is a minimum at the X point. Indeed this model is not frustrated; the corresponding planewave is shown in Fig. 2.9(a), which corresponds to a hard spin configuration with alternating stripes. However, the Néel state, which also satisfies the constraint, is not a ground state, which is clear since the M point is not a minimum of the band. Both of these states satisfy the zero-total-spin constraint on each crossed square, but one is a ground state and the other is not. Furthermore, note that in the stripe ground state all of the vertical bonds are unsatisfied, meaning that this state satisfies the naive definition of being "frustrated", in the sense that all pairwise interactions are not simultaneously satisfied, but the system does not exhibit an extensive or even subextensive ground state degeneracy, and so is not a frustrated spin model.

The reason that the $J_1 = J_2$ square lattice is not frustrated is because the Hamiltonian does not



(a) (b) (c)

Figure 2.9: (a) when $J_1 = J_2$ the ground state of the system a plane wave corresponding to the X point. The system is unfrustrated despite each crossed square individually is a frustrated unit. (b) The Néel state is not a ground state when $J_1 = J_2$ despite satisfying $\sum_{i\in c} S_i = 0$ on each cluster. It is a ground state when $J_2 = J_1/2$, because at that point the Hamiltonian splits into a sum over frustrated units. (c) An example of a frustrated ground state satisfying the zero-net-spin condition on each crossed square.

split into a sum over clusters. We would like the Hamiltonian to split as

$$H = E_0 + \frac{J}{2} \sum_c \left(\sum_{i\in c} S_i \right)^2 \tag{2.82}$$

for the system to be frustrated. Because the frustrated units touch on their edges, the sum over clusters *double counts* each shared edge, but the interior edges are not double counted. This means that the Hamiltonian splits into this frustrated form if we take $J_2 = 2J_2$. An example of a frustrated ground state is shown in Fig. 2.9(c), satisfying the zero-total-spin condition on every square. The interaction matrix band for this ratio is shown in Fig. 2.7(c), and we can directly observe the frustration as the flat degeneracy of the band all along the zone boundary. As we saw in the last section, for frustrated units with an even number of spins the hard spin ground states will lie in the SCGA ground state subspace. The lesson is that for frustrated units with an even number of spins, the frustration will always be visible as a degeneracy at the bottom of the bands of the interaction matrix (in this case the SCGA ground state subspace is subextensive, since the band is only flat along the one-dimensional zone boundary). This contrasts with the triangular lattice antiferromagnet, whose Hamiltonian also splits into a sum over clusters in the form of Eq. (2.82), but with an odd number of spins per cluster, and whose interaction matrix band has a minimum at the zone corner.

We can extrapolate to say that *any* lattice of edge sharing fully connected frustrated units requires the tuning of all overlapping edge couplings (J_1) to be exactly twice the coupling of all edges internal to the frustrated unit, so that the Hamiltonian will split into a sum of individually frustrated units as in Eq. (2.82). Another example of this form of frustration is the honeycomb lattice with second and third neighbor interactions, which can be be tuned to a fully frustrated point where the Hamiltonian splits into a sum over edge-sharing hexagons, a model studied by Ref. [79]. This is experimentally problematic, as such fine tuning is almost never guaranteed in materials. In the next section, we discuss ice models, which are constructed from *corner sharing* frustrated units. The corner sharing nature prevents the double counting issue on edges and allows us to achieve frustration with only nearest-neighbor interactions, without fine tuning. These models are frustrated by an extensive SCGA ground state subspace, which manifests in the form of a flat band in the spectrum of $\mathcal{J}$.

2.5 Corner Sharing Frustrated Units: Ice Models and Line Graphs

In this section, we introduce and discuss the ice models, which are Ising antiferromagnets on lattices made of corner sharing frustrated units. We will see that the ice models are all frustrated by the presence of a flat band in the spectrum of the interaction matrix. They are all related to the so-called "line graph" construction [80], a connection we will make extensive use of in Chapter 3. The line graph[12] construction is illustrated in Fig. 2.10 for the (a) checkerboard, (b) kagome, and (c) pyrochlore lattices; the construction is as follows, moving from left to right:

1. Start with a "parent lattice" (far left), here we show the (a) square, (b) honeycomb, and (c) diamond lattices.
2. Add a new site at the center of each edge of the parent lattice, these are the vertices of the line graph lattice.
3. Add an edge (dashed lines) between any two of the new vertices if their corresponding bonds in the parent lattice meet at a vertex in the parent lattice.
4. Remove edges and vertices of the parent lattice.

[12] Line graph lattices are also referred to by Henley as "medial latttices" [55]. The relation between line graphs and frustrated magnetism is touched on briefly in [81].

This process naturally constructs a fully connected cluster of spins surrounding each of the parent lattice vertices, and these clusters naturally meet at their corners. Every vertex of the line graph is shared between exactly two clusters. The line graph construction can be performed for any parent lattice, but these three are the most ubiquitous line graph geometries studied with experimental realizations. Note that in these three examples, the underlying or "parent" lattice was bipartite, meaning that the clusters can be separated into A and B clusters. This has the consequence that for any cluster of the line graph, its neighboring clusters are not themselves neighbors of each other.

By an "ice model," we specifically mean an Ising model with nearest-neighbor antiferromagnetic interactions on a line graph lattice. The Hamiltonian of such a system can be written as

$$H = \frac{J}{2} \sum_{ij} S_i \mathcal{A}^{\text{lg}}_{ij} S_j, \tag{2.83}$$

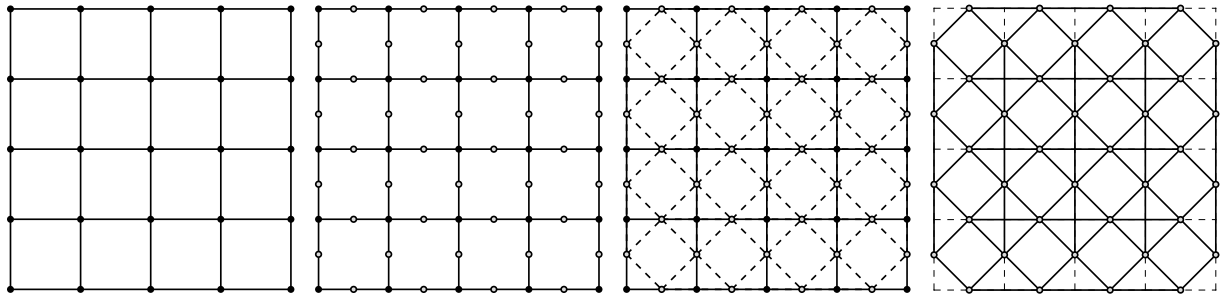
where $\mathcal{A}^{\text{lg}}$ is the adjacency matrix of the line graph lattice, whose matrix elements are

$$\mathcal{A}^{\text{lg}}_{ij} = \begin{cases} 1 & i, j \text{ are nearest-neighbors,} \\ 0 & \text{otherwise.} \end{cases} \tag{2.84}$$

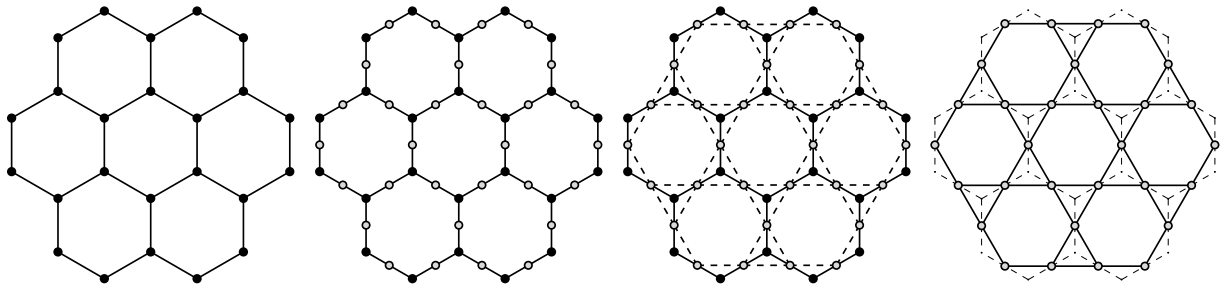
Much of the discussion of Chapter 3 will be devoted to understanding the adjacency matrices of line graph lattices and their unique spectrum. In this section, we will summarize the SCGA perspective on the frustration of these lattices.

While the work presented in the remainder of this thesis will focus almost entirely on the physics of pyrochlore magnetism, with an emphasis on spin ice (i.e. the ice model in the pyrochlore lattice), we note that there is a very large field studying two-dimensional or "artificial" spin ice constructed from nanometer-sized magnet island [82–85] forming a square or honeycomb array, and whose nearest-neighbor interactions can be described by the corresponding checkerboard and kagome lattices. The two-dimensional ice models are closely related to so-called vertex models [86, 87], which were originally studied as a model for layered ferroelectrics, where much of the early insights about the physics of ice models was made [88, 89]. Vertex models have a rich theoretical structure, admitting exact solutions [90, 91] and have deep mathematical structure connected to topics as far flung as knot theory [92, 93].

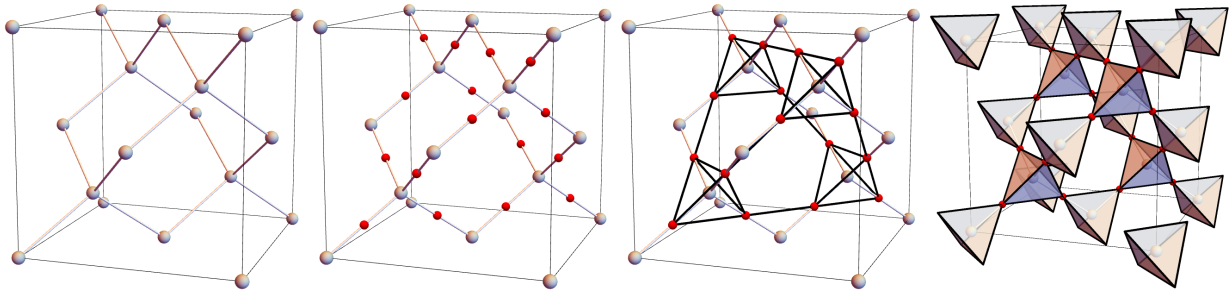
We quickly note other corner sharing frustrated geometries. First, Henley provides a list of line graph lattices of interest in Ref. [55]. In Ref. [94], the ruby and "fully connected kagome" lattices were considered, which are the line graphs of the kagome and triangular lattices, respectively. In three dimensions other corner sharing geometries exist, in particular the hyperkagome lattice, which is a network of corner sharing triangles [95–99], the so-called hyper-hyperkagome lattice, which has three triangles meeting at each corner [100, 101], and a related structure called the trillium lattice [102]. We will not be studying these lattices, but their physics is amenable to the same analyses discussed in this chapter.



(a)

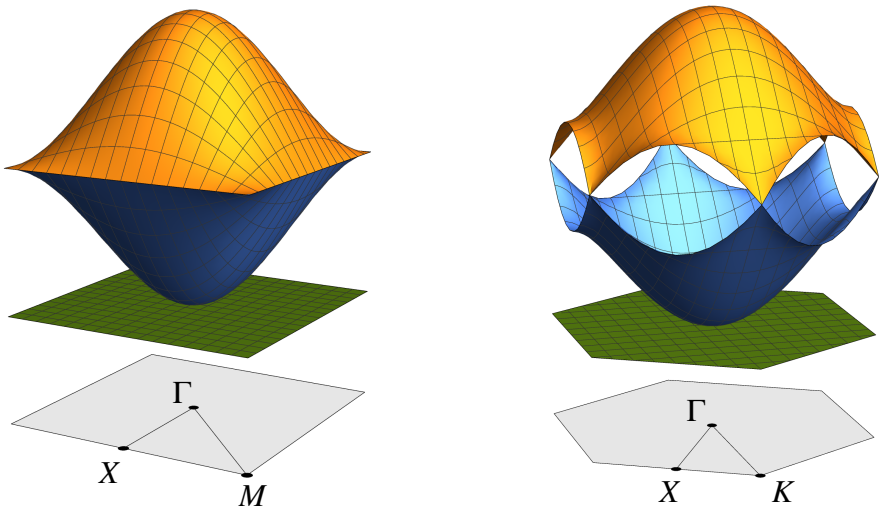


(b)



(c)

Figure 2.10: Line graph construction of the (a) checkerboard, (b) kagome, and (c) pyrochlore lattices.



(a) (b)

Figure 2.11: Bands of the adjacency matrices for the (a) checkerboard and (b) kagome lattices, demonstrating flat bands at the bottom of their spectrum, with a band touching point at zone center.

2.5.1 Line Graph Lattices and Flat Bands

Line graph lattices have generated a large amount of interest, and not only in frustrated magnetism. Various studies have looked at tight binding models on line graphs [103, 104], the Hubbard model on line graphs [105, 106], and topological physics on line graphs [107, 108]. This research falls within the broader scope of interest in strongly correlated physics driven by the presence of flat momentum bands in a variety system, which has been driven forward by the novel properties of twisted bilayer graphene driven by nearly flat bands [109]. In electronic systems, the primary idea of a flat band is that it corresponds to a collection of states which are highly localized in direct space, i.e. can be seen to have either no kinetic energy or very large mass. These states are not necessarily interesting, since they could simply refer to the atomic limit of a band insulator where electrons are fully localized, but flat bands are highly sensitive to perturbations which can drive these states to have novel correlated properties [110] such as many body localization [111]. For example, they may have non-trivial topology [112–114], and a significant amount of work has gone into classifying flat bands with non-trivial topology using topological quantum chemistry [17, 115] resulting in a large catalogue of flat band materials [116]. In frustrated magnetism, as we will see in much more detail in Chapter 3, the flat bands of line graph lattices correspond directly to states localized on loops of the lattice [103].

Consider the edge sharing cases of the triangular and $J_2/J_1 = 1/2$ square lattices discussed in

the previous section. Those are both Bravais lattices whose interaction matrix has a single band. Both have Hamiltonians that split into a sum of frustrated clusters given in Eq. (2.82). Each cluster imposes one constraint on the low energy manifold of states, that $\sum_{i\in c} S_i = 0$. Since these are Bravais lattices, this amounts to one constraint per unit cell, resulting in a subextensive SCGA ground state subspace despite the frustration.

Line graph lattices are composed of corner sharing frustrated clusters, and their nearest-neighbor antiferromagnetic Hamiltonians always split into a sum over frustrated units of the form of Eq. (2.82). Line graph lattices are never Bravais lattices, always having $N_{\text{sl}} > 1$ sublattices, and so their interaction matrix always has multiple bands. The number of clusters per cell N_c is necessarily less than the number of sublattices. Each cluster provides one energetic constraint $\sum_{i\in c} S_i = 0$. The SCGA ground state subspace is therefore orthogonal to the $L^d N_c$ vectors $\sum_{i\in c} |i\rangle \ \forall\, c$, and has dimension $L^d(N_{\text{sl}} - N_c)$, which is extensive. This means that the interaction matrix $\mathcal{J}$ has a corresponding extensive degeneracy at the bottom of its spectrum, which manifests as $(N_{\text{sl}} - N_c)$ flat bands.

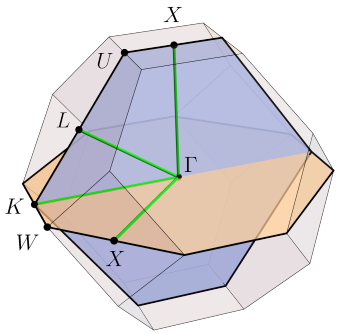
The checkerboard ($N_{\text{sl}} = 4$), kagome ($N_{\text{sl}} = 3$), and pyrochlore ($N_{\text{sl}} = 4$) line graph lattices are depicted in Fig. 2.10, consisting of corner sharing squares, triangles, and tetrahedra, respectively, all of which have $N_c = 2$ (since their parent lattices are bipartite). The corresponding bands for their nearest-neighbor antiferromagnetic interaction matrices (equivalently, their adjacency matrices) are shown in Fig. 2.11 and Fig. 2.12.

2.5.2 Long Wavelength: Emergent Field Theory and Gauss' Law

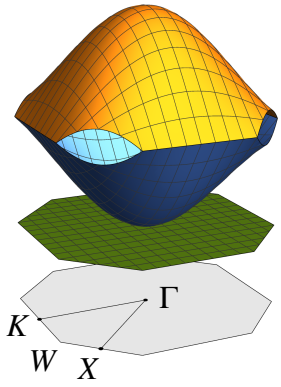
In this section we will analyze the long-wavelength description of ice models utilizing the method discussed in Section 2.3.3, by applying it to the nearest-neighbor pyrochlore Ising antiferromagnet. This model is referred to as nearest-neighbor spin ice, and will be abbreviated NNSI. The derivation was originally demonstrated by [76, 81], though the connection to the point group was not demonstrated at that time. The same procedure can be performed equivalently for other line graph lattices with similar results. This method will also be utilized in Chapter 5 for a model complementary to spin ice, and in Chapter 6 for $N = 3$ anisotropic Heisenberg spins.

The space group of the pyrochlore lattice is denoted $Fd\bar{3}m$, space group 227 in the International Tables of Crystallography [117], which is the same as for its parent diamond lattice. The point group of the pyrochlore lattice unit cell is O_h. It has two important subgroups, the first is the symmetry group of a single tetrahedron T_d, and the second is the symmetry group for two tetrahedra touching at a corner (a tetrahedral antiprism) D_{3d}. We will call these "site symmetry groups". D_{3d} is the site symmetry group for the pyrochlore lattice, meaning that it contains the symmetry elements which leave a single site of the lattice fixed, shown in Fig. 2.13(a). T_d is

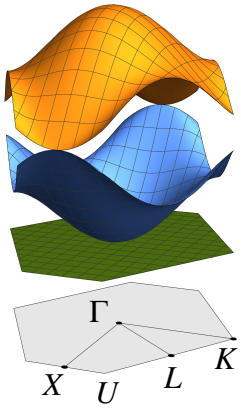
Γ	$\{000\}$
X	$\{001\}$
U	$\{\frac{1}{4}\frac{1}{4}1\}$
L	$\{\frac{1}{2}\frac{1}{2}\frac{1}{2}\}$
K	$\{\frac{3}{4}\frac{3}{4}0\}$
W	$\{\frac{1}{2}10\}$



(a)

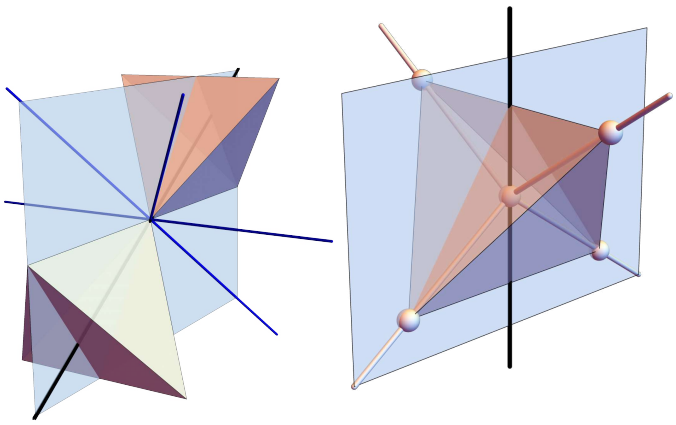


(b)



(c)

Figure 2.12: (a) The first Brillouin zone (FBZ) of the face centered cubic (FCC) lattice with high-symmetry points labeled, along with a table of their positions as fractions of the cubic reciprocal lattice vectors, with the $(hk0)$ (tan) and (hhl) (purple) high-symmetry reciprocal planes indicated. See Appendix A for detailed figures showing the structure of the pyrochlore lattice and its relation to the FCC lattice. (b) The bands of the pyrochlore lattice in the $(hk0)$ plane. (c) The bands of the pyrochlore lattice in the (hhl) plane. There are two flat bands at the bottom of the spectrum which are flat in the entire zone.



(a) (b)

T_d irrep	E	$8C_3$	$3C_2$	$6S_4$	$6\sigma_d$
A_1 (**1**)	1	1	1	1	1
A_2 ($\bar{\mathbf{1}}$)	1	1	1	-1	-1
E (**2**)	2	-1	2	0	0
T_1 ($\bar{\mathbf{3}}$)	3	0	-1	1	-1
T_2 (**3**)	3	0	-1	-1	1

(c)

Figure 2.13: Points group symmetries of the pyrochlore lattice. (a) The site symmetry group D_{3d} of the pyrochlore lattice includes a single C_3 easy axis (black), three mirror planes (one shown in light blue), three C_2 axes (dark blue) each orthogonal to one mirror plane, and inversion through the site. (b) The site symmetry group of the diamond lattice T_d contains four C_3 axes (white), three C_2 axes (black), three fourfold roto-inversion axes (black, coinciding with the C_2 axes), and six mirror planes (light blue). (c) The character table for T_d; columns are equivalence classes and rows are irreducible representations.

the site symmetry group of the parent diamond lattice, and is the symmetry group of a single tetrahedral frustrated unit, shown in Fig. 2.13(b).[13]

The character table for T_d is shown in Fig. 2.13(c). Vertical columns list the equivalence classes or conjugacy classes of symmetry elements of T_d, where E indicates the identity, C_n indicates n-fold rotation, S_n indicates n-fold rotation followed by a mirror in the plane orthogonal to that axis, and σ_d indicates mirror planes. There are 24 total elements of this group. The horizontal rows of the table label the irreducible representations. There are always an equal number of irreducible representations as conjugacy classes. The numbers in the table are the characters of the symmetry operations in the different irreps, which is the trace of the corresponding matrix. The letters indicate dimension of the representation, which can be read off from the character of the identity: A is 1-dimensional, E is 2-dimensional, and T is three-dimensional. There are two A and two T irreps, and they differ only by the sign of the characters of symmetry elements containing an inversion (improper rotations and mirror planes). A_1 is called the trivial representation where every symmetry element acts trivially as the identity, while A_2 is the sign representation, where inversion acts as multiplication by -1. We can call these respectively the "scalar" and "pseudoscalar" representations, and we will also indicate them as $\mathbf{1}$ and $\bar{\mathbf{1}}$, respectively. The two T irreps are both rotation representations. To see this, consider four three-dimensional vectors z_μ, $\mu = 1, 2, 3, 4$, which point from the center of the tetrahedron to each of its four corners. Under a symmetry transformation of the tetrahedron, the four corners are permuted,[14] which is equivalent to a rigid rotation of the four vectors z_μ. The T_2 irrep is precisely this three dimensional rotation representation, where the z_μ transform as vectors. the T_1 irrep differs by the sign of the parity changing elements, and therefore corresponds to a *pseudo*vector representation, where the z_μ transform as angular momentum vectors.[15] We will therefore indicate these two representations

[13] A note on reading symmetry notations: The notation $Fd\bar{3}m$ is called Hermann-Mauguin or International notation, where F indicates FCC Bravais lattice symmetry, d indicates the presence of a non-symmorphic glide symmetry along half an FCC lattice vector, $\bar{3}$ indicates a threefold roto-inversion axis (black axis in Fig. 2.13(a)), and m indicates a mirror plane. A roto-inversion is a rotation followed by an inversion, also called an improper rotation. Note that a mirror plane is equivalent to a 2-fold rotoinversion, i.e. $m \equiv \bar{2}$. The notations T_d and D_{3d} are Shönflies notation for the point groups, which in International notation are $\bar{4}3m$ and $\bar{3}2/m$ respectively, where $2/m$ indicates that the mirror plane is orthogonal to the 2-fold rotation axis (the blue mirror plane in Fig. 2.13(a) is orthogonal to one of the C_2 axes). The symmetry equivalence classes in the column headers of Fig. 2.13(c) of E, C_3, C_2, S_4, and σ_d in Schönflies notation are equivalently 1, 3, 2, $\bar{4}$, and m in International notation.

[14] The group T_d has exactly 4! elements, and is in fact isomorphic to the symmetric group of permutations of four elements.

[15] Under an inversion, a normal vector remains invariant, while a pseudovector changes sign. Given a basis $\boldsymbol{e}_i$ for a vector space, a vector can be expanded as $\boldsymbol{v} = \sum_i v^i \boldsymbol{e}_i$, where v^i are its components in this basis. Under an inversion, *both* the basis vectors $\boldsymbol{e}_i$ and the vector components v^i change sign, so that $\boldsymbol{v}$ remains invariant (it still "points in the same direction"). On the other hand, a pseudovector's component *do not* change sign, so that under inversion a pseudovector obtains an overall change in sign. The most common pseudovectors in physics are angular momenta

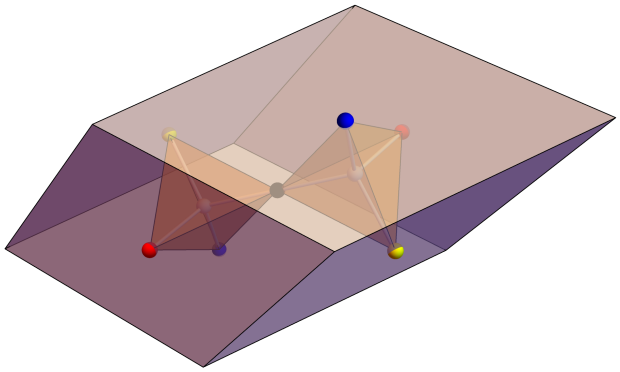


Figure 2.14: A single FCC unit cell of the pyrochlore lattice, which contains two tetrahedra (two diamond lattice sites, white spheres) and four pyrochlore sites, colored black, red, blue, and yellow. One site sits in the middle (black). The remaining sites sit in the center of a face of the cell and are each shared between two FCC cells, thus yielding only four sites per cell. Additional detailed figures for the pyrochlore lattice can be found in Appendix A.

by **3** and $\bar{\mathbf{3}}$ respectively. We will discuss the two-dimensional E and **2** irrep in Chapter 6.

As we laid out in Section 2.3.3, the long-wavelength expansion of the Hamiltonian begins by studying the $\boldsymbol{q} = \mathbf{0}$ sector of the interaction matrix, which is equivalent to the interaction matrix for a single unit cell. For the pyrochlore lattice a single unit cell contains two tetrahedra, shown in Fig. 2.14. According to the method laid out there, we should perform our analysis with respect to the full point group O_h, but this is unnecessary for our purposes, because O_h decomposes into two copies of the single tetrahedron site symmetry group T_d, which are mapped to each other by the C_2 and inversion symmetries of the D_{3d} pyrochlore site symmetry group. The irreps of O_h are equivalent to two copies of the irreps of T_d, but with an "even" and "odd" realization, depending how their sign change under the D_{3d} inversion. We are only interested here in the physics which does not spontaneously break the inversion symmetry between the two tetrahedron (aka breathing anisotropy), so we can perform our analysis just with respect to T_d. Equivalently, we can perform the analysis for a single tetrahedron, which is a common first order approximation method for pyrochlore physics.

vectors. As a basic example, consider a rotating disk. We can utilize the right hand rule to assign to the rotation an angular velocity vector $\boldsymbol{\omega}$ which points orthogonal to the disk surface. Applying an inversion means either reversing the direction of rotation (active transformation) or changing to a left hand rule (passive transformation), and in either case the direction of $\boldsymbol{\omega}$ is reversed. As a vector, $\boldsymbol{\omega}$ indicates the axis about which the disk is rotating, but the assigning of a direction to $\boldsymbol{\omega}$ is dependent on whether we choose a right handed or left handed convention and is not intrinsic, because the rotation itself occurs in the two-dimensional plane orthogonal to ω. Therefore psedovectors are often called axial vectors, while normal vectors are often called polar vectors.

For a single tetrahedron with four Ising spins S_i, the soft spin configuration space is four dimensional, spanned by vectors $|1\rangle$, $|2\rangle$, $|3\rangle$, $|4\rangle$. Under the action of T_d the four spins are permuted. Since spins are angular momenta, they transform as pseudovectors, meaning that their components do not change sign under inversion. This is a four dimensional permutation representation of T_d, which we denote $\mathbf{4}$, given by the 24 four-dimensional permutation matrices. This representation is reducible, since according to the table in Fig. 2.13(c) the largest irreducible representations of T_d are three-dimensional. To compute its decomposition into irreps, we utilize the orthogonality theorem [118],

$$n_I = \frac{1}{24} \sum_{g \in T_d} \chi_I(g)^* \chi_{\mathbf{4}}(g), \tag{2.85}$$

where n_I is the number of times irrep I appears in the decomposition, $\chi_I(g)$ is the character of g in irrep I, which are given in the table in Fig. 2.13(c), and $\chi_{\mathbf{4}}(g)$ is the character of g in the representation $\mathbf{4}$, which is the trace of the corresponding matrix. The result is

$$\mathbf{4} = \mathbf{1} \oplus \mathbf{3}. \tag{2.86}$$

This result could have been anticipated by the fact that every permutation representation of a symmetric group always contains the trivial representation in its irrep decomposition.

In this case the trivial representation corresponds to the subspace spanned by $\sum_{\mu=1}^{4} |\mu\rangle$. We already saw in Section 2.4.2 that for a single frustrated cluster this is always the excited eigenvector of the interaction matrix. The remaining three-dimensional subspace is orthogonal to this vector and corresponds to the ground state for a single frustrated cluster. To obtain a basis for the remaining three dimensional subspace, we will use a neat geometric observation. Let $\hat{\boldsymbol{z}}_\mu$ be the four unit vectors pointing from the center of a tetrahedron to its corners. Under the action of T_d, these four vectors are permuted, but this is permutation is equivalent to a three-dimensional rotation and/or reflection of the vectors $\hat{\boldsymbol{z}}_\mu$. Therefore the $\hat{\boldsymbol{z}}_\mu$ transform in the $\mathbf{3}$ irreducible representation under the action of T_d, meaning that the vectors $\sum_\mu z_\mu^\alpha |\mu\rangle$ ($\alpha = x, y, z$) span the space transforming in the $\mathbf{3}$ irrep, where we define $\boldsymbol{z}_\mu = \sqrt{3}\hat{\boldsymbol{z}}_\mu$. These vectors determine the $\boldsymbol{q} = \mathbf{0}$ normal modes by Eq. (2.72), which reads

$$\begin{pmatrix} S_{\mathbf{1}} \\ S_{\mathbf{3},x} \\ S_{\mathbf{3},y} \\ S_{\mathbf{3},z} \end{pmatrix} = \underbrace{\frac{1}{2} \begin{pmatrix} 1 & 1 & 1 & 1 \\ {z_1}^x & {z_2}^x & {z_3}^x & {z_4}^x \\ {z_1}^y & {z_2}^y & {z_3}^y & {z_4}^y \\ {z_1}^z & {z_2}^z & {z_3}^z & {z_4}^z \end{pmatrix}}_{\mathcal{T}_{(I,i_I)}{}^\mu} \begin{pmatrix} S_1 \\ S_2 \\ S_3 \\ S_3 \end{pmatrix}. \tag{2.87}$$

We will give these modes the more suggestive names $Q \equiv S_{\mathbf{1}}$, $B^x \equiv S_{\mathbf{3},1}$, $B^y \equiv S_{\mathbf{3},2}$, $B^z \equiv S_{\mathbf{3},3}$, where $\boldsymbol{B} = \sum_\mu S_\mu \boldsymbol{z}_\mu$ is a three-component vector. Physically, since the magnetic moments of the

ions are given $\boldsymbol{m}_\mu = \mu_B g_\parallel S_\mu \hat{\boldsymbol{z}}_\mu$, we can identify $(\mu_B g_\parallel/\sqrt{3})\boldsymbol{B}$ as the net magnetic moment of the tetrahedron. Q on the other hand is maximum when all of the $S_\mu = 1$ (or -1), and we will eventually see can be interpreted as the "charge" of the tetrahedron. Following the procedure, we construct the long-wavelength ansatz modes according to Eq. (2.74),

$$Q(\boldsymbol{q}) \equiv \frac{1}{2}\sum_\mu S_\mu(\boldsymbol{q}), \tag{2.88}$$

$$B^\alpha(\boldsymbol{q}) \equiv \frac{1}{2}\sum_\mu S_\mu(\boldsymbol{q}) z_\mu{}^\alpha. \tag{2.89}$$

Performing a long-wavelength expansion of the Hamiltonian to quadratic order in $a\boldsymbol{q}$, where a is the pyrochlore nearest-neighbor distance, and applying the unitary transformation to the irreducible basis (see Refs. [76, 81] for details), we obtain the long-wavelength Hamiltonian[16]

$$H = E_0 + \frac{J}{2}a^2 \begin{pmatrix} Q(\boldsymbol{q}) \\ B^x(\boldsymbol{q}) \\ B^y(\boldsymbol{q}) \\ B^z(\boldsymbol{q}) \end{pmatrix}^\dagger \begin{pmatrix} 8a^{-2} - |\boldsymbol{q}|^2 & -q_y q_z & -q_x q_z & -q_x q_y \\ -q_y q_z & q_x^2 & q_x q_y & q_x q_z \\ -q_x q_z & q_x q_y & q_y^2 & q_y q_z \\ -q_x q_y & q_x q_z & q_y q_x & q_z^2 \end{pmatrix} \begin{pmatrix} Q(\boldsymbol{q}) \\ B^x(\boldsymbol{q}) \\ B^y(\boldsymbol{q}) \\ B^z(\boldsymbol{q}) \end{pmatrix} + O\Big(|a\boldsymbol{q}|^4\Big). \tag{2.90}$$

We see that the $|Q(\boldsymbol{q})|^2$ term (top left matrix element) is gapped, corresponding to the maximum of the upper band in Fig. 2.12(b,c) (yellow), plus a portion which quadratically disperses downward away from $\boldsymbol{q} = \boldsymbol{0}$. Due to this gap this mode can be safely integrated out at low temperatures, leaving only the lower three bands, described by the three components of $\boldsymbol{B}(\boldsymbol{q})$. Doing so means we ignore all matrix elements in the first row and column, focusing only on the 3×3 matrix coupling the components of $\boldsymbol{B}(\boldsymbol{q})$, whose eigenvalues must yield the lower three bands, one of which is flat and one of which quadratically disperses upwards. Indeed, we can immediately recognize this matrix as $\boldsymbol{q}\boldsymbol{q}^T$, which is the projector along $\boldsymbol{q}$. We therefore obtain the following effective long-wavelength SCGA Hamiltonian,

$$\beta H_{\text{SCGA}} = \frac{1}{2}\int d^3\boldsymbol{q}\, (\lambda_0 |\boldsymbol{B}(\boldsymbol{q})|^2 + \xi^2 |\boldsymbol{q}\cdot\boldsymbol{B}(\boldsymbol{q})|^2), \tag{2.91}$$

where $\xi = a\sqrt{\beta J}$ is a lengthscale which diverges at zero temperature. We learn two things from this: first, λ_0 appears here as a "mass" for the field $\boldsymbol{B}$, controlling the Gaussian fluctuations. It is often referred to as the "flux stiffness" for this reason. Second, the *divergence* of the $\boldsymbol{B}$, $\boldsymbol{q}\cdot\boldsymbol{B}(\boldsymbol{q})$, is quadratically suppressed, and this corresponds to the quadratically dispersing band at the threefold band touching point in Fig. 2.12. In direct space this Hamiltonian describes a coarse grained

[16] We use the conventions for the $\hat{z}_\mu$ given in Appendix A.

field theory of a three-component vector field, which at zero temperature obeys an emergent Gauss' law $\nabla \cdot \boldsymbol{B} = 0$. At finite temperature, violations of the Gauss' law constraint are energetic excitations, corresponding to the $Q \neq 0$ on a tetrahedron. These are the so-called "monopoles" of spin ice [119].

Notice that for a single frustrated tetrahedron there is one excitation mode, Q, and three ground state modes $\boldsymbol{B}$, corresponding to the degeneracies at $\boldsymbol{q} = \boldsymbol{0}$ in the band structure of the interaction matrix. When these tetrahedra are coupled together into a corner-sharing lattice, one of the three ground state modes becomes dispersive, specifically the component along $\boldsymbol{q}$. The origin of this dispersion is the sharing of a single spin between neighboring tetrahedra and therefore the spatial coupling of $\boldsymbol{B}$ between neighboring tetrahedra. This can be formally understood as a continuity constraint between neighboring tetrahedra, which is described in more detail in Refs. [73, 120].

Much of Chapter 3 will be devoted to developing a rigorous mathematical framework through which the physics of this effective Hamiltonian can be formulated in direct space on the discrete lattice, without resorting to a long-wavelength approximation, and thereby refining this intuitive but simple effective field theory. For now, the key observation to make is that, if the single quadratically dispersing band near $\boldsymbol{q} = \boldsymbol{0}$ corresponds to the *divergence* of $\boldsymbol{B}$, $\boldsymbol{q} \cdot \boldsymbol{B}$ then the two flat bands must correspond to the *curl* of $\boldsymbol{B}$, $\boldsymbol{q} \times \boldsymbol{B}$, which is orthogonal to $\boldsymbol{q}$. Since the curl of $\boldsymbol{B}$ does not appear in the long-wavelength Hamiltonian, it is energetically free to fluctuate without energetic penalty, and this corresponds to the presence of flat bands in the interaction matrix. We will make this observation precise in Chapter 3. In Chapter 4 we will demonstrate the importance of this observation for the interpretation of neutron scattering cross sections, and in Chapter 5 we will use it to construct a complementary model to spin ice. We will also utilize the same method of long wavelength expansion and irrep decomposition for the case of $N = 3$ component spins in Chapter 6.

2.6 Correlations and Neutron Scattering: Pinch Points

We now turn to the analysis of the spin-spin correlation functions Eq. (2.8),

$$\langle S_i^\alpha S_j^\beta \rangle = \langle i, \alpha | \, \mathcal{G} \, | j, \beta \rangle \, . \tag{2.92}$$

In practice, the correlation functions are analyzed in reciprocal space. They can be visualized defining structure factors, which are scalar functions of $\boldsymbol{q}$ derived from the matrix elements of $\mathcal{G}$. There are a few natural ones we will introduce in this section. The most important of these are the neutron structure factors, which are proportional to the scattering cross sections observed in neutron scattering.

2.6.1 Structure Factors

For each irrep we defined a basis $|\boldsymbol{q}, I, i_I\rangle$ for the long wavelength expansion (Eq. (2.73)). We define a structure factor for each irrep as[17]

$$S_I(\boldsymbol{q}) \equiv \sum_{i_I} \langle \boldsymbol{q}, I, i_I | \, \mathcal{G} \, | \boldsymbol{q}, I, i_I \rangle = \left\langle \left| S_{\boldsymbol{q},I,i_I} \right|^2 \right\rangle . \tag{2.93}$$

For spin ice, this means the correlation functions

$$S_Q(\boldsymbol{q}) \equiv \langle |Q(\boldsymbol{q})|^2 \rangle = \frac{1}{4}\frac{1}{L^d} \sum_{ij} \langle S_i^z S_j^z \rangle e^{-i\boldsymbol{q}\cdot(\boldsymbol{r}_j - \boldsymbol{r}_i)}, \tag{2.94}$$

$$S_{\boldsymbol{B}}(\boldsymbol{q}) \equiv \langle |\boldsymbol{B}(\boldsymbol{q})|^2 \rangle = \frac{3}{4}\frac{1}{L^d} \sum_{ij} \langle S_i^z S_j^z \rangle (\hat{\boldsymbol{z}}_i \cdot \hat{\boldsymbol{z}}_j) \, e^{-i\boldsymbol{q}\cdot(\boldsymbol{r}_j - \boldsymbol{r}_i)}. \tag{2.95}$$

Since spins carry a magnetic moment $\boldsymbol{m}_i$ it is natural to consider the moment-moment correlators $\langle m_i^\alpha m_j^\beta \rangle$ (for $N = 1, 2, 3$). The spin vector $\boldsymbol{S}_i$ is related to the magnetic moment by the so-called g-tensor,

$$m_i^\alpha = \sum_\beta g_i{}^{\alpha\beta} S_i^\beta, \tag{2.96}$$

which is generically of the form [28, 29, 121]

$$g_i = g_\parallel \hat{\boldsymbol{z}}_i \hat{\boldsymbol{z}}_i^T + g_\perp (\mathbb{1}_{3\times 3} - \hat{\boldsymbol{z}}_i \hat{\boldsymbol{z}}_i^T), \tag{2.97}$$

where $g_\parallel$ and $g_\perp$ are usually determined experimentally. We then define the generalized susceptibility tensor

$$\tilde{\chi}^{\alpha\beta}(\boldsymbol{q}) \equiv \sum_{\mu\nu} \sum_{\gamma\delta} \langle \boldsymbol{q}, \mu, \gamma | \, \mathcal{G} \, | \boldsymbol{q}, \nu, \delta \rangle \, g_\mu{}^{\alpha\gamma} g_\nu{}^{\beta\delta} = \frac{1}{L^d} \sum_{ij} \langle m_i^\alpha m_j^\beta \rangle e^{-i\boldsymbol{q}\cdot(\boldsymbol{r}_j - \boldsymbol{r}_i)}. \tag{2.98}$$

In particular for Ising spins where the moment is along a local easy axis $\boldsymbol{z}_i$, this simplifies to $(g_\parallel)^2 \sum_{\mu\nu} \langle \boldsymbol{q}, \mu | \, \mathcal{G} \, | \boldsymbol{q}, \nu \rangle \, \hat{z}_\mu^\alpha \hat{z}_\nu^\beta$. Using the generalized susceptibility tensor, we can define the moment-moment structure factor,

$$S_{\boldsymbol{m}}(\boldsymbol{q}) \equiv \mathrm{Tr}\, \tilde{\chi}(\boldsymbol{q}) = \frac{1}{L^d} \sum_{ij} \langle \boldsymbol{m}_i \cdot \boldsymbol{m}_j \rangle \, e^{-i\boldsymbol{q}\cdot(\boldsymbol{r}_j - \boldsymbol{r}_i)}. \tag{2.99}$$

[17] Note we can think of the correlation matrix in a basis-free way as $\big\langle \, |S\rangle\langle S| \, \big\rangle$, where the outer angle brackets indicate thermal average.

For spin ice, this is proportional to Eq. (2.95) since $\boldsymbol{B}$ is proportional to the net moment of a tetrahedron. It may also be of interest to study the pure spin-spin correlators, defining the so-called spin structure factor

$$S(\boldsymbol{q}) \equiv \sum_{\mu\nu}\sum_{\alpha} \langle \boldsymbol{q}, \mu, \alpha | \, \mathcal{G} \, | \boldsymbol{q}, \nu, \alpha \rangle = \frac{1}{L^d} \sum_{ij} \langle \boldsymbol{S}_i \cdot \boldsymbol{S}_j \rangle \, e^{-i\boldsymbol{q}\cdot(\boldsymbol{r}_j - \boldsymbol{r}_i)}. \tag{2.100}$$

2.6.2 Neutron Scattering

The susceptibility tensor is also directly related to neutron scattering cross sections. For a complete derivation of the scattering cross section see Ref. [73], we will only summarize the salient details. By Fermi's golden rule, the probability for a system to transition from initial state $|i\rangle$ to final state $|f\rangle$ given a Hamiltonian V is proportional to $|\langle f | V | i \rangle|^2$. For neutron scattering the initial state is a neutron with momentum $\boldsymbol{k}$ and spin σ and our bulk system in a state $|\psi\rangle$, the final state is a neutron with momentum $\boldsymbol{k}'$ and spin σ' with the bulk system in a state $|\psi'\rangle$.[18] The interaction potential is $V = -\boldsymbol{\mu}_n \cdot \boldsymbol{B}$, where $\boldsymbol{\mu}_n$ is the neutron magnetic moment and $\boldsymbol{B}$ is the bulk magnetic field. Assuming that the magnetic moments of the ions is well localized, the bulk field is that of an array of magnetic dipoles,

$$\boldsymbol{B}(\boldsymbol{r}) = \sum_i \frac{\mu_0}{4\pi} \frac{3(\boldsymbol{m}_i \cdot \hat{\boldsymbol{R}}_i)\hat{\boldsymbol{R}}_i - \boldsymbol{m}_i}{R_i^3}, \tag{2.101}$$

where $\boldsymbol{r}$ is the field point, $\boldsymbol{r}_i$ is the position of the i'th ion, and $\boldsymbol{R}_i = \boldsymbol{r} - \boldsymbol{r}_i$. Due to the crystal translation invariance of the system the scattering intensity is only a function of the scattering momentum $\boldsymbol{q} = \boldsymbol{k}' - \boldsymbol{k}$. The magnetic field $\boldsymbol{B}$ has zero divergence, so when we Fourier transform the potential $V(\boldsymbol{r})$, one obtains

$$V(\boldsymbol{q}) \propto \sum_{ab} \mu_n^a (\delta^{ab} - \hat{q}^a \hat{q}^b) m^b(\boldsymbol{q}). \tag{2.102}$$

The projector in the middle simply projects out the component of $\boldsymbol{m}(\boldsymbol{q})$ along $\boldsymbol{q}$, which we also saw for the emergent zero-divergence field describing the long wavelength limit of spin ice in Eqs. (2.110) and (2.111). Thus, applying Fermi's golden rule, after thermally averaging over intermediate states of the bulk system and summing over the neutron spin states, the (equal time)

[18] The magnetic field $\boldsymbol{B}$ should not be confused with the irrep normal mode for spin ice, which is denoted with the same symbol. They are directly related of course, since the irrep normal mode $\boldsymbol{B}$ is proportional to the net moment of a single tetrahedron.

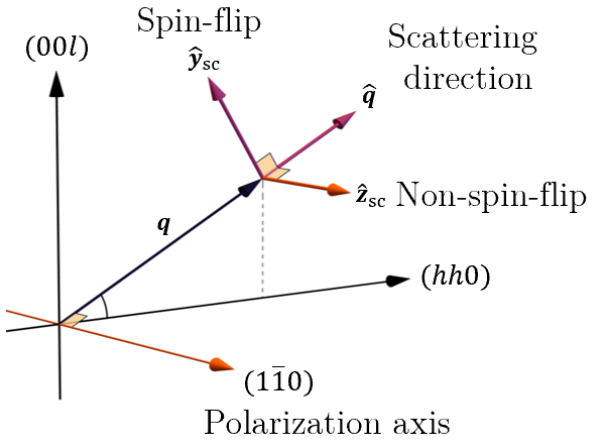


Figure 2.15: When performing polarization analysis, the neutron beam is polarized along a fixed direction $\hat{\boldsymbol{z}}_{\rm sc}$. For each scattering wavevector $\boldsymbol{q}$ in the plane orthogonal to $\hat{\boldsymbol{z}}_{\rm sc}$ (here spanned by $(hh0)$ and $(00l)$), there is a natural orthonormal basis spanned by $\hat{\boldsymbol{z}}_{\rm sc}$, $\hat{\boldsymbol{x}}_{\rm sc} \equiv \hat{\boldsymbol{q}}$, and $\hat{\boldsymbol{y}}_{\rm sc} \equiv \hat{\boldsymbol{z}}_{\rm sc} \times \hat{\boldsymbol{x}}_{\rm sc}$. The neutron is insensitive to the components of the magnetic moments along $\boldsymbol{q}$. Components of the magnetic moments along $\hat{\boldsymbol{y}}_{\rm sc}$, orthogonal to the neutron polarization, can induce spin flipping, which components along $\hat{\boldsymbol{z}}_{\rm sc}$ do not induce spin flipping.

scattering intensity at momentum $\boldsymbol{q}$ is proportional to the following structure factor,

$$S_{\rm n}(\boldsymbol{q}) \equiv \sum_{\mu\nu}\sum_{ab}(\delta^{ab} - \hat{q}^a\hat{q}^b)\langle m^a_\mu(-\boldsymbol{q})m^b_\nu(\boldsymbol{q})\rangle. \tag{2.103}$$

There is also the possibility of performing a polarization analysis of the neutron scattering cross section, resolving neutrons which flip their spin from those which do not by polarizing the incoming beam [121, 122]. The initial beam is polarized along a fixed polarization axis $\hat{\boldsymbol{z}}_{\rm sc}$. One then analyzes neutrons for scattering wavevectors $\boldsymbol{q} \perp \hat{\boldsymbol{z}}_{\rm sc}$. For each $\boldsymbol{q}$ in this plane, there is a natural orthonormal basis spanned by $\hat{\boldsymbol{z}}_{\rm sc}$, $\hat{\boldsymbol{x}}_{\rm sc} \equiv \hat{\boldsymbol{q}}$, and $\hat{\boldsymbol{y}}_{\rm sc} \equiv \hat{\boldsymbol{z}}_{\rm sc} \times \hat{\boldsymbol{x}}_{\rm sc}$, see Fig. 2.15 for an illustration. The components of the local magnetic field orthogonal to the neutron moment, i.e. along $\hat{\boldsymbol{x}}_{\rm sc}$ and $\hat{\boldsymbol{y}}_{\rm sc}$, can induce the neutron spin to flip, while the component along $\hat{\boldsymbol{z}}_{\rm sc}$ cannot. However, the neutron is entirely insensitive to the component along $\hat{\boldsymbol{x}}_{\rm sc}$ since the magnetic field has zero divergence.[19] Therefore the polarized neutron scattering cross sections will be proportional

[19]Since there are no true magnetic monopoles (that we know of), the physical magnetic field $\boldsymbol{B}$ which the neutron interacts with has zero divergence, $\nabla \cdot \boldsymbol{B} = 0$, or in reciprocal space $\boldsymbol{q} \cdot \boldsymbol{B}(\boldsymbol{q}) = 0$. Therefore the neutron is "blind" to the longitudinal (along $\boldsymbol{q}$) moment-moment correlations, which need not be zero.

to the following "spin flip" (SF) and "non-spin flip" (NSF) structure factors,

$$S_{\mathrm{SF}}(\boldsymbol{q}) \equiv \sum_{\mu\nu} \langle (\boldsymbol{m}_\mu(-\boldsymbol{q}) \cdot \hat{\boldsymbol{y}}_{\mathrm{sc}})(\boldsymbol{m}_\nu(\boldsymbol{q}) \cdot \hat{\boldsymbol{y}}_{\mathrm{sc}}) \rangle \equiv \left\langle \left| \sum_\mu \boldsymbol{m}_\mu(\boldsymbol{q}) \cdot \hat{\boldsymbol{y}}_{\mathrm{sc}} \right|^2 \right\rangle, \tag{2.104}$$

$$S_{\mathrm{NSF}}(\boldsymbol{q}) \equiv \sum_{\mu\nu} \langle (\boldsymbol{m}_\mu(-\boldsymbol{q}) \cdot \hat{\boldsymbol{z}}_{\mathrm{sc}})(\boldsymbol{m}_\nu(\boldsymbol{q}) \cdot \hat{\boldsymbol{z}}_{\mathrm{sc}}) \rangle \equiv \left\langle \left| \sum_\mu \boldsymbol{m}_\mu(\boldsymbol{q}) \cdot \hat{\boldsymbol{z}}_{\mathrm{sc}} \right|^2 \right\rangle. \tag{2.105}$$

The unpolarized scattering cross section is then the sum of the SF and NSF cross sections,

$$S_{\mathrm{n}}(\boldsymbol{q}) = S_{\mathrm{SF}}(\boldsymbol{q}) + S_{\mathrm{NSF}}(\boldsymbol{q}). \tag{2.106}$$

For completeness, it is worthwhile to also define a structure factor for the correlations along $\boldsymbol{q}$, which are generally not observable via standard neutron scattering,

$$S_{\boldsymbol{q}}(\boldsymbol{q}) \equiv \sum_{\mu\nu} \langle (\boldsymbol{m}_\mu(-\boldsymbol{q}) \cdot \hat{\boldsymbol{x}}_{\mathrm{sc}})(\boldsymbol{m}_\nu(\boldsymbol{q}) \cdot \hat{\boldsymbol{x}}_{\mathrm{sc}}) \rangle \equiv \left\langle \left| \sum_\mu \boldsymbol{m}_\mu(\boldsymbol{q}) \cdot \hat{\boldsymbol{x}}_{\mathrm{sc}} \right|^2 \right\rangle. \tag{2.107}$$

We also have that

$$S_{\boldsymbol{m}}(\boldsymbol{q}) = S_{\mathrm{SF}}(\boldsymbol{q}) + S_{\mathrm{NSF}}(\boldsymbol{q}) + S_{\boldsymbol{q}}(\boldsymbol{q}). \tag{2.108}$$

We note that, by performing polarization analysis for scattering in different Brillouin zones away from $\boldsymbol{q} = \boldsymbol{0}$, it is in fact possible, at least in some cases and with considerable effort, to reconstruct all components of $\tilde{\chi}^{\alpha\beta}$. For example, Ref. [123] considers a two-dimensional square lattice system; since $\chi^{\alpha\beta}(\boldsymbol{q})$ is periodic in reciprocal space, they can probe $\chi^{\alpha\beta}((100) + \boldsymbol{q})$, for which the components along $\hat{\boldsymbol{x}}$ are inaccessible but the components along $\boldsymbol{q}\hat{y}$ are accessible, then probe $\chi^{\alpha\beta}((010) + \boldsymbol{q})$, for which the component along $\hat{\boldsymbol{y}}$ are inaccessible but the components along $\hat{\boldsymbol{x}}$ are accessible. In this way all components of $\chi^{\alpha\beta}(\boldsymbol{q})$ can in principle be reconstructed.

2.6.3 Pinch Points in Spin Ice

Consider the low-temperature behavior of the correlation matrix for spin ice via its spectral decomposition:

$$\mathcal{G} = \sum_{\boldsymbol{q},n} \frac{1}{\lambda_0 + \beta(\epsilon_n(\boldsymbol{q}) - \epsilon_0)} \left|\psi_{\boldsymbol{q},n}\right\rangle \left\langle\psi_{\boldsymbol{q},n}\right| \to \begin{cases} \mathbb{1} & T \to \infty, \\ \sum\limits_{\substack{\text{flat} \\ \text{bands}}} 2 \left|\psi_{\boldsymbol{q},n}\right\rangle \left\langle\psi_{\boldsymbol{q},n}\right| & T \to 0. \end{cases} \tag{2.109}$$

The factor of 2 comes from $\lambda \to 1/2$ at low temperature, corresponding to the depopulation of the two dispersive bands to maintain the self-consistency condition $\mathrm{Tr}\,\mathcal{G} = N_{\mathrm{spins}}$, Eq. (2.10). We say that the spectral weight moves from the dispersive bands to the flat bands, leading the flat band mode correlators to double in intensity. The most important observation about Eq. (2.109) is that *the zero temperature correlation matrix is the projector to the flat band subspace times* λ_0^{-1}. In the high temperature paramagnetic phase the correlation matrix is just the identity matrix among all modes, while in the limit of zero temperature the correlation matrix becomes the proportional to the identity in the flat band subspace and zero in the excited mode subspace. Thus we say that the low temperature phase is a "cooperative paramagnet" [55].

The effective field theory of Eq. (2.91) allows us to directly read off the long-wavelength correlation functions for the irrep modes $\boldsymbol{B}(\boldsymbol{q})$, by writing the integrand as a quadratic form using Eq. (2.90),

$$\beta H_{\mathrm{eff}} = \frac{1}{2}\int \mathrm{d}^3\boldsymbol{q} \sum_{\alpha\beta} B^\alpha(-\boldsymbol{q})\lambda_0 \left[\delta^{\alpha\beta} + \frac{\xi^2}{\lambda_0}\frac{q^\alpha q^\beta}{|\boldsymbol{q}|^2}\right] B^\beta(\boldsymbol{q}), \tag{2.110}$$

such that the two-point correlators for $\boldsymbol{B}$ are given by the inverse of the kernel,

$$\langle B^\alpha(-\boldsymbol{q})B^\beta(\boldsymbol{q})\rangle \approx \frac{1}{\lambda_0}\left(\delta^{\alpha\beta} - \frac{q^\alpha q^\beta}{|\boldsymbol{q}|^2 + \lambda_0\xi^{-2}}\right) \to \begin{cases} \delta^{\alpha\beta} & T \to \infty, \\ 2(\delta^{\alpha\beta} - \hat{q}^\alpha\hat{q}^\beta) & T \to 0. \end{cases} \tag{2.111}$$

Here we utilized the fact that at high temperature $\lambda_0 \to 1$ and $\xi \to 0$, and at low temperature $\lambda_0 \to 1/2$ and $\xi \to \infty$. In the high temperature limit correlations are isotropic and $\boldsymbol{q}$-independent, while at zero temperature the correlations are anisotropic. In Eq. (2.110) (or Eq. (2.91)) we can view the isotropic stiffness term as a Gaussian background, while the second term yields the constraint that $\nabla \cdot \boldsymbol{B} = 0$ when $\xi \to \infty$. If we Fourier transform the zero-temperature limit of Eq. (2.111) to direct space, the correlation function looks like that of a dipole Eq. (2.101), decaying as a power law. For this reason the low temperature phase is referred to as a Coulomb phase [55].

The low temperature correlations are highly anisotropic due to their dipolar nature, and exhibit characteristic anisotropic singularities called pinch points which reflect the constraint $\boldsymbol{q} \cdot \boldsymbol{B}(\boldsymbol{q}) = 0$ near the zone center. In Fig. 2.16, we show various structure factors for nearest-neighbor spin ice computed in the SCGA, at an intermediate and at low temperature. For each temperature we are showing two dimensional cuts in the high symmetry $(hk0)$ and (hhl) planes, with FCC Brillouin zone outlines indicated by thin dashed lines, compare to the two planes in Fig. 2.12(a).

At low temperature sharp pinch points form indicating the zero-divergence constraint. At finite temperature the pinch points broaden, indicating violations of the constraint. At low temperature, regions of diffuse intensity indicate paramagnetic correlations of the modes forming the flat bands, while regions of zero intensity indicate the suppression of modes forming the dispersive bands.

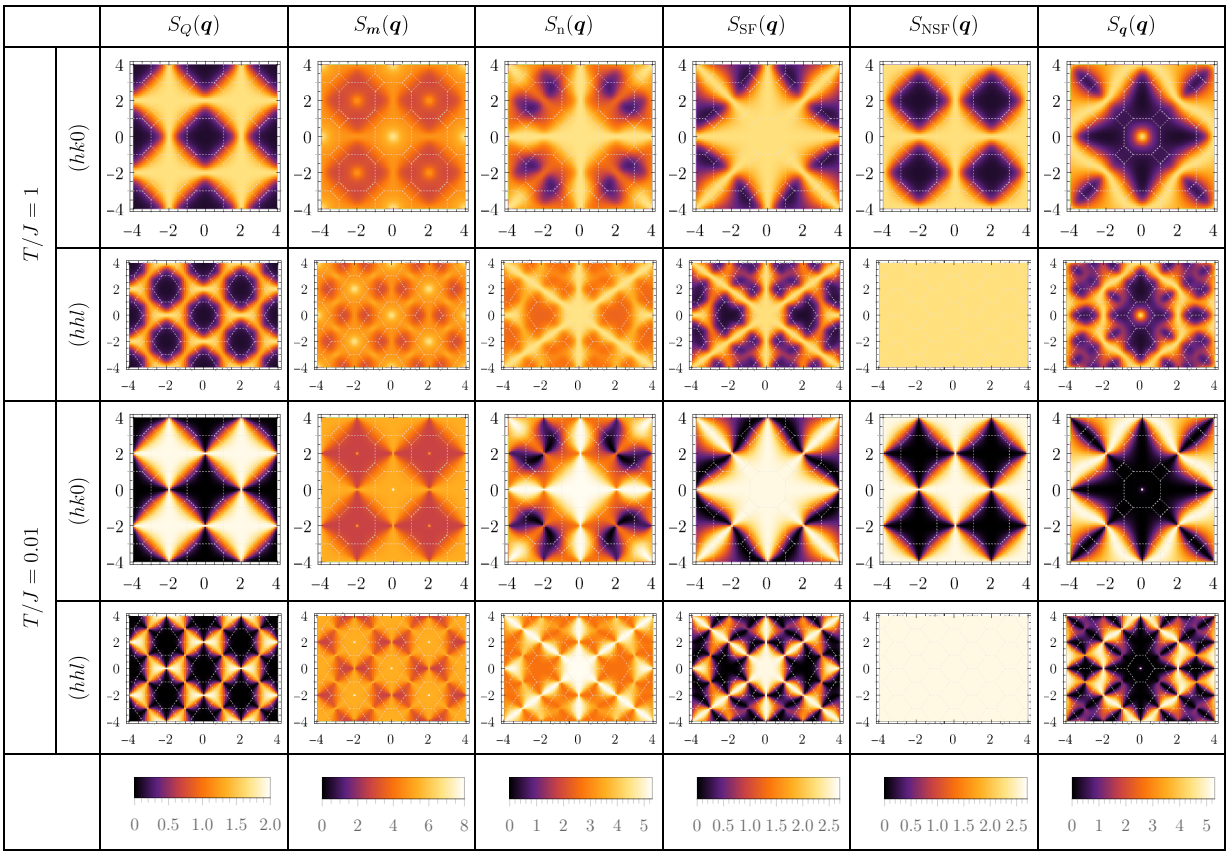
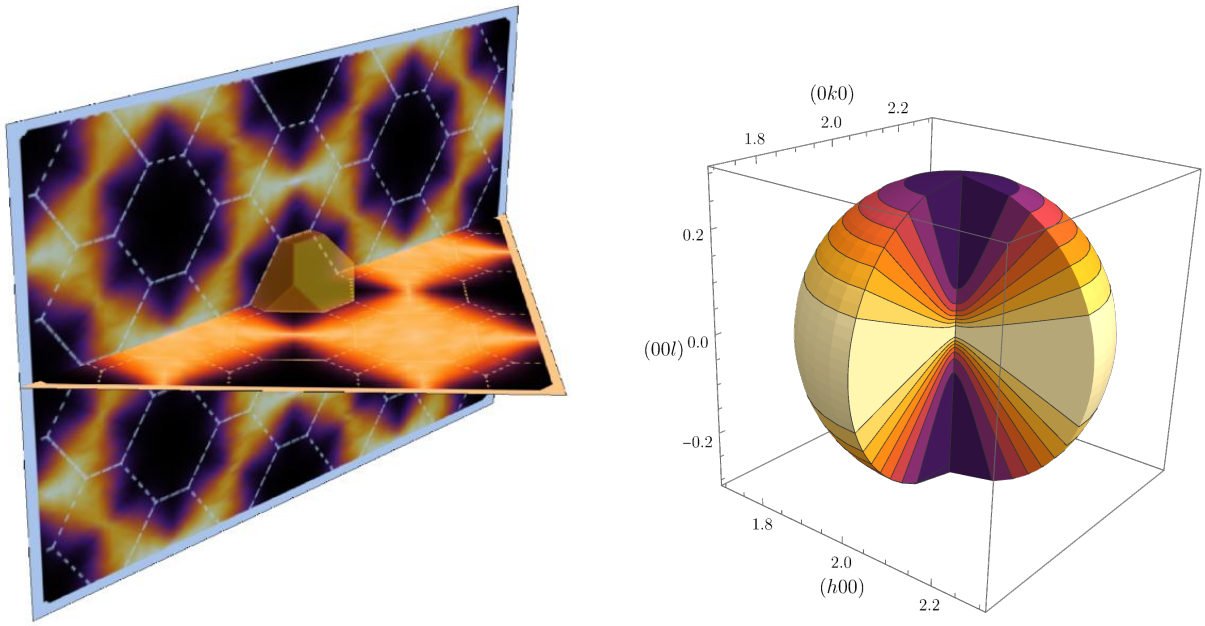


Figure 2.16: Structure factors for nearest-neighbor spin ice at high and low temperature, showing the development of singular pinch points features reflecting singularities in the spin-spin correlation matrix.



(a) (b)

Figure 2.17: In attempting to visual three dimensional correlations with two dimensional cuts it is important to not lose track of the three dimensional structure. (a) The $S_Q(\boldsymbol{q})$ structure factor (first column of Fig. 2.16) showing the intersection of the $(hk0)$ and (hhl) planes and their relation to the FBZ. (b) The pinch point in $S_Q(\boldsymbol{q})$ centered at (220), showing that the intensity is pinched along the $(00l)$ (vertical) direction and remains along the two orthogonal directions.

Regions of intermediate intensity probe a mixture of dispersive and flat band modes. Note in particular the NSF structure factor in the (hhl) plane, which shows no dispersion, this will be the topic of Chapter 4. In Fig. 2.17(a), we show the intersection of the $(hk0)$ and (hhl) planes for the $S_Q(\boldsymbol{q})$ structure factor to demonstrate the three dimensional nature of the correlations. In Fig. 2.17(b) we show the three dimensional intensity of the pinch point located at the (220) zone center, which is pinched along the $(00l)$ direction. This is because there is a zone center at (222) and another at $(22\bar{2})$. Relative to these two zone centers, (220) is along the longitudinal $\hat{\boldsymbol{q}}$ direction.

2.7 Classical Monte Carlo

The SCGA soft spin approach is extremely useful, as discussed throughout this chapter, but it is a significant approximation which does not always accurately reflect the hard spin physics (in particular for the triangular lattice Ising antiferromagnet). It is generally a strong starting point for a basic analysis of a model, but can only take us so far. To obtain information about the hard spin model, we turn to numerical Markov Chain Monte Carlo methods. We will not describe these in detail since they are standard methods. For a general Ising model, we utilize either single spin flip updates or cluster flip updates with the Metropolis acceptance condition, where the probability of accepting a proposed update is given by $\min(1, \exp(-\beta\Delta E))$, where ΔE is the change in energy of the system before and after the flip. For spin ice we supplement single spin flips with loop moves [124–127], which flips a cluster of head-to-tail spins. Such a cluster is generated by a random walk starting at a random spin, until the random walk intersects itself, at which point the "starting tail" of the walk is discarded and the remaining closed loop is flipped. For nearest-neighbor spin ice we utilize the extremely efficient cluster algorithm introduced in Ref. [128], which naturally breaks the spin configurations into clusters of closed loops and open strings of spins oriented head-to-tail, which can then be flipped with 50/50 probability. In Chapter 6 we consider classical Heisenberg spins, in which case we utilize Gaussian updates [129], where we generate a random vector $\boldsymbol{\nu}$ whose three Cartesian components are Gaussian distributed random numbers with fixed variance, and we update a single spin by $\boldsymbol{S}_i \to \boldsymbol{S}_i + \epsilon\boldsymbol{\nu}$, then normalize to unit length, again accepted according to the Metropolis condition. The value of ϵ is then updated to maintain an approximately 50% acceptance rate for proposed updates under the Metropolis condition.

Chapter 3

Hodge Theory for Geometrically Frustrated Magnetism

In the previous chapter we discussed frustration on line graph lattices, all of which have at least one flat band at the base of the spectrum of their adjacency matrix $\mathcal{A}$ with a band touching point at the zone center from which another band disperses upwards. We saw that the emergent long-wavelength field theory for such a model is described by a vector field obeying a zero-divergence constraint, which is the coarse-grained consequence of the microscopic lattice constraint that $\sum_{i\in c} S_i = 0$ on each cluster. In this chapter our goal is to formulate the zero-divergence field theory at the *lattice* level, without resorting to the long-wavelength approximation. We will see that all such models can be reformulated as a theory of discrete 1-forms on the parent lattice. We will mainly focus on the spin ice model on the three-dimensional pyrochlore lattice, though our results are applicable to any line graph lattice, and the formalism we develop based on Hodge theory describes any such model.

As discussed previously, the pyrochlore lattice is the line graph of the diamond lattice—pyrochlore sites sit at the midpoint of the diamond lattice edges, and the monopole excitations of spin ice sit on the vertices of the diamond lattice. These facts are essential to the mapping of quantum spin ice to a lattice gauge theory [25], where the local z-components of spins map to electric fields on the edges of the diamond lattice, and to the construction of the gauge mean field theory [130], where the spinons (monopoles) are added as slave bosons on the diamond lattice vertices. Our goal is therefore to develop this lattice field theory picture in a way that allows us to recast classical geometrically frustrated models as 1-form fields, which will allow us to develop various tools to understand their physics more clearly, in particular the Helmholtz-Hodge decomposition, which will underpin our understanding of these systems.

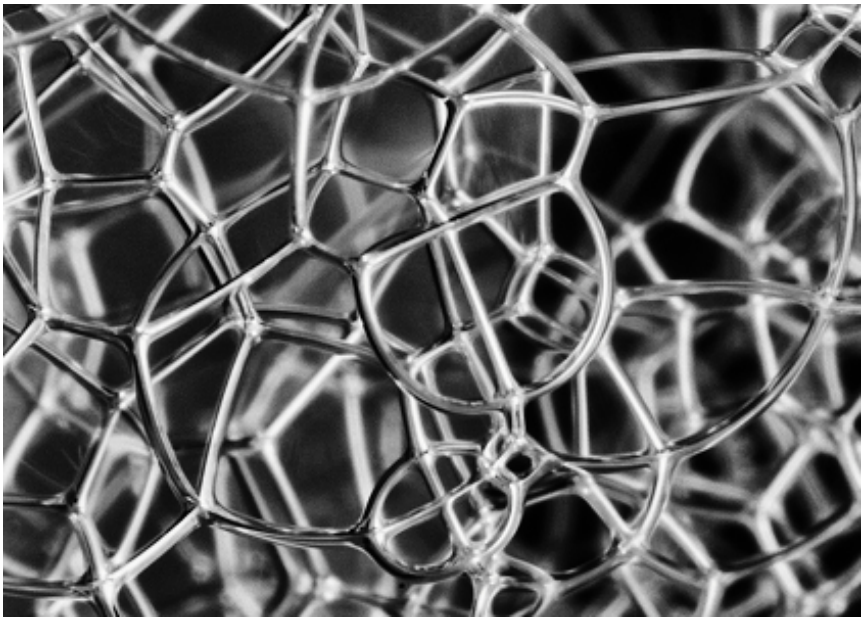


Figure 3.1: A foam of soap bubbles forms a natural 3-dimensional cellular complex. Photograph by Paul VanDerWerf, distributed under a CC-BY 2.0 license.

3.1 Hodge Theory: Discrete Differential Forms

A differential k-form is a function which can be integrated over a k-dimensional domain to obtain a number. When studying spin models, we generically are interested in discrete degrees of freedom living on a rigid lattice, rather than a continuous manifold. A continuum field is a fluctuating k-form, so we would like the discrete analogue of differential forms.

3.1.1 Cellular Complexes

The starting point for defining discrete forms is to define the space on which they live. In differential geometry, differential forms can be understood to take their values at each point on a manifold. Discrete differential forms should take values on the discretization of a manifold, which in algebraic topology is called a cellular complex. For our purposes, a cellular complex is a discretization of flat 3-dimensional Euclidean space into 3-dimensional cells, which are separated by 2-dimensional surfaces. The surfaces meet along 1-dimensional edges, and these terminate at 0-dimensional vertices. An intuitive and physically relevant example of a cellular complex is a foam of soap bubbles, see Fig. 3.1. Each bubble is a 3-dimensional cell, and the soap films

form 2-dimensional surfaces. We call these minimal k-dimensional domains k-cells. We will not consider a k-cell to have any additional microscopic structure, e.g. we do not consider our 2-cells (such as the soap film walls) to have a curvature that varies smoothly across its surface. They are the smallest building blocks for a cellular complex.

Mathematically, a cellular complex as we have described is a special case of the more general and abstract construction of a "CW complex". We will very briefly describe the concept here, in order to be precise about what data is required to specify a cellular complex (e.g. in a computer program). The basic idea is to start by specifying the vertices, then to "glue in" the edges between pairs of vertices, then to glue in the faces, etc. The gluing is data that we need to specify to define the complex. Most importantly, it is data that we need to define the boundary relation between k-cells and $(k-1)$-cells, which will play a central role in the definition of discrete differential operators. A k-cell is always homeomorphic to a k-disk, whose boundary is a $(k-1)$-sphere.[1] Denote the collection of vertices $\mathcal{K}_0$, which we call the 0-skeleton. The 1-cells (which we will also call edges or bonds) are glued to the 0-skeleton by a collection of maps: for each $\tau_1 \in \mathcal{K}_1$, there is a map (a piece of data) which glues the two endpoints of the 1-cell to two vertices in the 0-skeleton. Once all of the 1-cells are glued in, we have formed the 1-skeleton, the collection of all vertices and edges glued together, which we denote $\mathcal{K}_1$.[2] Next we glue in the 2-cells (which we will also call faces or plaquettes). Each 2-cell is homeomorphic to a disk, and we have a gluing map which attaches its boundary, homeomorphic to a circle, to the edges in the 1-skeleton. Once that is done, we have formed the 2-skeleton (all vertices, edges, and faces glued together), denoted $\mathcal{K}_2$. Similarly for the 3-skeleton.[3] Once complete, we have formed the entire cellular complex, and we denote the collection of all the cells by $\mathcal{K}$. While we will not need to use this construction explicitly, these gluing maps are the precise data we need to keep track of relating each cell to its boundary.

3.1.2 Orientations

A k-cell can be equipped with one of two orientations. For example, a 1-cell is a 1-dimensional interval between two points, and can be equipped with two orientations. Let τ_k denote an oriented

[1] A 0-disk is a point, a 1-disk is an interval with two endpoints, a 2-disk is a regular "frisbee" disk with an edge, and a 3-disk is a solid ball with a surface. "Homeomorphic to" means for our purposes "continuously deformable to". A 0-sphere is a set of two points, a 1-sphere is a circle, and a 2-sphere is the usual sphere.

[2] Notice that the 1-skeleton of the complex is a graph.

[3] Note that at no point did we require that the cells don't intersect each other, or that we can't glue to the entire boundary of a k-cell to a single $(k-1)$-cell, or make any stronger stipulations on the structure of the complex. This is why CW complexes are extremely general and abstract, while the cellular complexes we will consider, coming from crystalline lattices, will have much more sensible properties than a fully general CW complex.

k-cell, and let $\bar{\tau}_k$ denote the opposite orientation. An orientation of the entire complex is denoted o, meaning that we have selected an orientation for every cell of the complex, and we denote the collection of all cells with the orientation given by o by $\mathcal{K}^o$.

3.1.3 Chains and Cochains

We will now introduce discrete differential forms, aka "cochains". Note that the word "differential" implies *smoothness* (i.e. differentiability), meaning that differential forms are infinitesimal continuum objects, taking a value at each point in a continuous space. Thus differential forms are not the precise objects we wish to discretize. Given a differential k-form ω, we can define a function f_ω, called a *cochain*, which takes as its argument a *macroscopic* (non-infinitesimal) k-dimensional integration domain U and returns the integral of ω over U,

$$f_\omega(U) := \int_U \omega. \tag{3.1}$$

This is the object which our discrete forms will be analogous to. Note in particular that the value of f_ω reverses sign if the orientation of U is reversed,

$$f_\omega(\bar{U}) = -f_\omega(U). \tag{3.2}$$

We now introduce the concept of a k-chain: a *formal linear combination of k-cells*.[4] Let ξ_k be a k-chain, then it can be written as

$$\xi_k = \sum_{\tau_k \in \mathcal{K}^o} c(\tau_k)\tau_k \tag{3.3}$$

where the $c(\tau_k)$ are real numbers.[5] The sum is over every k-cell τ_k with fixed orientation given by o, i.e. each k-cell appears only once and with a fixed orientation.[6] Chains serve as our discrete oriented integration domains, i.e. they represent (linear combinations of) oriented curves, surfaces, volumes, etc. For each oriented cell τ_k there is an associated "indicator chain",

$$\tau_k \equiv \cdots + 0 + 0 + 0 + 1\tau_k + 0 + 0 + 0 + \cdots, \tag{3.4}$$

[4] In more formal language, we define the k-chain group to be the free Abelian group generated by the k-cells of the complex.

[5] More generally one can consider the coefficients to be in any Abelian (commutative) group. We will also be interested in complex coefficients. In the study of singular homology it is standard to take integer coefficient [2].

[6] In the case of infinite complexes we would only consider finite linear combinations, though we will always consider finite complexes.

with coefficient +1 on that cell and zero on all others. We have used a slight abuse of notation to denote this chain τ_k, i.e. on the left τ_k is a chain and on the right τ_k is an oriented cell. Thus we say $\tau_k \equiv +1\tau_k$. We then identify the indicator chain of the same cell with reversed orientation, $\bar{\tau}_k$, with the chain that has the opposite sign,

$$\bar{\tau}_k \equiv \cdots + 0 + 0 + 0 - 1\tau_k + 0 + 0 + 0 + \cdots . \tag{3.5}$$

Thus we say $\bar{\tau}_k \equiv -1\tau_k$.

Finally, we define our discrete differential forms, or cochains. The formal definition is *a k-cochain is a linear functional k-chains*. More plainly, a k-cochain α is a function which takes a value on each oriented k-cell τ_k, and which distributes over linear combinations,

$$\alpha\left(\sum_{\tau_k \in \mathcal{K}_k^o} c_k \tau_k\right) := \sum_{\tau_k \in \mathcal{K}_k^o} c_k \, \alpha(\tau_k). \tag{3.6}$$

In particular, this means that the value of a cochain changes sign under a reversal of orientation,

$$\alpha(\bar{\tau}_k) = \alpha(-\tau_k) = -\alpha(\tau_k), \tag{3.7}$$

just like Eq. (3.2). A cochain is fully specified by its value on each oriented cell of the complex.

3.1.4 The Boundary of a Chain

In the construction of the cellular complex, we had to specify the gluing data, which tells us which $(k-1)$-cells are associated to the boundary of a k-cell. We will now define the boundary operator ∂ as a linear operator on chains utilizing this data.[7] Since it is a linear operator, we only need to specify how it acts on the oriented k-cells, given by the following formula

$$\partial \tau_k = \sum_{\tau_{k-1} \in \partial \tau_k} \tau_{k-1} \tag{3.8}$$

This formula obviously appears recursive. The formula can be read "the boundary operator acting on the indicator chain for k-cell τ_k is equal to the sum over all cells τ_{k-1} which are in the boundary of τ_k". To make sense of this, we note that the ∂ which appears below the summation symbol must be understood as the boundary relationship defined by the gluing data which is part of the definition of the cellular complex. The orientations of the τ_{k-1} appearing in the sum are implicit

[7] As always, linear means that it distributes over linear combinations.

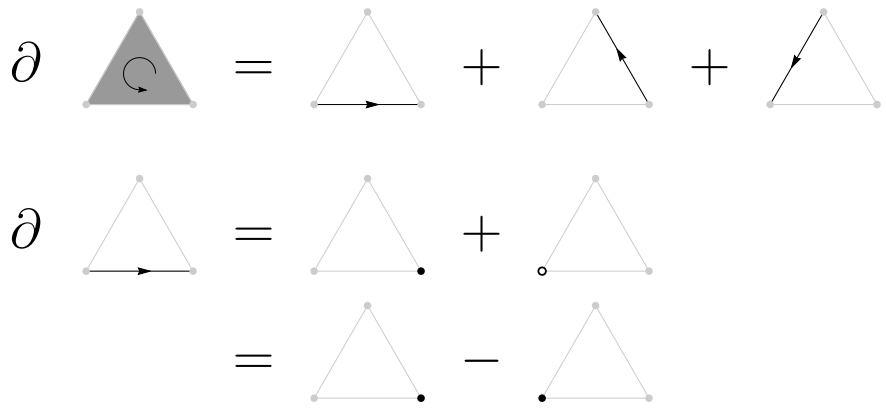


Figure 3.2: An example of the boundary operator acting on a chain. Top: the boundary of an oriented triangular 2-cell is the sum of its three 1-cell edges with orientations (indicated by arrows) determined by the orientation of the "bulk". Bottom: the boundary of an oriented 1-cell is the sum of the two endpoints with opposite orientation, where the terminating endpoint is oriented "in" (filled black circle) and the starting endpoint is oriented "out" (empty circle). Equivalently, it is the *difference* of the 0-cells with the *same* orientation, here both oriented "in". We take the convention that all 0-cells are always oriented "in" in our formulas.

from the gluing data: the boundary of τ_k inherits an orientation from the orientation of the bulk, and therefore determines the orientations of the τ_{k-1}.

This is illustrated in Fig. 3.2 for a plaquette and an edge. The boundary of an oriented plaquette is the sum of all oriented edges on its boundary with orientations determined to be compatible with the bulk orientation. Similarly, the boundary of an edge is the sum of its two endpoints with opposite orientation, where the "terminating" end is oriented "in" (black dot) and the "originating" end is oriented "out" (open dot). As a chain, a vertex oriented "out" is equivalent to the negative of the same vertex oriented "in", so we can equivalently write the boundary of an edge as the *difference* of the two endpoints, and in the last line of Fig. 3.2 we show the result if we take both ends oriented "in". *We will always take the convention that a vertex is oriented "in"*, which can be seen as a sign convention for our calculations. Such a global orientation is only possible for the 0-cells and the top-dimensional cells in a complex,[8] and there is no such canonical choice of global orientation in the general case for intermediate dimensional cells except in special cases which we will discuss later.

The boundary operator has an important property, which is that $\partial^2 = 0$, i.e. ∂ is nilpotent. This can be seen by applying the boundary twice to the triangle in Fig. 3.2. The result is a sum

[8]Top-dimensional means the highest-dimension of cells in the complex.

over the vertices of the triangle, and each vertex appears twice with opposite signs, thus summing to zero. Thus we have the classic maxim "the boundary of a boundary is empty".

3.1.5 Stokes' Theorem and the Differential

In order to define differential operators for discrete forms, we utilize Stokes' theorem, which can be stated as

$$\int_{\partial U} \omega = \int_U \mathrm{d}\omega. \tag{3.9}$$

On the left is the integral of a k-form ω over the k-dimensional boundary of a $(k+1)$-dimensional domain U. On the right is the integral of a $(k+1)$-form, called $\mathrm{d}\omega$, over U. The operator d is called the exterior derivative or differential. We will take Stokes' theorem as the *definition* for the operation d, which maps a k-form to a $(k+1)$-form. Stokes' theorem tells us that the boundary operation, which acts on geometric integration domains, is adjoint to the differential, which acts on differential forms. The differential of a k-cochain α is a $(k+1)$-cochain $\mathrm{d}\alpha$ whose value on each $(k+1)$-cell is defined precisely according to Stokes' theorem:

$$\mathrm{d}\alpha(\tau_{k+1}) := \alpha(\partial\tau_{k+1}), \tag{3.10}$$

where the boundary operator appearing on the right hand side is the one defined in Eq. (3.8). Expanding the right hand side by linearity, we can express the value of $\mathrm{d}\alpha$ on a cell as

$$\mathrm{d}\alpha(\tau_{k+1}) = \sum_{\tau_k \in \partial\tau_{k+1}} \alpha(\tau_k). \tag{3.11}$$

Note that the nilpotency of the boundary operator ∂ implies that $\mathrm{d}^2 = 0$, i.e.

$$\mathrm{d}(\mathrm{d}\alpha(\xi)) = \mathrm{d}\alpha(\partial\xi) = \alpha(\partial^2\xi) = \alpha(0) = 0. \tag{3.12}$$

3.1.6 The de Rham Complex

Let Ω^k denote the space of k-cochains. For a finite cell complex this is just a finite-dimensional vector space. Let $\Omega = \oplus_{k=0}^{d} \Omega^k$ be the direct sum of these vector spaces.[9] Elements of Ω are linear combinations of k-cochains with different k. The differential is a linear map $\mathrm{d} : \Omega \to \Omega$.[10] Let

[9] A direct sum of vector spaces is akin to a disjoint union of sets. Given two vector spaces U and V containing respectively vectors $\boldsymbol{u}$ and $\boldsymbol{v}$, then $\boldsymbol{u} + \boldsymbol{v}$ is an element of $U \oplus V$. The dimension of $U \oplus V$ is the dimension of U plus the dimension of V.

[10] The notation $f : A \to B$ is to be read "f is a function from A to B".

d_k denote the differential restricted to act on $\Omega^k \subset \Omega$. These maps raise the degree of forms, $\mathrm{d}_k : \Omega^k \to \Omega^{k+1}$. We thus have the following sequence of maps[11]

$$0 \xrightarrow{i} \Omega^0 \xrightarrow{\mathrm{d}_0} \Omega^1 \xrightarrow{\mathrm{d}_1} \Omega^2 \xrightarrow{\mathrm{d}_2} \Omega^3 \xrightarrow{\mathrm{d}_3} 0. \tag{3.13}$$

This sequence is called the de Rham complex.[12] Because $\mathrm{d}^2 = 0$, the de Rham complex is said to be an *exact sequence*. We will discuss this further below when we introduce de Rham cohomology.

3.1.7 Connection to Vector Calculus: Integration on Chains

We quickly review the connection of the Hodge formalism to the vector calculus which is standard in physics and engineering pedagogy. In three dimensional vector calculus there are three primary differential operators: gradient, curl, and divergence. The gradient acts on scalar functions to product a vector field; the curl acts on a vector field to produce another vector field; the divergence acts on a vector field to produce a scalar function. All three of these are naturally encoded in the differential, summarized as follows:

$$\mathrm{d}_0 : \Omega^0 \to \Omega^1 \quad \leftrightarrow \quad \text{grad : scalar functions} \to \text{vector fields}, \tag{3.14}$$

$$\mathrm{d}_1 : \Omega^1 \to \Omega^2 \quad \leftrightarrow \quad \text{curl : vector fields} \to \text{vector fields}, \tag{3.15}$$

$$\mathrm{d}_2 : \Omega^2 \to \Omega^3 \quad \leftrightarrow \quad \text{div : vector fields} \to \text{scalar functions}. \tag{3.16}$$

We can identify 0- and 3-forms as analogs of scalar functions, and 1- and 2-forms as analogs of vector fields.[13]

Given a k-cochain ω and a k-chain ξ, we say that $\omega(\xi)$, which is a number, is the integral of ω over ξ, in accordance with our original definition of cochains in the continuum Eq. (3.1). We think of a chain as a weighted linear combination of integration domains. For example, consider if we wanted to perform a contour integral. We represent the oriented contour C by a 1-chain

$$\xi_C = \sum_{\tau_1 \in C} \tau_1, \tag{3.17}$$

[11] The i is to be read "inclusion".

[12] Note that there are two uses of the word "complex" here: first, we discussed a cellular complex (CW complex), which is a discrete (or piecewise linear) topological space. The de Rham complex is instead an exact sequence of maps, and is an example of a "cochain complex". Both uses are standard and the meaning should be clear from context.

[13] More precisely, 3-forms map to pseudoscalars and 2-forms map to pseudovectors, a distinction which is not necessary when working with the forms directly.

which is a linear combination of a sequence of 1-cells oriented head-to-tail. Then the integral of ω along the contour is simply

$$\int_C \omega \equiv \omega(\xi_C) = \sum_{\tau_1 \in C} \omega(\tau_1). \tag{3.18}$$

Now consider the case where ω is an *exact* 1-cochain, $\omega = \mathrm{d}\phi$ for a 0-cochain ϕ. In the vector calculus language we would treat ω as a vector field given by $\boldsymbol{\omega} = \nabla\phi$. We can integrate this vector field along a contour, C from a point x to a point y, and we know that by the fundamental theorem of calculus we obtain

$$\int_C d\boldsymbol{\ell} \cdot \nabla\phi = \phi(y) - \phi(x). \tag{3.19}$$

In the discrete form language, we have

$$\int_C \mathrm{d}\phi \equiv \mathrm{d}\phi(\xi_C) = \sum_{\tau_1 \in C} \mathrm{d}\phi(\tau_1) \overset{!}{=} \sum_{\tau_1 \in C} \phi(\partial\tau_1) = \phi(\partial\xi_C) \equiv \int_{\partial C} \phi, \tag{3.20}$$

where the ! indicates the use of Stokes' theorem. For every vertex τ_0 along the curve's length, we obtain two contributions in the second sum with opposite orientations which cancel each other, $\phi(\tau_0) + \phi(\bar{\tau}_0) = 0$, except at the two endpoints. Thus the value of $\mathrm{d}\phi(\xi_C)$ is equal to the difference of the values of ϕ at the two endpoints, which is again the fundamental theorem of calculus. In particular, if the curve is closed the result is zero. This result was guaranteed by the use of Stokes' theorem to define the differential.

The same result holds for d_1 and d_2 when integrating over an oriented surface and volume, respectively: the contribution of all cells which are not on the boundary of the integration domain cancel, leaving only the boundary, thus reproducing the curl and divergence theorems of vector calculus. For example, if A is a 1-form and $\boldsymbol{A}$ a corresponding vector field, then $\mathrm{d}_1 A$ is a 2-form corresponding to the pseudovector field $\nabla \times \boldsymbol{A}$. Integrating $\mathrm{d}_1 A$ over a surface S, the curl theorem states that

$$\iint_S (\nabla \times \boldsymbol{A}) \cdot d\boldsymbol{S} = \oint_{\partial S} \boldsymbol{A} \cdot d\boldsymbol{\ell} \leftrightarrow \int_S \mathrm{d}_1 A = \int_{\partial S} A. \tag{3.21}$$

Similarly, if D is a 2-form, corresponding to pseudovector field $\boldsymbol{D}$, $\mathrm{d}_2 D$ is a 3-form corresponding to a pseudoscalar function $\nabla \cdot \boldsymbol{D}$. Integrating over a volume V we have

$$\iiint_V (\nabla \cdot \boldsymbol{D}) dV = \oiint_{\partial V} \boldsymbol{D} \cdot d\boldsymbol{S} \leftrightarrow \int_V \mathrm{d}_2 D = \int_{\partial V} D. \tag{3.22}$$

These formulas are reproduced in the discrete case automatically, because the fundamental theorem of calculus in one dimension, the curl theorem, and divergence theorem are all subsumed into

Stokes' theorem, Eq. (3.9), which we used to define the differential d. Indeed we can view Stokes' theorem as generalizing the fundamental theorem of calculus to arbitrary dimension in the statement that d (the derivative) and ∂ (the boundary relation) are adjoint. Furthermore, the fact that the boundary of a boundary is empty then implies that $\mathrm{d}_k \mathrm{d}_{k-1} = 0$, which automatically reproduces the vector calculus identities

$$\nabla \times (\nabla \phi) = 0 \quad \leftrightarrow \quad \mathrm{d}_1 \mathrm{d}_0 = 0, \tag{3.23}$$

$$\nabla \cdot (\nabla \times \boldsymbol{A}) = 0 \quad \leftrightarrow \quad \mathrm{d}_2 \mathrm{d}_1 = 0. \tag{3.24}$$

3.1.8 Homology

A k-chain is a weighted superposition of oriented k-dimensional integration contours. k-chains can be added, and so form an Abelian group, which we denote $\mathcal{C}_k$.[14] The homology groups H_k of a cellular complex classify topologically distinct chains utilizing the boundary data of the complex. We have the following terminological definitions:

- A k-chain ξ_k is a *k-cycle* if its boundary is empty, $\partial \xi_k = 0$. The k-cycles form an Abelian subgroup of $\mathcal{C}_k$, denoted $\mathcal{Z}_k \equiv \ker \partial_k$.
- A k-chain ξ_k is a *k-boundary* if $\xi_k = \partial \xi_{k+1}$ for some $k+1$-chain. The k-boundaries form an Abelian subgroup of $\mathcal{C}_k$, denoted $\mathcal{B}_k \equiv \mathrm{im}\, \partial_{k+1}$.

Every boundary is a cycle, since $\partial^2 = 0$, but not every cycle need be a boundary, such that

$$\mathcal{B}_k \subseteq \mathcal{Z}_k \subset \mathcal{C}_k. \tag{3.25}$$

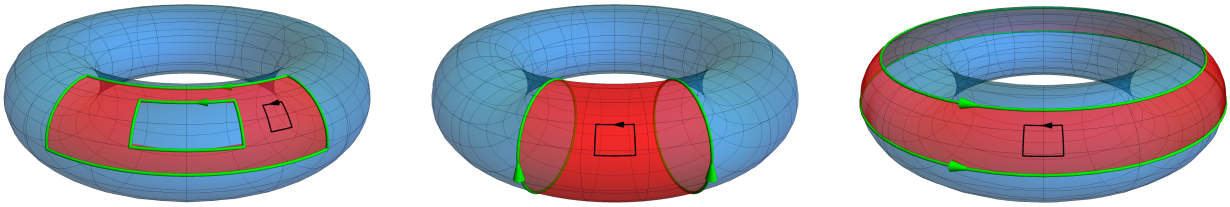
The k-cycles which are not boundaries must represent closed k-dimensional surfaces which wrap around the holes of the complex and are thus not contractible.

The homology groups H_k are defined to capture the failure of a k-cycle to be a k-boundary, as the quotient groups[15]

$$H_k := \frac{\mathcal{Z}_k}{\mathcal{B}_k} \equiv \frac{\ker \partial_k}{\mathrm{im}\, \partial_{k+1}}. \tag{3.26}$$

[14] When computing homology it is standard to take the coefficients of the linear combinations to be integers, which allows homology to detect "twisted" cycles, such as occur in real projective spaces which are non-orientable manifolds [2]. We will allow for real-valued chains as well when representing our weighted integration contours.

[15] The symbol ker means "kernel" and im means "image." For any function $f : A \to B$, the kernel of f is the set of elements of A which are mapped to zero in B. The image of f is the subset of B which are mapped to by f, i.e. the range of f.



(a) (b) (c)

Figure 3.3: Depiction of homologous 1-dimensional closed contours (cycles) on the 2-dimensional torus. Two 1-cycles are homologous if their difference is the boundary of an oriented 2-dimensional region (red). The two cycles in (a) are both homologous to 0, i.e. the trivial cycle, because they are both independently boundaries of a 2-dimensional region on the surface of the torus. The cycles in (b) and (c) are not homologous to the trivial cycle because each cycle by itself does not bound a 2-dimensional region, but the two cycles in (b) and homologous to each other, and the two cycles in (c) are homologous to each other.

Elements of H_k are equivalence classes of k-cycles which differ by a k-boundary, where every boundary cycle is homologous to the trivial zero cycle by definition. Thus H_k is a topological invariant which classifies the distinct k-dimensional contours which do not mutually bound a $(k+1)$-dimensional bulk. This is illustrated for a 2-torus T^2 in Fig. 3.3, (a) shows two cycles which are both homologous to the trivial cycle, (b) and (c) each show a pair of homologous cycles which mutually bound an oriented two-dimensional region, but are non-trivial since each cycle by itself does not bound a two-dimension region. The cycles in (b) are not homologous to the cycles in (c), and we conclude there are two non-trivial homology classes: $H_1(T^2) = \mathbb{Z} \oplus \mathbb{Z}$.

3.1.9 Cohomology

Whereas chains represent oriented integration surfaces, cochains represent functions on those surfaces, i.e. some data attached to the surfaces. Cohomology is a topological invariant of the complex which parallels homology. For real- and complex-valued cochains, the space of cochains Ω^k is a vector space. We have the following terminological definitions:

- A k-cochain ω^k is said to be *closed* if its exterior derivative ("gradient") is zero, $\mathrm{d}\omega^k = 0$. A closed cochain may also be called a *cocycle*. The cocycles form a linear subspace of Ω^k, denoted $\mathcal{Z}^k \equiv \ker \mathrm{d}_k$.
- A k-cochain ω^k is said to be *exact* if $\omega^k = \mathrm{d}\omega^{k-1}$ for some $(k-1)$-cochain. The exact k-cochains form a linear subspace of Ω^k, denoted $\mathcal{B}^k \equiv \mathrm{im}\, \mathrm{d}_{k-1}$.

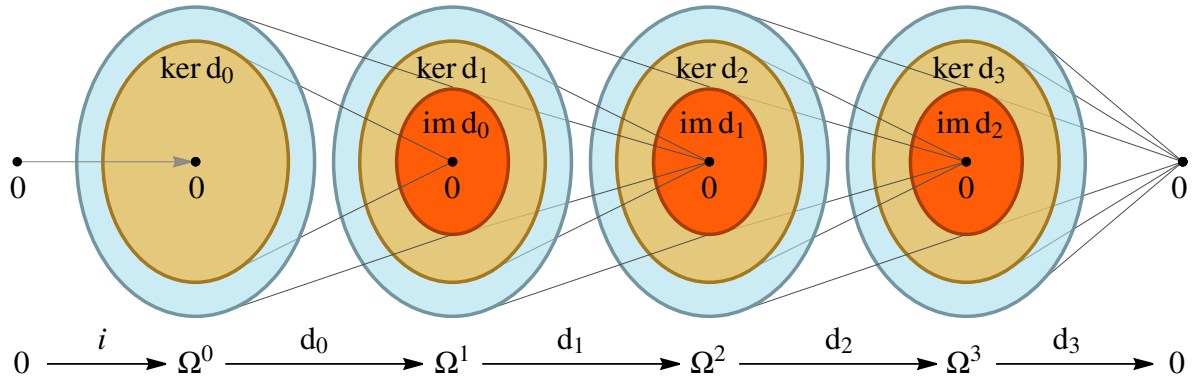


Figure 3.4: A depiction of the de Rham complex, which is an exact sequence. The kernel of d_k is a subspace of Ω^k, and the image of d_{k-1} is a subspace of the kernel of d_k, which follows from $\mathrm{d}_k\mathrm{d}_{k-1} = 0$. Elements of $\operatorname{im}\mathrm{d}_{k-1}$ are *exact* k-forms and elements of $\ker\mathrm{d}_k$ are *closed* k-forms. Every exact k-form is closed, but not every closed k-form need be exact. Cohomology measures the failure of a closed form to be exact. Elements of H^k are equivalence classes of closed forms which differ by an exact form.

Every exact cochain is trivially closed, since $\mathrm{d}^2 = 0$, but not every closed cochain need be exact. This is depicted in Fig. 3.4, which pictorially shows the de Rham complex Eq. (3.13). The image of d_{k-1}, the space of exact k-cochains, is entirely contained within the kernel of d_k, the space of closed k-cochains.[16] The failure of a closed cochain to be exact is classified by cohomology.[17] The k'th de Rham cohomology space H^k is defined just like homology but with the differential,

$$H^k := \frac{\mathcal{Z}^k}{\mathcal{B}^k} \equiv \frac{\ker\mathrm{d}_k}{\operatorname{im}\mathrm{d}_{k-1}}. \tag{3.27}$$

The elements of H^k are equivalence classes of k-cochains under the equivalence relation[18]

$$\alpha^k \sim \beta^k \iff \alpha^k - \beta^k = \mathrm{d}_{k-1}\gamma^{k-1} \tag{3.28}$$

for some $(k-1)$-cochain γ^{k-1}. The de Rham cohomology of the complex is denoted $H^\bullet := \oplus_{k=0}^{d} H^k$. It is often said that the de Rham cohomology is an *obstruction* the the global exactness of forms.

[16] Such a sequence of maps is called an exact sequence. In particular, the boundary maps induce an analogous exact sequence for the chains, but the boundary maps form a descending sequence (i.e. going right to left), generically called a chain complex, while the differentials form an ascending sequence, generically called a cochain complex.

[17] Cohomology is a vast and highly generalizable concept from algebraic topology which can be applied in far more abstract cases than what we are discussing here. The version of cohomology we are using is called de Rham cohomology, since it derived from the de Rham complex.

[18] The symbol $\sim$ can be read as "equivalent to". The symbol $\iff$ is to be read "if and only if".

The Poincaré lemma states that every form is *locally* exact (or "locally a gradient"), and cohomology is essentially therefore a non-local topological invariant that measures the failure of closed forms to be *globally* exact [2]. Cohomology will be important because it classifies solutions of the Laplace equation, the so-called harmonic forms, which have non-zero integral over the non-trivial cycles of the complex. From here on we will use the words "cochain" and "form" interchangeably.

3.1.10 Inner Product of Cochains

We will equip the vector spaces Ω^k with an inner product $\langle -, - \rangle_k$ defined as

$$\langle \alpha, \beta \rangle_k = \sum_{\tau_k} \alpha(\tau_k)\beta(\tau_k). \tag{3.29}$$

We will often omit the subscript k, since it is implicit based on the degree of the forms appearing within. Other choices of inner product are of course possible, but this is the only one we will use in this work.

3.1.11 The Codifferential and Coboundary

We define the codifferential δ to be the adjoint of the differential with respect to the inner product Eq. (3.29),

$$\langle \mathrm{d}\alpha, \beta \rangle = \langle \alpha, \delta\beta \rangle. \tag{3.30}$$

We can obtain a concrete expression for $\delta\beta$ by first expanding both sides,

$$\sum_{\tau_k} \mathrm{d}\alpha(\tau_k)\beta(\tau_k) = \sum_{\tau_{k-1}} \alpha(\tau_{k-1})\delta\beta(\tau_{k-1}). \tag{3.31}$$

Utilizing Stokes' theorem, in the form Eq. (3.11), the left hand side of Eq. (3.31) can be written as

$$\sum_{\tau_k} \mathrm{d}\alpha(\tau_k)\beta(\tau_k) = \sum_{\tau_k} \left(\sum_{\tau_{k-1} \in \partial\tau_k} \alpha(\tau_{k-1}) \right) \beta(\tau_k). \tag{3.32}$$

The first sum is over all oriented k-cells, and the second is over the oriented boundary of each k-cell. We can rearrange this double sum into a sum over all oriented $(k-1)$-cells τ_{k-1} and a sum over all k-cells whose boundary includes τ_{k-1},[19]

$$\sum_{\tau_k} \sum_{\tau_{k-1} \in \partial\tau_k} \alpha(\tau_{k-1})\beta(\tau_k) = \sum_{\tau_{k-1}} \sum_{\{\tau_k | \tau_{k-1} \in \partial\tau_k\}} \alpha(\tau_{k-1})\beta(\tau_k). \tag{3.33}$$

[19] The notation $\{x|P(x)\}$ is to be read "the set of all x such that $P(x)$ is true.

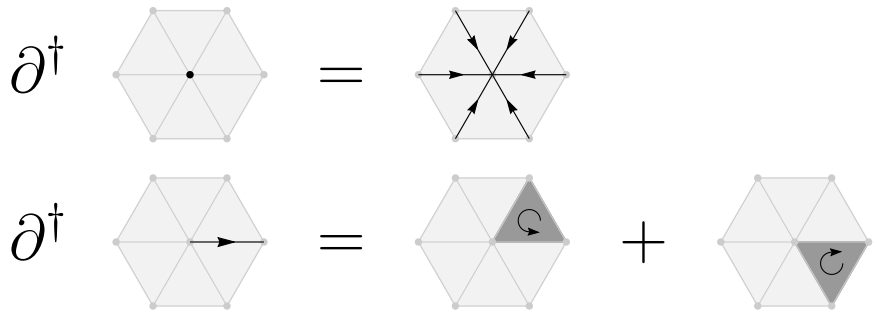


Figure 3.5: Top: coboundary of a 0-cell in a two dimensional cell complex, which is the sum of all six oriented edges ending terminating at that cell. Bottom: The coboundary of a 1-cell is the sum of all 2-cells containing that cell in their boundary.

Comparing this to the right hand side of Eq. (3.31), we can conclude that

$$\delta\beta(\tau_{k-1}) = \sum_{\{\tau_k|\tau_{k-1}\in\partial\tau_k\}} \beta(\tau_k), \tag{3.34}$$

where the sum runs over all k-cells τ_k such that τ_{k-1} is on the boundary of τ_k and their orientations match. This espression should be compared to Eq. (3.11). Indeed it is convenient to introduce a "coboundary" operator, $\partial^\dagger$, which acts on chains as

$$\partial^\dagger\tau_k = \sum_{\{\tau_{k+1}|\tau_k\in\partial\tau_{k+1}\}} \tau_{k+1},$$

so that we can write Eq. (3.34) as a "co-Stokes' theorem",

$$\delta\alpha(\tau_{k-1}) = \alpha(\partial^\dagger\tau_{k-1}). \tag{3.35}$$

This is notationally compact and nicely parallels Stokes' theorem. Note that the coboundary $\partial^\dagger$ has no continuum analog, but will be an essential conceptual tool for discrete cellular complexes. In Fig. 3.5, we illustrate the coboundary of a vertex and an edge for a two-dimensional cellular complex. The coboundary of a vertex oriented "in" is the sum of all oriented edges terminating at that vertex, while the coboundary of an oriented edge is the sum of all oriented plaquettes containing that edge in their boundary.

3.1.12 The Hodge Laplacian and Harmonic Forms

We can now define the Hodge Laplacian (also called the Laplace-de Rham operator) for cochains as

$$\Delta = (\mathrm{d} + \delta)^2 = \mathrm{d}\delta + \delta\mathrm{d}. \tag{3.36}$$

We can also define the Laplacian for k-cochains as

$$\Delta_k = \mathrm{d}_{k-1}\delta_k + \delta_{k+1}\mathrm{d}_k. \tag{3.37}$$

Note that while d raises the degree of forms and δ lowers the degree of forms, Δ leaves the degree unchanged, i.e. $\Delta_k : \Omega^k \to \Omega^k$. A k-cochain γ satisfying Laplace's equation $\Delta_k\gamma = 0$ is said to be *harmonic*, and satisfies both $\mathrm{d}\gamma = 0$ and $\delta\gamma = 0$. The space of harmonic cochains is by definition the kernel of the Laplace operator.

3.1.13 The Helmholtz-Hodge Decomposition

We can now define the Hodge decomposition for closed cellular complexes (compact and without boundary). In vector calculus this is called the Helmholtz decomposition, and we will use the common moniker Helmholtz-Hodge decomposition (HHD). It states that every differential k-form ω can be written as a sum

$$\omega = \mathrm{d}\alpha + \delta\beta + \gamma\,, \tag{3.38}$$

where $\alpha \in \Omega^{k-1}$, $\beta \in \Omega^{k+1}$, and $\gamma \in \ker\Delta_k$. In 3D Euclidean vector calculus this is often written for a vector field as

$$\boldsymbol{E} = \operatorname{grad}\phi + \operatorname{curl}\boldsymbol{A} + \boldsymbol{E}_0\,, \tag{3.39}$$

where ϕ is a scalar function, $\boldsymbol{A}$ is another vector field,[20] and $\boldsymbol{E}_0$ is a harmonic vector field (which in Euclidean space is a uniform constant vector field). The first term in Eq. (3.38) is an *exact* form, and the second term is a *coexact* form, while the last term is *harmonic*. The word "exact" should evoke "local": the operators d and δ are differentials or gradients, the value of $\mathrm{d}\alpha$ and $\delta\beta$ on individual cells of the complex are determined by the values of α and β on nearby cells which are connected by the boundary data. The harmonic portion of the decomposition is precisely the portion which is non-local and cannot be expressed in terms of purely local data.

[20] Note that if $\boldsymbol{E}$ is a vector field then $\boldsymbol{A}$ must be a pseudovector field, so that $\nabla \times \boldsymbol{A}$ is a vector field. The distinction between vectors and pseudovectors (and between scalars and pseudoscalars) is not necessary when working directly with forms: ϕ is a 0-form and A is a 2-form, E is a 1-form.

The HHD can equivalently be stated as a decomposition of Ω^k into a direct sum of subspaces,

$$\Omega^k = \underbrace{\operatorname{im} \mathrm{d}_{k-1} \oplus \overbrace{\ker \Delta_k}^{\ker \delta_k}}_{\ker \mathrm{d}_k} \oplus \operatorname{im} \delta_{k+1} . \tag{3.40}$$

Note that the space of harmonic k-forms is precisely the intersection $\ker \mathrm{d}_k \cap \ker \delta_k$. Also note that the cohomology class of ω in Eq. (3.38) is equivalent to the cohomology class of γ.

3.1.14 "Gauge Invariance" and the HHD

In the HHD Eq. (3.38), the fact that $\mathrm{d}^2 = \delta^2 = 0$ implies that $\alpha \in \Omega^{k-1}$ and $\beta \in \Omega^{k+1}$ are not uniquely defined. In particular, we can always make the substitution $\alpha \longrightarrow \alpha + \mathrm{d}\alpha'$, with $\alpha' \in \Omega^{k-2}$, and $\beta \longrightarrow \beta + \delta\beta'$, with $\beta' \in \Omega^{k+2}$. In other words, α is only defined up to addition of an exact form and β is only defined up to addition of a coexact form. This is a type of gauge freedom, as the term is often used by physicists.

Consider the classic example of a magnetic field $\boldsymbol{B}$. The magnetic field has zero divergence, $\nabla \cdot \boldsymbol{B} = 0$, and therefore can be written as the curl of a vector potential, $\boldsymbol{B} = \nabla \times \boldsymbol{A}$. The vector potential is not uniquely defined, because we can add to it the gradient of any smooth function, $\boldsymbol{A} \longrightarrow \boldsymbol{A} + \nabla\lambda$. The vector potential $\boldsymbol{A}$ (and sometimes the magnetic field $\boldsymbol{B}$) is generally called a "gauge field", and the shifting of $\boldsymbol{A}$ by a gradient as a gauge transformation. This can be understood entirely in terms of the in terms of the HHD, Eq. (3.38). Let B be a 2-form and A be a 1-form, then

$$B = \mathrm{d}A + 0 + B_0. \tag{3.41}$$

The fact that the coexact term is zero follows from the fact that $\mathrm{div} B \equiv \mathrm{d}B = 0$, which will see in more detail shortly. The harmonic portion B_0 corresponds to a uniform magnetic field filling space. The gauge transformation $\boldsymbol{A} \longrightarrow \boldsymbol{A} + \nabla\lambda$ then corresponds to the substitution $A \longrightarrow A + \mathrm{d}\lambda$, where λ is a 0-form. The "gauge invariance" under this substitution comes from the fact that $\mathrm{d}^2 = 0$.

What this "gauge invariance" or "gauge freedom" means is the following. In the HHD, Eq. (3.38), $\alpha \in \Omega^{k-1}$, and $\mathrm{d}\alpha \in \operatorname{im} \mathrm{d}_{k-1} \subset \Omega^k$. Eq. (3.40) says that we can express the map $\mathrm{d}_{k-1} : \Omega^{k-1} \longrightarrow \Omega^k$ as

$$\mathrm{d}_{k-1} : \underbrace{\ker \mathrm{d}_{k-1} \oplus \operatorname{im} \delta_k}_{\Omega^{k-1} \ni \alpha} \longrightarrow \operatorname{im} \mathrm{d}_k. \tag{3.42}$$

Anything contained in $\ker \mathrm{d}_{k-1}$ gets mapped to $0 \in \operatorname{im} \mathrm{d}_k$, and therefore the unique part of α which appears in the HHD as $\mathrm{d}\alpha$ is entirely contained in $\operatorname{im} \delta_k$. Thus α is only defined up to addition of

an element of $\ker \mathrm{d}_{k-1}$. This is the so-called "gauge invariance" in the choice of α: the "physical" quantity of interest is $\mathrm{d}\alpha$, which knows nothing about the portion of α in $\ker \mathrm{d}_{k-1}$.

In order to remove this ambiguity, we can enforce a "gauge condition". Since $\operatorname{im} \delta_k \subset \ker \delta_{k-1}$ (see Eq. (3.40)), we can always choose the gauge condition $\delta\alpha = 0$, which ensures that $\alpha \in \ker \delta_{k-1}$, uniquely specifying α up to a harmonic $(k-1)$-form. We call this the Coulomb gauge condition since it generalizes $\nabla \cdot \boldsymbol{A} = 0$ from electromagnetism: the magnetic field B is a closed 2-form, locally of the form $B = \mathrm{d}A$ where A is the vector potential 1-form, and the Coulomb gauge condition $\nabla \cdot \boldsymbol{A} = 0$ is equivalent to $\delta A = 0$. Similarly, for the term $\delta\beta$ in the HHD we can write

$$\delta_{k+1} : \underbrace{\ker \delta_{k+1} \oplus \operatorname{im} \mathrm{d}_k}_{\Omega^{k+1} \ni \beta} \to \operatorname{im} \delta_{k+1}, \tag{3.43}$$

such that the unique portion of β is the part contain in $\operatorname{im} \mathrm{d}_k$. Thus we can choose the "co-Coulomb" condition $\mathrm{d}\beta = 0$, since $\operatorname{im} \mathrm{d}_k \subset \ker \mathrm{d}_{k+1}$.

3.1.15 Computing the HHD: Poisson's Equation

We now consider the problem of actually computing the HHD from a given k-form ω in the form of Eq. (3.38), which can be obtained from $\mathrm{d}\omega$ and $\delta\omega$. The complementary problem is that we are given $\mathrm{d}\omega$ and $\delta\omega$ and wish to determine ω, which can be done by determining its HHD. First, applying d and δ to Eq. (3.38), using $\mathrm{d}^2 = \delta^2 = 0$ and $\mathrm{d}\gamma = \delta\gamma = 0$, we obtain

$$\delta\omega = \delta\mathrm{d}\alpha, \tag{3.44}$$
$$\mathrm{d}\omega = \mathrm{d}\delta\beta, \tag{3.45}$$

From the definition of the Laplacian, $\Delta = \mathrm{d}\delta + \delta\mathrm{d}$, we can write

$$\delta\omega = \delta\mathrm{d}\alpha \tag{3.46}$$
$$= (\Delta - \mathrm{d}\delta)\alpha \tag{3.47}$$
$$= \Delta\alpha - \mathrm{d}(\delta\alpha) \tag{3.48}$$

and similarly $\mathrm{d}\omega = \Delta\beta - \delta(\mathrm{d}\beta)$. Applying the Coulomb gauge condition $\delta\alpha = 0$ ($\mathrm{d}\beta = 0$), these reduce to Poisson's equation

$$\Delta\alpha = \delta\omega \quad (\delta\alpha = 0), \tag{3.49}$$
$$\Delta\beta = \mathrm{d}\omega \quad (\mathrm{d}\beta = 0). \tag{3.50}$$

Thus, given $\mathrm{d}\omega$ and $\delta\omega$, α and β are given in Coulomb gauge by the formal solutions

$$\alpha = \Delta_{k-1}^{-1}(\delta\omega), \tag{3.51}$$

$$\beta = \Delta_{k+1}^{-1}(\mathrm{d}\omega). \tag{3.52}$$

Thus the inverse of the Laplacian plays a central role in the HHD, as it does in electrodynamics. Note that the Laplacian inverse requires careful definition when the cohomology $H^\bullet$ is non-trivial, i.e. when $\ker\Delta$ is non-empty.

Using Eqs. (3.51) and (3.52), we can rewrite Eq. (3.38) as

$$\omega = (\mathrm{d}\Delta^{-1}\delta)\omega + (\delta\Delta^{-1}\mathrm{d})\omega + \gamma \tag{3.53}$$

Comparing to Eq. (3.40), we observe that the operators appearing in Eq. (3.53) are the linear projectors

$$\mathcal{P}^{\mathrm{im}\,\mathrm{d}} \equiv \mathrm{d}\Delta^{-1}\delta : \Omega \to \mathrm{im}\,\mathrm{d}, \tag{3.54}$$

$$\mathcal{P}^{\mathrm{im}\,\delta} \equiv \delta\Delta^{-1}\mathrm{d} : \Omega \to \mathrm{im}\,\delta, \tag{3.55}$$

Note that the harmonic component cannot be determined from the local data $\mathrm{d}\omega$ and $\delta\omega$, since $\mathrm{d}\gamma = \delta\gamma = 0$. The projector to the harmonic subspace can be defined by subtracting off the exact and coexact parts,

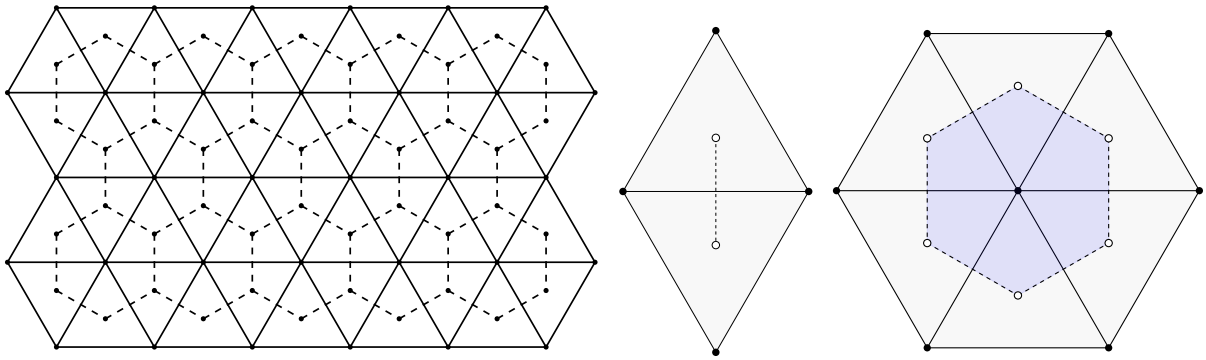
$$\mathcal{P}^{\ker\Delta} \equiv 1 - \mathcal{P}^{\mathrm{im}\,\mathrm{d}} - \mathcal{P}^{\mathrm{im}\,\delta}, \tag{3.56}$$

such that $\gamma = \mathcal{P}^{\ker\Delta}\omega$. Therefore to summarize, if we know the projectors, then the HHD of a k-form ω is given by

$$\omega^k = \underbrace{\mathcal{P}^{\mathrm{im}\,\mathrm{d}_{k-1}}\omega}_{\text{exact}} + \underbrace{\mathcal{P}^{\mathrm{im}\,\delta_{k+1}}\omega}_{\text{co-exact}} + \underbrace{\mathcal{P}^{\ker\Delta_k}\omega}_{\text{harmonic}}. \tag{3.57}$$

3.1.16 Hodge Duality and the Dual Cell Complex

For any cellular complex, we can define its dual cell complex where each k-cell of the original complex is identified with a $(d-k)$-cell of the dual complex. To form the dual complex, we follow the same method as described for the construction of CW complex in Section 3.1.1, but using the data from the original cell complex as a guide: form the dual 0-skeleton $\tilde{\mathcal{K}}_0$ by placing a new 0-cell at the center of each d-cell of the original complex. Next we form the dual 1-skeleton $\tilde{\mathcal{K}}_1$, then the dual 2-skeleton $\tilde{\mathcal{K}}_2$, and so on. To build $\tilde{\mathcal{K}}_k$, we glue the boundary of a k-cell to the $(k-1)$-cells of the dual complex corresponding to the *coboundary* of a $(d-k)$ cell of the original complex.



(a) (b) (c)

Figure 3.6: (a) The triangular lattice cell complex and its dual cell complex, the honeycomb lattice. (b) To construct the dual complex we start by placing a dual vertex (open circle) inside every top-dimensional cell (triangle) of the original complex. We then insert the dual 1-cells (dashed line) by gluing the boundary to two vertices whose original 2-cells form the coboundary of an original 1-cell. (c) We then insert the dual 2-cells (hexagon) by gluing the boundary along the dual 1-cells (dashed lines) whose original 1-cells form the coboundary of an original 0-cell. In this way the boundary relation of the dual complex reflects the coboundary relation of the original complex.

This is demonstrated in Figure Fig. 3.6. In (a), solid lines form the edges of a triangular lattice, which is a $d = 2$ dimensional cellular complex, where the 2-cells are triangles, and dashed lines form the edges of the dual cell complex, which is a honeycomb lattice. In (b), we have placed a single 0-cell (open circle) at the center of each original 2-cell (triangle). We then glue the boundary of a 1-cell (dashed line) to two 0-cells (open circles) corresponding to the coboundary of a 1-cell (horizontal edge) of the original complex. Finally in (c), we glue the boundary of a 2-cell to the 1-cells of the dual complex corresponding to the coboundary of a 0-cell (center point) of the original complex.

The end result is that each k-cell of the original complex is mapped to a $(d-k)$-cell of the dual complex, and the boundary relation of the dual complex corresponds to the coboundary relation of the original lattice. This construction is indeed a duality, since the dual complex of the dual complex is the original complex. But note that the two complexes need not be the same, i.e. the triangular and honeycomb lattice complexes are distinct. The same construction for a square lattice complex would yield another square lattice, and thus we say that the square lattice is self-dual.

The dual complex construction also allows us to define an isomorphism between the k-forms Ω^k of the original complex and the $(d-k)$-forms $\Omega^{d-k}_{\text{dual}}$ on the dual complex. The linear map is

called the Hodge star,

$$\star : \Omega^k \to \Omega^{d-k}_{\text{dual}} \tag{3.58}$$

Given a k-form ω on the original complex, the value of $\star\omega$ on a dual $(d-k)$-cell is proportional to the value on the corresponding k-cell of the original complex, encoded in the linear map $\star$. The proportionality is not uniquely defined,[21] and we are free to choose it based on relevant physical intuition. In the particular case of electromagnetism, one may have noticed that there are both 1-form and 2-form realizations of the electric and magnetic fields in Table 1.1. They are related in the continuum theory by $H = \star B$ and $D = \star E$, i.e. $\star$ encodes the permittivity and permeability of the medium.

The important point is that any theory formulated in terms of k-forms on the original complex can be equivalently formulated in terms of $(d-k)$-forms on the dual complex. This will be utilized in Chapter 5 to construct a 2-form dual of spin ice. In particular, note that the codifferential δ is essentially equivalent to the differential on the dual complex, and indeed it can be defined as

$$\delta = \star^{-1} \mathrm{d}^{\text{dual}} \star \,. \tag{3.59}$$

This immediately tells us that the codifferential reproduces gradient, divergence, and curl in the "opposite direction" of d, i.e.

$$\delta_3 : \Omega^3 \to \Omega^2 \sim \text{grad}, \tag{3.60}$$

$$\delta_2 : \Omega^2 \to \Omega^1 \sim \text{curl}, \tag{3.61}$$

$$\delta_1 : \Omega^1 \to \Omega^0 \sim \text{div}. \tag{3.62}$$

We will see these relations are true up to a sign.

3.2 The Laplacian and the Adjacency Matrix

In this section we will apply the formalism just introduced to the analysis of the spectrum of adjacency matrices of lattices, which arise as the interaction matrix for nearest-neighbor antiferromagnetic models discussed extensively in the previous chapter. In particular, we will see that the adjacency matrix for a lattice is directly related to the Laplacian Δ_0, a classic result from graph theory. More importantly, we will see that for line graph lattices (cf. Section 2.5.1),

[21] In the continuum case of differential forms, the $\star$ operator is defined if there is a uniquely defined volume form (a volume form is a top-dimensional form), denoted vol^d, which is the case in particular when the manifold is equipped with a metric. Then the Hodge star is defined by $\omega \wedge \star\omega = \text{vol}^d$. Since we have note defined a wedge product, we will not use this.

if we interpret the spin degrees of freedom as 1-forms on the parent lattice, treated as a cellular complex, then the adjacency matrix of the line graph lattice is precisely $\mathrm{d}_0\delta_1$, which is "half" of the 1-form Laplacian $\Delta_1 = \mathrm{d}_0\delta_1 + \delta_2\mathrm{d}_1$. As we will see, the absence of the other "half", $\delta_2\mathrm{d}_1$, is the reason for the presence of flat bands in the spectrum of the line graph adjacency matrix.

3.2.1 The Graph Laplacian

Here we quickly demonstrate that our definition of the 0-form Laplacian $\Delta_0 = \delta_1\mathrm{d}_0$ is equivalent to the standard definition of the Laplacian of a graph. Any graph can be considered a CW complex with only 0-cells and 1-cells, so the definition of d_0 and δ_1 are unambiguous, and conversely given any cellular complex, its 1-skeleton is a graph. When dealing with only the 1-skeletons, it is extremely convenient to utilize a simple notation for the cells, chains, and cochains. First, we will use i to index 0-cells, which can be written as v_i (v for "vertex"). A 0-chain can be written as $\sum_i c_i v_i$, and a 0-cochain (scalar function) f has value on cell i of $f(v_i) \equiv f_i$. Every 1-cell has exactly two endpoints, and its orientation is fully specified by the order of the endpoints, so an oriented 1-cell can be specified by e_{ij} (e for "edge") oriented from v_i to v_j, such that $\partial e_{ij} = v_j - v_i$. The value of a 1-cochain α on a 1-cell is then denoted $\alpha(e_{ij}) \equiv \alpha_{ij} = -\alpha_{ji}$.

Utilizing this notation, the differential of f is a 1-cochain $\mathrm{d}_0 f$ whose value on a 1-cell is

$$f_{ij} \equiv \mathrm{d}_0 f(e_{ij}) = f(\partial e_{ij}) = f(v_j - v_i) = f(v_j) - f(v_i) \equiv f_j - f_i, \tag{3.63}$$

which simply demonstrates that d_0 computes the finite difference gradient of the scalar function f across each 1-cell: $\mathrm{d}_0 \sim \mathrm{grad}$. Similarly, for α a 1-cochain, its codifferential is a 0-cochain $\delta_1\alpha$, whose value on a 0-cell can be expressed as[22]

$$(\delta_1\alpha)_j \equiv \delta_1\alpha(v_j) = \alpha(\partial^\dagger v_j) = \sum_{\{e|v_j\in\partial e\}} \alpha(e) \overset{!}{\equiv} \sum_{i@j} \alpha(e_{ij}) \equiv \sum_{i@j} \alpha_{ij}, \tag{3.64}$$

where in the indicated equivalence we rewrote the "sum over all 1-cells e terminating at 0-cell v_j" as the "sum over all 0-cells v_i adjacent to (@) 0-cell v_j", where adjacent to here means connected by a 1-cell. If we interpret α_{ij} as the "flow of α from i to j", then the final expression is the negative of the divergence of α:[23]

$$\delta_1 \leftrightarrow -\mathrm{div}. \tag{3.65}$$

[22] The notation $i@j$ is to be read "all vertices i adjacent to vertex j".

[23] That is, $-(\delta_1\alpha)_j = \sum_{i@j} \alpha_{ji} > 0$ if the net flow is "out" of j

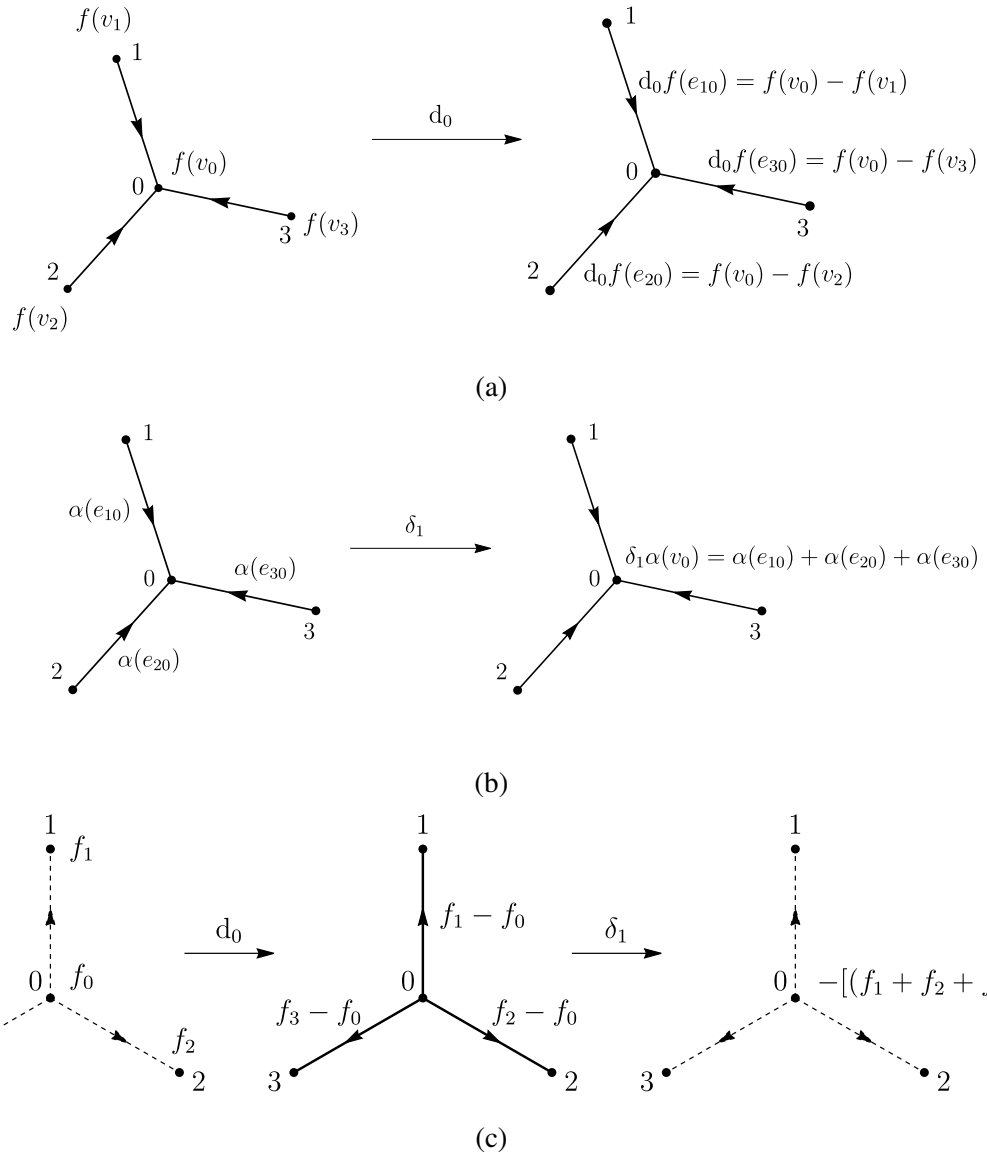


Figure 3.7: Depiction of (a) the differential d_0, (b) the codifferential δ_1, and (c) the 0-form Laplacian $\Delta_0 = \delta_1 \mathrm{d}_0$.

The above two formulas are depicted in Fig. 3.7(a,b). This notation is very helpful, compact, and transparent when working only with the 1-skeleton, but does not generalize very nicely to higher-dimensional cells.[24]

Now consider the 0-form Laplacian $\Delta_0 = \delta_1 \mathrm{d}_0$, which sends a 0-cochain to a 0-cochain,

$$(\Delta_0 f)_j = (\delta_1 \mathrm{d}_0 f)_j = \sum_{i@j} (\mathrm{d}_0 f)_{ij} = \sum_{i@j} (f_j - f_i) = z_j f_j - \sum_{i@j} f_i, \tag{3.66}$$

where z_j is the number of neighbors of site j. This formula is depicted in Fig. 3.7(c) for $z_0 = 3$. Since Δ_0 is a linear operator we can express the above as

$$(\Delta_0 f)_i = \sum_j (\Delta_0)_{ij} f_j = \sum_j \left[z_i \delta_{ij} - \mathcal{A}_{ij} \right] f_j, \tag{3.67}$$

where $\mathcal{A}$ is the adjacency matrix for the graph, with matrix elements $\mathcal{A}_{ij} = +1$ for neighboring vertices and 0 otherwise. Thus we conclude that

$$(\Delta_0)_{ij} = D_{ij} - \mathcal{A}_{ij}, \tag{3.68}$$

where $D_{ij} = z_i \delta_{ij}$ counts the number of nearest-neighbors. This expression is the standard graph Laplacian. The Laplacian is positive-definite since

$$\langle f, \delta_1 \mathrm{d}_0 f \rangle = \langle \mathrm{d}_0 f, \mathrm{d}_0 f \rangle = \sum_i |f_i|^2 > 0. \tag{3.69}$$

Note also that $\Delta_0 = \delta_1 \mathrm{d}_0$ is precisely analogous to the vector calculus statement $\Delta_0 = -\mathrm{div}\,\mathrm{grad}$.

3.2.2 Line Graphs and 1-Forms

We now consider the 1-form Laplacian $\Delta_1 = \mathrm{d}_0 \delta_1 + \delta_2 \mathrm{d}_1$. In vector calculus this would be

$$\Delta_1 = -\mathrm{grad}\,\mathrm{div} + \mathrm{curl}\,\mathrm{curl}. \tag{3.70}$$

Our aim here is to demonstrate that the operator $\mathrm{d}_0 \delta_1 \sim -\mathrm{grad}\,\mathrm{div}$ is related to the adjacency matrix of the *line graph* of the 1-skeleton of the complex. In particular we wish to show that,

[24] For *simplicial* complexes—those for which every k-cell is a k-simplex (point, line, triangle, tetrahedron, etc.)—it is in fact useful to denote a simplex by a list of its vertices, because a k-simplex has exactly $k+1$ vertices, and deleting one vertex from the list leaves the opposing $(k-1)$-simplex face. Thus simplicial complexes have a special combinatorial structure that makes them particularly useful as abstract constructions for studying topological spaces via triangulations (or "simplicializations"). However, our cell complexes based on lattices of atoms are generically not simplicial, with the sole exception being the triangular lattice.

under certain key assumptions, if we treat a 1-cochain α on the cell complex as a 0-cochain taking values on the vertices of the line graph of the 1-skeleton, then the following formula holds true:

$$(\mathrm{d}_0\delta_1\alpha)_e = \sum_{e'}[\mathcal{A}^{\mathrm{lg}}_{ee'} + 2\delta_{ee'}]\alpha_{e'}, \tag{3.71}$$

where $\mathcal{A}^{\mathrm{lg}}_{ee'}$ is the adjacency matrix of the line graph.

First we need to fix an ambiguity, which is how to assign the value of α on an oriented edge e to a value α_e on the corresponding vertex of the line graph. To do so, we must *fix a reference orientation*, o. We pick one orientation for every edge in the complex, which we denote e_o, then to each vertex of the line graph we assign the value $\alpha_e \equiv \alpha(e_o)$. Technically we had to do the same for the 0-form Laplacian, but in that case we can always choose every vertex to have the *same* orientation. This is not possible for the 1-cells.[25] In the special but important case where the parent lattice is bipartite, there are two canonical orientations, where every edge is oriented from sublattice A to B or vice-versa. This is the case for most line graph lattices of interest, in particular the pyrochlore, kagome, and checkerboard lattices. In this bipartite case, Eq. (3.71) follows from Fig. 3.8, where the fixed orientations are indicated by arrows. In the case that the underlying 1-skeleton of the parent lattice is not bipartite, then Eq. (3.71) takes the more general form

$$(\mathrm{d}_0\delta_1\alpha)_e = \sum_{e'}[\langle \partial e_o, \partial e'_o\rangle + 2\delta_{ee'}]\alpha_{e'}, \tag{3.72}$$

where $\langle \partial e_o, \partial e'_o\rangle = 1$ if e_o and e'_o share a vertex on their boundary with the same orientation, -1 if the orientation is opposite, and 0 otherwise.

Frustration on line graph lattices with non-bipartite parents was studied by Ref [94] and found to have qualitatively distinct behavior to that of line graphs with a bipartite parent lattice. This is an interesting point that warrants further study in relation to the current formalism, but we will not do so in this thesis, instead focusing on the bipartite case where Eq. (3.71) holds exactly. We note however that expressing the Hamiltonian in terms of the adjacency matrix to begin with was dependent on the choice of orientation of the quantization axis. If we were to flip one the quantization axis of the one of the spins we would correspondingly have to change the sign of its coupling with all of its neighbors, so in a sense as we will see expressing the Hamiltonian in terms of the matrix Eq. (3.72) is more intrinsic as it is independent of the orientations that we

[25] Only the 0-cells and top-dimensional cells can have a global orientation in the general case, so long as the underlying manifold is orientable. The case where the underlying manifold is non-orientable is interesting in its own right, but we will not study it here. Standard examples of orientable manifolds are spheres, tori, planes, Euclidean space, etc. Standard examples of non-orientable manifolds are the Möbius strip, the Klein bottle, and projective spaces.

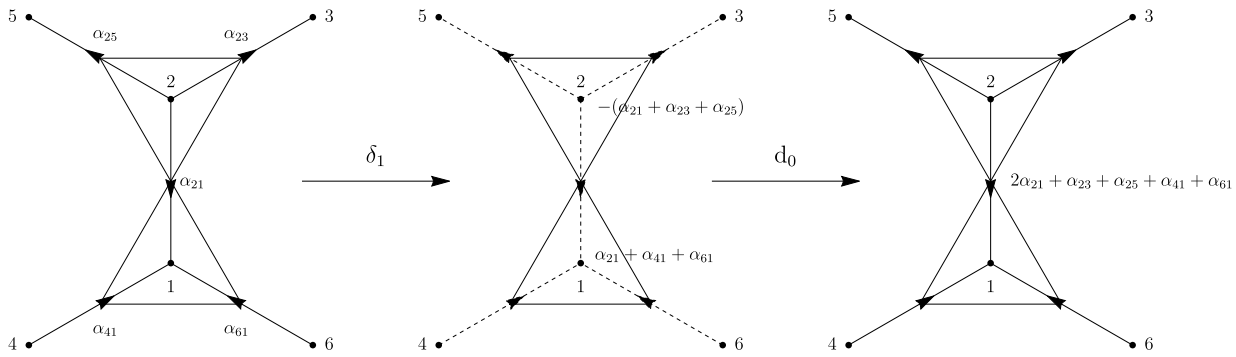


Figure 3.8: Demonstration of Eq. (3.71) when edges can be oriented consistently "in" and "out" at every vertex. Triangles indicate edges of the line graph lattice of the underlying parent lattice.

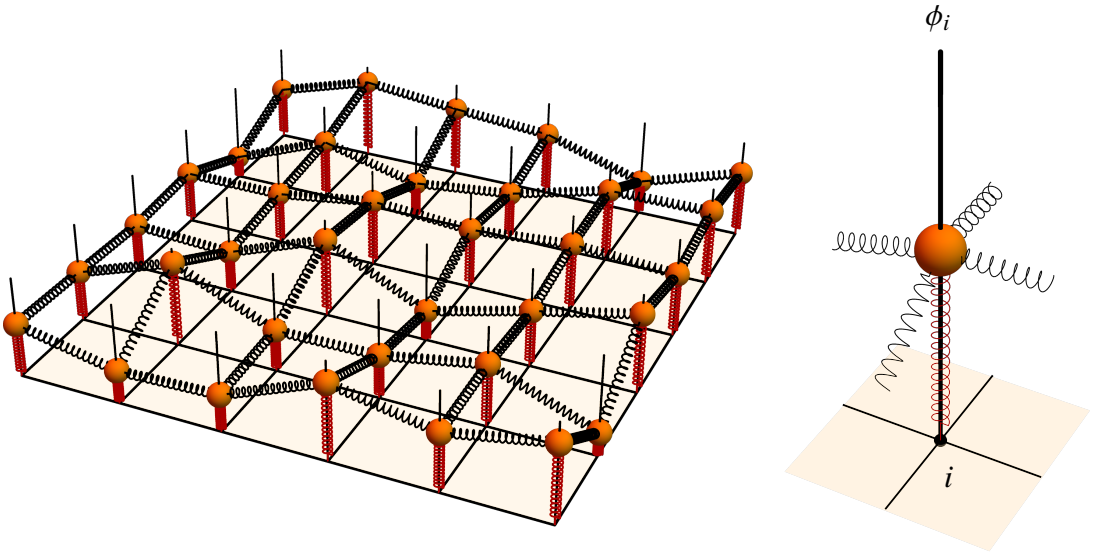
choose. Of course, since swapping the quantization axis orientation is a gauge transformations, it has no effect on the spectrum of the Hamiltonian since it corresponds to a choice of coordinates. With a bipartite parent lattice the choice of orientation is simply convenient to be taken uniformly throughout the lattice. Thus the 1-form description of the spin variables $S^z_e \equiv S^z(e)$ which change sign under reversal of the bond orientation (reversal of the quantization axis orientation) is in fact more intrinsic than treating the S^z as simply localized on sites of the line graph lattice.

3.3 Spin Models as Discrete Field Theories

In this section we will utilize the formalism developed so far to demonstrate that spin models can be considered discrete field theories. In particle, in the soft-spin approximation discussed in Chapter 2 the analogy between spin models and discrete field theories is exact. We show first that a square lattice Ising ferromagnet is equivalent to a 0-form field theory, then demonstrate that ice models map to 1-form field theories.

3.3.1 Discrete 0-Form (Scalar) Field Theory: Ising Ferromagnet

Now that we have discrete differential operators, we will use the SCGA to describe spin models as discrete field theories. The simplest field theory is a 0-form scalar field, which we denote ϕ. The



(a) (b)

Figure 3.9: (a) A configuration for discrete 0-form field on a 2D cellular complex. (b) A single fiber (internal space) of the 0-form field attached to a vertex i of the cell complex.

discrete field equivalent is a scalar-valued 0-cochain. An example of what we mean by a discrete 0-form field is shown in Fig. 3.9(a). The field is a collection of harmonic oscillators attached to the 0-cells (vertices) of cell complex (here a 2-dimensional square lattice), coupled together at nearest-neighbor separation. In (b) we show a single oscillator, which moves in a 1-dimensional "internal field space", or "fiber". Here the oscillators are represented by masses attached by springs, with each mass attached to its own vertex by a spring (red), and coupled to each neighbor by another spring (black). Accounting for only vertical displacement, the Hamiltonian for this system of coupled oscillators for this system can be written

$$H = \frac{1}{2}\sum_i m^2\phi_i^2 + \frac{1}{2}\sum_{ij} J\mathcal{A}_{ij}|\phi_j - \phi_i|^2, \tag{3.73}$$

where m^2 is the spring constant for the red springs and J is the spring constant for the coupling (black) springs. We recognize $\phi_j - \phi_i$ as the gradient $(\mathrm{d}\phi)_{ij}$. We can therefore express this Hamiltonian as a quadratic form in terms of the inner products on Ω^0 and Ω^1,

$$H = \frac{1}{2}m^2\langle\phi,\phi\rangle_0 + \frac{1}{2}J\langle\mathrm{d}\phi,\mathrm{d}\phi\rangle_1, \tag{3.74}$$

where subscripts indicate which space the inner product is in. This is the analog of the continuum Hamiltonian

$$H = \frac{1}{2}\int \mathrm{d}^3\boldsymbol{r}(J\nabla\phi\cdot\nabla\phi + m^2\phi^2). \tag{3.75}$$

Using the adjointness of d and δ, Eq. (3.30), we can express Eq. (3.74) as a single inner product using the Laplacian for scalar fields,

$$H = \frac{1}{2}\langle\phi, (m^2 + J\delta_1\mathrm{d}_0)\phi\rangle_0 = \frac{1}{2}\langle\phi, (m^2 + J\Delta_0)\phi\rangle. \tag{3.76}$$

Consider the equivalent for Eq. (3.75): integrating by parts the gradient we obtain

$$\int \mathrm{d}^3\boldsymbol{r}\nabla\phi\cdot\nabla\phi = \int \mathrm{d}^3\boldsymbol{r}\phi(-\nabla^2)\phi, \tag{3.77}$$

assuming either no boundary or that the fields go to zero sufficiently fast at infinity. This is the analogue of $\langle \mathrm{d}_0\phi, \mathrm{d}_0\phi\rangle_1 = \langle\phi, \delta_1\mathrm{d}_1\phi\rangle_0$. The eigenfunctions of the Laplacian are plane waves with wavevector $\boldsymbol{q}$, satisfying the eigenvalue equation

$$-\nabla^2 e^{i\boldsymbol{q}\cdot\boldsymbol{r}} = |\boldsymbol{q}|^2 e^{i\boldsymbol{q}\cdot\boldsymbol{r}}. \tag{3.78}$$

Fourier transforming Eq. (3.75), we diagonalize the integral kernel,

$$H = \frac{1}{2}\int \mathrm{d}^3\boldsymbol{q}\phi(-\boldsymbol{q})[m^2 + J\boldsymbol{q}^2]\phi(\boldsymbol{q}). \tag{3.79}$$

The integral kernel forms a quadratically dispersing band with minimum at m^2.[26] In the discrete case Eq. (3.76), the eigenvalues of Δ_0 will form N_{sl} bands, with one quadratically dispersing away from the minimum eigenvalue of m^2, with a finite bandwidth due to the periodicity of crystal momenta.

This discrete scalar field theory is precisely equivalent to the SCGA soft spin theory of of an Ising ($N = 1$) ferromagnet (cf. Section 2.4.1), if we treat the classical Ising soft spins as a 0-form field S. The SCGA Hamiltonian for a nearest-neighbor Ising ferromagnet is

$$H_{\mathrm{FM,SCGA}} = \frac{1}{2}\sum_{ij} S_i(\lambda T\delta_{ij} - J\mathcal{A}_{ij})S_j \tag{3.80}$$

$$\equiv \frac{1}{2}\sum_i \lambda_0 T|S_i|^2 + \sum_{ij} J\mathcal{A}_{ij}|S_j - S_i|^2 \tag{3.81}$$

$$\equiv \frac{1}{2}\langle S, (\lambda_0 T + J\Delta_0)S\rangle, \tag{3.82}$$

[26] In a relativistic theory then, m would correspond to the rest mass of the field quanta.

where in the second equality we absorbed the minimum eigenvalue of $\mathcal{A}$ into λ (see Sections 2.2.2 and 2.2.3). The nearest-neighbor Ising ferromagnet Hamiltonian is therefore equivalent to the discrete scalar field Hamiltonian Eq. (3.76), if we identify the mass-squared with $\lambda_0 T$. Thus we can view a lattice spin model with nearest-neighbor ferromagnetic interactions in the soft spin approximation as a 0-form field theory on a discrete cellular complex, where spins are localized on the vertices of the complex. The soft spin configuration space Ω_{spin} is then the space of 0-form configurations Ω^0. The quadratic spin Hamiltonian describes a free (Gaussian) discrete 0-form scalar field theory in the SCGA.

In the $N > 1$ case, the fiber or internal space of the field is an N-dimensional vector space in the soft spin approximation. The field would still be a 0-form field, i.e. the field degrees of freedom are still associated to the 0-cells of the complex, just with multiple components. When considering the hard-spin case, where the classical spins are treated as fixed-length vectors in the internal space, it is generally the case that the continuum analogue is taken to be a nonlinear σ model [131]. Such a field is still a 0-form field, i.e. the field still takes its value at 0-dimensional points, but the internal space of the field is taken to a be a manifold (a sphere) rather than a vector space.

3.3.2 Ice Models as Discrete 1-Form Field Theories

In the previous section we wrote down the Hamiltonian for a discrete free 0-form field theory, and concluded that it is equivalent to the SCGA model of a nearest-neighbor ferromagnet, where $\lambda_0 T$ plays the role of the mass squared. We will now demonstrate that ice models are naturally described as a discrete 1-form field theory. We will note here that this idea is not entirely novel, it is well known that quantum spin ice can be mapped through the ring exchange model to a lattice gauge field theory on the diamond lattice [25], and this method was extended through the introduction of the gauge mean field approach to quantum spin ice [130]. Our goal here is to clarify and develop the lattice field theory content of the *classical* theory utilizing the Hodge formalism discussed in this chapter.

Recall that in Section 2.5 we defined an ice model as a nearest-neighbor Ising antiferromagnet on a line graph lattice, whose Hamiltonian can be written in terms of the line graph adjacency matrix as in Eq. (2.83). Since every spin lives at the midpoint of an edge of the parent lattice, it is natural to take the quantization axis of the spin along the one dimensional edge, which is the local high-symmetry easy axis, which corresponds to the choice of an orientation for each bond. In the $N = 1$ Ising case, the resulting classical soft spin variable S_i^z is then a 1-form field on the parent lattice cellular complex. Utilizing Eq. (3.71), for the line graph of a *bipartite* parent lattice where we choose the canonical orientation with every edge oriented from the A to B sublattice of the

parent lattice, the ice model Hamiltonian can be written

$$H = \frac{J}{2}\sum_{ij} S_i \mathcal{A}^{\text{lg}}_{ij} S_j \tag{3.83}$$

$$= -2J\sum_{e} |S_e|^2 + \frac{J}{2}\sum_{ee'} S_e (\mathrm{d}_0\delta_1)_{ee'} S_{e'} \tag{3.84}$$

$$= E_0 + \frac{J}{2}\langle S, \mathrm{d}_0\delta_1 S\rangle, \tag{3.85}$$

where E_0 is the ground state energy in the spherical approximation, which is the true classical ground state energy when the frustrated units have an even number of spins (Section 2.4.2). Equivalently, the SCGA effective Hamiltonian for the ice models is

$$H_{\text{eff}} = \frac{1}{2}\langle S, (\lambda_0 + \mathrm{d}_0\delta_1) S\rangle. \tag{3.86}$$

Utilizing the adjointness relation, Eq. (3.30), we can equivalently write this as

$$H_{\text{eff}} = \frac{1}{2}\lambda_0\langle S, S\rangle_1 + \frac{1}{2}J\langle \delta_1 S, \delta_1 S\rangle_0, \tag{3.87}$$

where $\delta_1 S \sim -\mathrm{div} S$, Eq. (3.65). This is the discrete analogue of[27]

$$H_{\text{eff}} \sim \frac{1}{2}\int d^3\boldsymbol{r}\left(\lambda_0 |S|^2 + J(\mathrm{div} S)^2\right), \tag{3.88}$$

which is precisely the long-wavelength expansion of the Hamiltonian, Eq. (2.91), if we identify our 1-form field S with a vector field! This clarifies that the "emergent field theory" of a vector field obeying a zero-divergence constraint is not so surprising at all, it is simply a consequence of the fact that the underlying spin model describes a theory of a discrete 1-form field. Furthermore, if we expand Eq. (3.87) we can write it as

$$H_{\text{eff}} = \frac{1}{2}\lambda_0\langle S, S\rangle + \frac{1}{2}J\sum_i |(\delta_1 S)(v_i)|^2 \tag{3.89}$$

$$= \frac{1}{2}\lambda_0\langle S, S\rangle + \frac{1}{2}J\sum_i |S(\partial^\dagger v_i)|^2 \tag{3.90}$$

$$= \frac{1}{2}\lambda_0\langle S, S\rangle + \frac{1}{2}J\sum_i \Big|\sum_{j@i} S_{ji}\Big|^2, \tag{3.91}$$

[27] The precise continuum expression would be $\int \lambda_0 S\wedge\star S + J\delta_1 S\wedge\star\delta_1 S$, because the inner product for k-forms in the continuum is defined to be $\langle\alpha,\beta\rangle = \int \alpha\wedge\star\beta$, where the integral is over all space and $\alpha\wedge\star\beta$ is a top-dimensional (volume) form. We did not define our inner product using the Hodge star or define the wedge product for discrete forms, these are interesting topics that warrant further development but we will not do so here.

where the second term is a sum over the frustrated units of the line graph lattice, as in Eq. (2.82). Thus the frustration condition $\sum_{i\in c} S_i = 0$ which we discussed in Section 2.4.2 is equivalent to the zero-discrete-divergence condition $\delta_1 S = 0$. As such, ice models are naturally described as discrete 1-form field theories with Hamiltonian given by Eq. (3.87), where the discrete divergence $\delta_1 S$ is energetically suppressed and $\lambda_0 T$ describes the mass of the field. We identify the soft spin configuration space Ω_{spin} with the space of 1-forms Ω^1.

3.3.3 Helmholtz-Hodge Decomposition and Flat Bands: Rotational and Irrotational Modes

In Section 2.5.1 we discussed that the adjacency matrix of a line graph always has at least one flat band at the bottom of its spectrum. There we argued that these flat bands are a consequence of the number of constraints arising from each frustrated cluster of the lattice always being less than the number of spins, leaving an extensive SCGA ground state subspace. In this section we will understand this in terms of the HHD and the relation Eq. (3.71).

Since the spins are described as a discrete 1-form field S, we can apply the discrete HHD developed in Section 3.1.13, Section 3.1.14, and Section 3.1.15 to decompose them them as

$$S = \mathrm{d}_0\phi + \delta_2\alpha + \gamma, \tag{3.92}$$

where ϕ is a 0-form, α is a 2-form, and γ is harmonic 1-form. Equivalently, the space of soft spin configurations decomposes as

$$\Omega_{\text{spin}} \equiv \Omega^1 = \operatorname{im}\mathrm{d}_0 \oplus H^1 \oplus \operatorname{im}\delta_2 \tag{3.93}$$

We will adopt the term "irrotational" to refer to the $\mathrm{d}_0\phi$ portion of the decomposition Eq. (3.92), and "rotational" to refer to the $\delta_2\alpha$ portion, in accordance with the terminology used for the equivalent decomposition of vector fields. In particular, $\mathrm{d}_0\phi$ has zero curl, $\mathrm{d}_1(\mathrm{d}_0\phi) = 0$. This is equivalent to saying that $\mathrm{d}_0\phi$ integrates to zero on every closed cycle; that is, if C is a 1-cycle (a closed 1-dimensional contour), then

$$\mathrm{d}_0\phi(C) = \phi(\partial C) = \phi(0) = 0, \tag{3.94}$$

and therefore we say $\mathrm{d}_0\phi$ is an irrotational 1-form. On the other hand, the 2-form α represents a flux through each plaquette of the complex, and δ_2 maps this flux to a circulation around the boundary of each plaquette, thus guaranteeing that $\delta_2\alpha$ has zero divergence, i.e. $\delta_1(\delta_2\alpha) = 0$. Thus we can express Eq. (3.93) as

$$\Omega_{\text{spin}} \equiv \Omega^1 = \text{irrotational} \oplus \text{harmonic} \oplus \text{rotational}. \tag{3.95}$$

We will therefore also denote the decomposition Eq. (3.92) with the notation

$$S = S^{\text{irr}} + S^{\text{rot}} + S^{\text{har}}, \tag{3.96}$$

where

$$\mathrm{d}_1 S^{\text{irr}} \sim \operatorname{curl} S^{\text{irr}} = 0, \tag{3.97}$$

$$\delta_1 S^{\text{rot}} \sim \operatorname{div} S^{\text{rot}} = 0, \tag{3.98}$$

and the harmonic part has both zero divergence and zero curl.

We can immediately deduce using the HHD the meaning of the flat bands in the adjacency matrix for line graphs. Consider the bands of the adjacency matrix for the checkerboard, kagome, and pyrochlore line graph lattices with bipartite parent lattices, Figs. 2.11 and 2.12. In all cases there are exactly two dispersive bands, with the addition of flat bands at the bottom of the spectrum. Inserting the decomposition Eq. (3.92) into the Hamiltonian Eq. (3.86), using the fact that $\delta^2 = 0$, we see that α and γ drop out of the Hamiltonian entirely, i.e.

$$\langle S, \mathrm{d}_0\delta_1 S\rangle = \langle S, \mathrm{d}_0\delta_1(S^{\text{irr}} + S^{\text{rot}} + S^{\text{har}})\rangle = \langle S, \mathrm{d}_0\delta_1 S^{\text{irr}}\rangle = \langle S^{\text{irr}}, (\mathrm{d}_0\delta_1)S^{\text{irr}}\rangle. \tag{3.99}$$

This means that the rotational and harmonic components of the spin configuration have no energy cost other than the field stiffness cost coming from

$$\langle S, S\rangle = \langle S^{\text{irr}}, S^{\text{irr}}\rangle + \langle S^{\text{rot}}, S^{\text{rot}}\rangle + \langle S^{\text{har}}, S^{\text{har}}\rangle. \tag{3.100}$$

The absence of the rotational and harmonic modes in the Hamiltonian therefore corresponds to the presence of flat bands at the bottom of the adjacency matrix spectrum. In fact from Eq. (3.71) we can also deduce that these flat bands always have eigenvalue -2 since $\mathrm{d}_0\delta_1$ is positive semi-definite.[28] The subspace of Ω_{spin} spanned by the zero-energy flat band eigenvectors of the interaction matrix for ice models is therefore precisely the space of rotational and harmonic degrees of freedom. The dispersive band eigenvectors therefore must span the space of irrotational degrees of freedom. Therefore at zero temperature the irrotational modes are fully suppressed and all spin configurations obey the zero-divergence condition $\delta_1 S = 0$, but all rotational modes are completely degenerate, and we have a Coulomb phase.

To see this in some more detail, let's investigate the meaning of the two irrotational bands in Figs. 2.11 and 2.12 by considering the eigenvectors of $\mathrm{d}_0\delta_1$. First, note that the eigenvectors of the scalar Laplacian Δ_0 will be the discrete equivalent of plane waves. If ψ is an eigenvector of Δ_0 with eigenvalue ϵ, then

$$(\mathrm{d}_0\delta_1)(\mathrm{d}_0\psi) = \mathrm{d}_0(\delta_1\mathrm{d}_0)\psi = \mathrm{d}_0\Delta_0\psi = \epsilon(\mathrm{d}_0\psi), \tag{3.101}$$

[28] If we identify $\mathcal{A}^{\text{lg}} = \mathrm{d}_0\delta_1 - 2\mathbb{1}$, then since $\mathrm{d}_0\delta_1$ has minimum eigenvalue 0, $\mathcal{A}^{\text{lg}}$ has minimum eigenvalue -2.

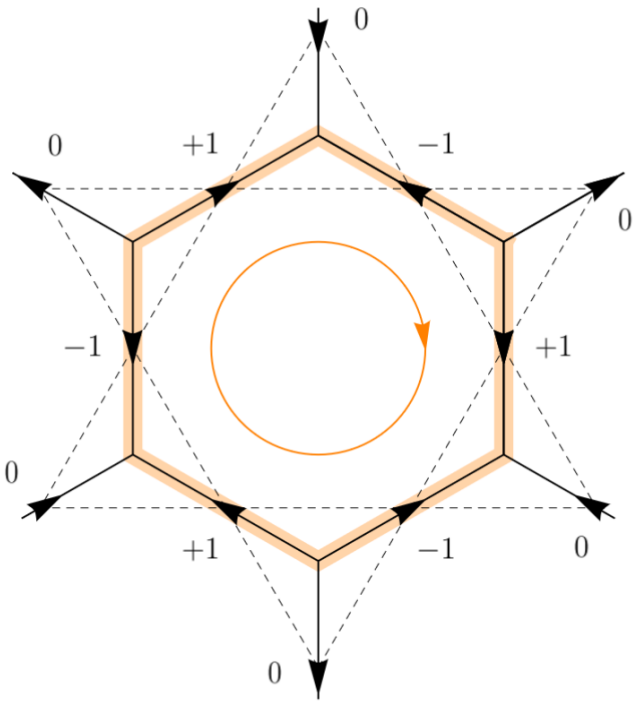


Figure 3.10: As a basis for the space of rotational modes we can take 1-forms $\delta_2 \alpha_p$, where $\alpha_p(p) = 1$ if p is oriented as shown by the orange arrow and 0 on all other plaquettes. Then $\delta_2 \alpha_p(e)$ on the boundary edges of p is ± 1 depending on the orientation taken for e and 0 on all other edges.

meaning that $\mathrm{d}_0\psi$ is an eigenvector of $\mathrm{d}_0\delta_1$ with the *same* eigenvalue. Therefore the spectrum of $\mathrm{d}_0\delta_1$ is exactly that of Δ_0, with the addition of flat bands with zero eigenvalue corresponding to the rotational modes. The eigenvectors of the dispersive bands are the gradients of the eigenvectors of Δ_0, of which there is one for each vertex of the parent lattice. The eigenvectors of the flat bands must be supported on 1-cycles since they correspond to rotational modes with zero divergence. These eigenvectors were investigated by Bergman et al. [103] in the context of hopping models on line graph lattices, who found that the flat band eigenvectors correspond to closed loops of lattice sites along which the sign of the wavefunction alternates, which we illustrate in Fig. 3.10 for the kagome lattice. We can reinterpret these alternating-sign wavefunctions as 1-forms representing a circulation around a single plaquette. In other words, we can choose a basis for the rotational modes given by $\{\delta_2\alpha_p\}$, where α_p is an "indicator 2-cochain" which represents a unit flux through oriented plaquette p, such that $\delta_2\alpha_p$ is an oriented circulation around the boundary ∂p. As Ref. [103] showed, however, this basis slightly overcounts the number of eigenvectors of the flat bands, because they are not all linearly independent. In particular, summing the value of $\delta_2\alpha_p$ over a non-trivial oriented surface that winds around a 2-dimensional hole yields zero, since such a surface is closed and has no boundary. Similarly, if we take as basis for the irrotational modes $\mathrm{d}_0\psi_i$, where ψ_i is non-zero on a single vertex i, we would overcount by one: summing over every i yields zero (every bond, having two endpoints, cancels). The remaining basis vectors are neither exact nor coexact, and they integrate to non-zero values on non-trivial 1-cycles of the complex, meaning that they are harmonic cochains which satisfy Laplace's equation for 1-forms,

$$\Delta_1\gamma = (\mathrm{d}_0\delta_1 + \delta_2\mathrm{d}_1)\gamma = 0. \tag{3.102}$$

The space of harmonic modes H^1 is therefore spanned by the eigenvectors of the band touching point in the spectrum of the line graph adjacency matrix.[29]

3.4 Spin Ice: 1-Form Field Theory on the Diamond Cell Complex

Consider the checkerboard, kagome, and pyrochlore lattices shown in Fig. 2.10. For the 2D line graphs it is easy to see that their parent lattices (square and hexagonal honeycomb) are cellular complexes. For the square lattice the 2-cells are squares and for the hexagonal honeycomb the 2-cells are hexagons. For the pyrochlore lattice the parent lattice is the diamond lattice, which is less obviously a cellular complex.

[29] It would be interesting for this reason to study ice models on cell complexes for which H^1 is trivial (e.g. on a sphere). Some work on this topic was done by Ref. [132] for a quantum spin-1/2 model on small complexes with a few frustrated clusters forming a polyhedron (topologically equivalent to a sphere).

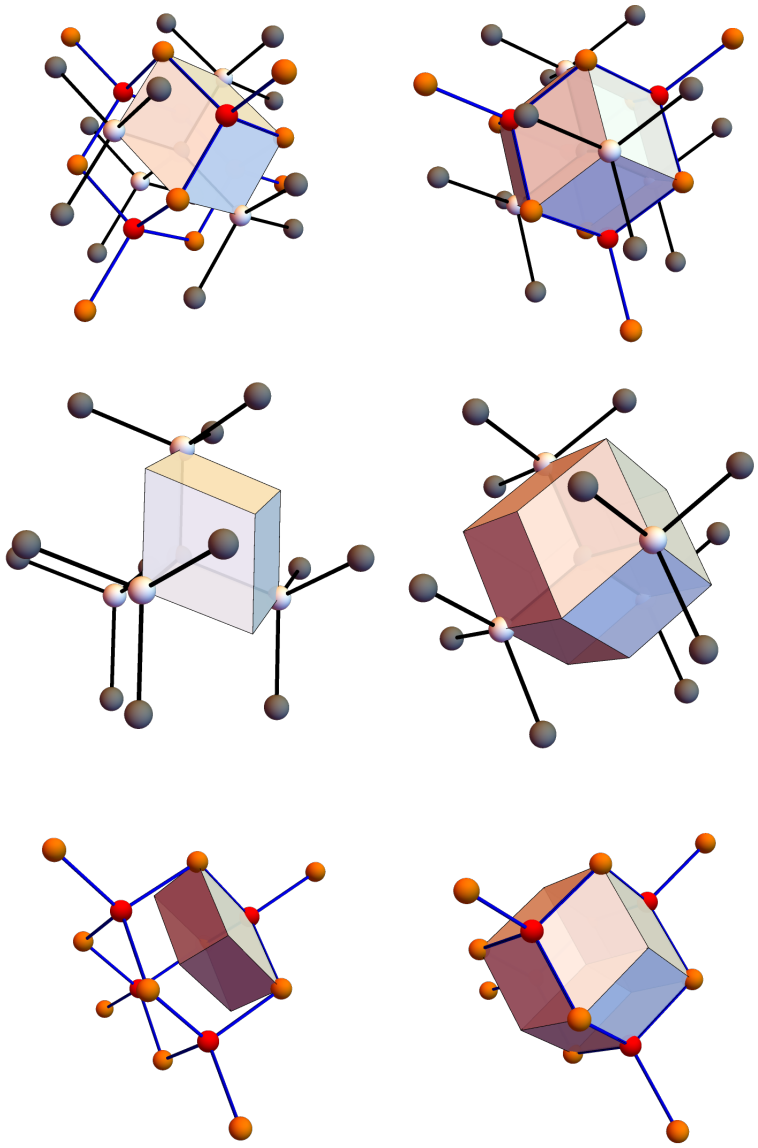


Figure 3.11: The bipartite diamond lattice is self-dual. Here we show the diamond lattice (red, orange spheres, blue edges) and its dual (gray, white spheres, blue edges). The diamond and its dual together form a BCC lattice. On the left we show a single BCC unit cell, and on the right we show four BCC unit cells which either surround a single diamond site or fill the diamond "cage." These four BCC cells together form the Wigner Seitz cell for the FCC lattice.

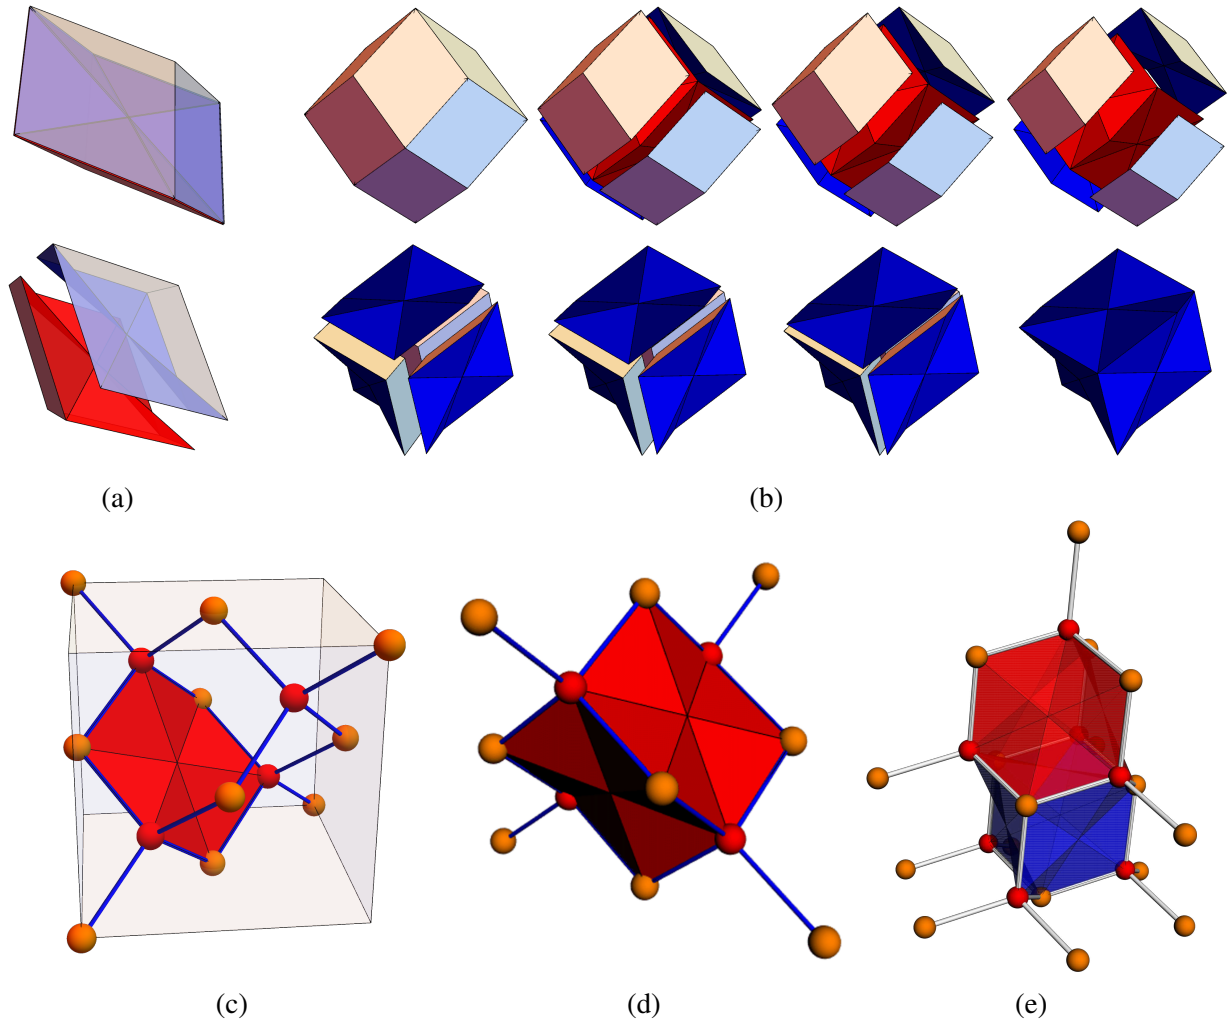


Figure 3.12: (a) The FCC wigner Seitz unit cell can be subdivided into two symmetric halves separated by a buckled hexagon (blue and red faces). (b) The Wigner Seitz cell cell splits into two equivalent cells (red, blue), cf. right hand column of Fig. 3.11. (c) The buckled hexagon forms a 2-cell of the diamond cell complex, shown here along with the 1-skeleton (blue edges, red and orange vertices) of the complex in a single cubic unit cell. (d) A single 3-cell of the diamond lattice cell complex, with four buckled hexagonal 2-cells on its boundary (two of which are not visible). (e) Two 3-cells of the diamond cellular complex formed by the separation shown in (a,b), which are inverted copies of each other. These two cells together for a unit cell for the FCC lattice, containing the same volume as the original Wigner Seitz cell.

Fig. 2.10 only shows the 1-skeleton of the diamond lattice cell complex. The geometry and topology of the complex is illustrated in more detail in Figs. 3.11 and 3.12. First, in Fig. 3.11, we show the bipartite diamond 1-skeleton (red, orange vertices and blue edges) as well as the dual 1-skeleton (gray, white vertices and black edges). The diamond lattice is self-dual, so its dual 1-skeleton is another diamond lattice. The diamond lattice has primitive description as an FCC lattice with a 2-atom basis; it is the three dimensional analog of the hexagonal honeycomb lattice, which is a triangular lattice with a 2-atom basis. Together, the diamond lattice and its dual form a BCC lattice, whose primitive unit cell is shown in the left column. In the right column we show that four diamond unit cells fit together to make the FCC Wigner Seitz cell, which fits perfectly inside the "cage" formed by the diamond lattice 1-skeleton (bottom right). In Fig. 3.12, we show in (a) that the BCC cell can be divided into two symmetric parts separated by a buckled hexagon. In (b) we show what happens when we divide the four BCC primitive cells forming the FCC Wigner Seitz cell. This divides the FCC cell into two identical cells, which is the diamond lattice 3-cell (red and blue). The 2-cells of the diamond lattice are the buckled hexagons, shown in (c) with respect to the 1-skeleton. The 3-cells, whose boundary is made of four buckled hexagons, is shown in (d), and in (e) we show how two of the diamond 3-cells fit together to form a primitive cell of the FCC lattice.

3.4.1 Eigenvectors of the Adjacency Matrix

Here we will demonstrate the exact calculation of the eigenvectors of the adjacency matrix of the pyrochlore lattice. The exact same calculation can be performed for any line graph on a bipartite parent lattice with translation symmetry. We note that a very similar calculation is done in [104] by defining "incidence matrices", adjacency matrices between the line graph vertices and parent lattice vertices. With hindsight we can observe that these incidence matrices are simply the direct space matrix elements of the linear maps d_0 and δ_1 (with appropriate choice of basis).

First we compute the eigenvectors of the diamond lattice Laplacian Δ_0. We introduce a complete basis set of plane wave 0-cochains denoted $|\boldsymbol{q}, \mu_0\rangle \in \Omega^0$, whose direct space components are

$$\langle \boldsymbol{r}, \mu_0 | \boldsymbol{q}, \nu_0 \rangle = \delta_{\mu_0,\nu_0} \frac{1}{\sqrt{L^3}} e^{-i\boldsymbol{q}\cdot(\boldsymbol{r}+\boldsymbol{c}_{\mu_0})}, \tag{3.103}$$

where $\mu_0, \nu_0 = A, B$ index the sublattices of the diamond lattice 0-skeleton and $\boldsymbol{c}_{\mu_0}$ is the sublattice position in the unit cell. These are plane waves which are non-zero only on a single sublattice. They are not themselves eigenvectors of the lattice Laplacian. We also will use the shorthand $|i\rangle_0$ for the position basis, referring to site i of the diamond lattice, where the subscript on the ket reminds us that this is a vector in Ω^0. Utilizing the relation between Δ_0 and the adjacency matrix $\mathcal{A}$ of the parent lattice, Eq. (3.69), noting that the diamond lattice has constant coordination, we

have $\Delta_0 = 4\mathbb{1} - \mathcal{A}$.[30] So we only need to find the eigenvectors of the adjacency of matrix of the diamond lattice. We have

$$\mathcal{A}\,|\boldsymbol{q}, A\rangle_0 = \frac{1}{\sqrt{L^3}} \sum_{i\in A} e^{-i\boldsymbol{q}\cdot\boldsymbol{r}_i} \mathcal{A}\,|i\rangle_0 \tag{3.104}$$

$$= \frac{1}{\sqrt{L^3}} \sum_{i\in A} e^{-i\boldsymbol{q}\cdot\boldsymbol{r}_i} \sum_j \mathcal{A}_{ij}\,|j\rangle_0 \tag{3.105}$$

$$= \frac{1}{\sqrt{L^3}} \sum_{i\in A} e^{-i\boldsymbol{q}\cdot\boldsymbol{r}_i} \sum_{j@i} |j\rangle_0 \tag{3.106}$$

In the final summation over the neighbors of i, notice that because the lattice is bipartite the neighbors of i are always on sublattice B. Therefore we can rewrite the two sums as a sum over sites in sublattice B and their neighbors,

$$\mathcal{A}\,|\boldsymbol{q}, A\rangle_0 = \frac{1}{\sqrt{L^3}} \sum_{i\in A} \sum_{j@i} e^{-i\boldsymbol{q}\cdot\boldsymbol{r}_i} |j\rangle_0 \tag{3.107}$$

$$= \frac{1}{\sqrt{L^3}} \sum_{j\in B} \sum_{i@j} e^{-i\boldsymbol{q}\cdot\boldsymbol{r}_i} |j\rangle_0 \tag{3.108}$$

$$= \frac{1}{\sqrt{L^3}} \sum_{j\in B} \left[\sum_{i@j} e^{-i\boldsymbol{q}\cdot(\boldsymbol{r}_i-\boldsymbol{r}_j)}\right] e^{-i\boldsymbol{q}\cdot\boldsymbol{r}_j} |j\rangle_0 \tag{3.109}$$

$$\equiv\equiv G(\boldsymbol{q}) \frac{1}{\sqrt{L^3}} \sum_{j\in B} e^{-i\boldsymbol{q}\cdot\boldsymbol{r}_j} |j\rangle_0 \tag{3.110}$$

$$= G(\boldsymbol{q})\,|\boldsymbol{q}, B\rangle_0 \tag{3.111}$$

where we defined

$$G(\boldsymbol{q}) \equiv \sum_{i@j} e^{-i\boldsymbol{q}\cdot(\boldsymbol{r}_i-\boldsymbol{r}_j)} \quad j \in B, \tag{3.112}$$

which is independent of j by translation symmetry. Indeed $G(\boldsymbol{q})$ can be written is

$$G(\boldsymbol{q}) = \sum_{\mu_1} e^{i\boldsymbol{q}\cdot\boldsymbol{\delta}_{\mu_1}}, \tag{3.113}$$

[30] Note that this formula is only meaningful under the convention that we take every 0-cell to have the same orientation. It is not strictly correct to treat $\mathcal{A}$ as a linear operator acting on Ω^0, because its matrix elements $\mathcal{A}_{ij} = 0, 1$ are not orientation-dependent, while those of Ω^0 are.

where $\boldsymbol{\delta}_{\mu_1}$ is the displacement vector from sublattice A to B of the diamond lattice corresponding to sublattice μ_1 of the pyrochlore lattice. Similarly we can show that

$$\mathcal{A}\,|\boldsymbol{q}, B\rangle_0 = G^*(\boldsymbol{q})\,|\boldsymbol{q}, A\rangle_0\,, \tag{3.114}$$

which follows from the inversion symmetry of sublattices A and B. The eigenvectors of $\mathcal{A}$ are then given by symmetric and antisymmetric linear combinations of $|\boldsymbol{q}, A\rangle$ and $|\boldsymbol{q}, B\rangle$, given up to normalization by

$$|\boldsymbol{q}, \pm\rangle_0 \propto \sqrt{G^*(\boldsymbol{q})}\,|\boldsymbol{q}, A\rangle_0 \pm \sqrt{G(\boldsymbol{q})}\,|\boldsymbol{q}, B\rangle_0\,, \tag{3.115}$$

which satisfy the eigenvalue equation

$$\mathcal{A}\,|\boldsymbol{q}, \pm\rangle_0 = \pm|G(\boldsymbol{q})|\,|\boldsymbol{q}, \pm\rangle_0\,. \tag{3.116}$$

The eigenvalues therefore come in positive-negative pairs, forming two bands symmetric about zero, with $G(\boldsymbol{0}) = 4$, the coordination number of the lattice. The scalar Laplacian $\Delta_0 = 4\mathbb{1} - \mathcal{A}$ has the same two bands but (i) inverted and (ii) lifted up so that the minimum eigenvalue is zero at $\boldsymbol{q} = \boldsymbol{0}$ and all other eigenvalues are positive, where the band containing the zero eigenvalue corresponds to $|\boldsymbol{q}, +\rangle_0$, the symmetric linear combination. The space H^0 of harmonic 0-forms, the kernel of Δ_0, is therefore the one-dimensional space spanned by the vector $|\boldsymbol{0}, +\rangle_0$. The components of $|\boldsymbol{0}, +\rangle_0$ are constant $1/\sqrt{2L^3}$ on every lattice site, corresponding to the fact that a harmonic scalar function is constant. On the other hand, $|\boldsymbol{0}, -\rangle$ alternates sign on the A and B sublattices.

The eigenvectors of the dispersive bands of $\mathcal{A}^{\mathrm{lg}}$, i.e. the eigenfunctions of $\mathrm{d}_0\delta_1$, are then given by $\mathrm{d}_0\,|\boldsymbol{q}, \pm\rangle_1 \in \Omega^1$. Each 1-cell can be indexed by unit cell position $\boldsymbol{r}$ and sublattice index μ. We define a basis for Ω^1 given by the vectors $|\boldsymbol{r}, \mu_1\rangle$, which is a 1-cochain which takes value +1 on edge μ in unit cell $\boldsymbol{r}$ when it is oriented from diamond sublattice A to B, -1 if it is oriented from B to A, and zero on all other 1-cells. We will also denote these basis cochains as $|ij\rangle_1$ for the cochain which is +1 on the edge oriented from i to j. We can conveniently express the differential and codifferential as[31]

$$\mathrm{d}_0 = \sum_{e_{ij}} |ij\rangle_1 \otimes ({}_0\langle j| - {}_0\langle i|\,), \tag{3.117}$$

$$\delta_1 = \sum_j |j\rangle_0 \otimes \sum_{i@j} {}_1\langle ij|\,. \tag{3.118}$$

[31] Note: the "kets" $|\tau_k\rangle_k$ are understood to be "indicator cochains" which are zero everywhere except on the oriented k-cell τ_k, on which they take the value +1. The "bras" $\langle\tau_k|_k$ should then be understood as the *evaluation functional*, which takes a cochain (a ket) and evaluates it on the oriented cell τ_k. Thus if α is a k-cochain, represented by a ket as $|\alpha\rangle$, then its value on a k-cell τ_k is $\langle\tau_k|\alpha\rangle \equiv \alpha(\tau_k)$.

where in the first line each edge is summed only once. Note that which orientation we pick for each edge is irrelevant, since $|ij\rangle_1 = -|ji\rangle_1$. We then define the plane wave basis for 1-forms by

$$|\boldsymbol{q}, \mu_1\rangle = \frac{1}{\sqrt{L^3}} \sum_{\boldsymbol{r}} e^{-i\boldsymbol{q}\cdot(\boldsymbol{r}+\boldsymbol{c}_{\mu_1})} |\boldsymbol{r}, \mu_1\rangle , \tag{3.119}$$

where we take $\boldsymbol{r} + \boldsymbol{c}_{\mu_1}$ to be the position of the midpoint of the edge, where the physical ion is located. Taking $i \in A$ and $j \in B$, the eigenvectors of $\mathrm{d}_0\delta_1$ have components

$$\begin{aligned}
{}_1\langle ij|\, \mathrm{d}_0 \,|\boldsymbol{q}, \pm\rangle_0 &= ({}_0\langle j| - {}_0\langle i|)\, |\boldsymbol{q}, \pm\rangle_0 \\
&= \left(\sqrt{G^*(\boldsymbol{q})}\, \langle j|\boldsymbol{q}, A\rangle \pm \sqrt{G(\boldsymbol{q})}\langle j|\boldsymbol{q}, B\rangle\right) - \left(\sqrt{G^*(\boldsymbol{q})}\, \langle i|\boldsymbol{q}, A\rangle \pm \sqrt{G(\boldsymbol{q})}\langle i|\boldsymbol{q}, B\rangle\right) \\
&= \frac{1}{\sqrt{L^3}} \left[\pm\sqrt{G(\boldsymbol{q})} e^{-i\boldsymbol{q}\cdot\boldsymbol{r}_j} - \sqrt{G^*(\boldsymbol{q})} e^{-i\boldsymbol{q}\cdot\boldsymbol{r}_i}\right] \\
&= \frac{1}{\sqrt{L^3}} e^{-i\boldsymbol{q}\cdot(\boldsymbol{r}_i+\boldsymbol{r}_j)/2} \underbrace{\left[\pm\sqrt{G(\boldsymbol{q})} e^{-i\boldsymbol{q}\cdot\boldsymbol{\delta}_{\mu_1}/2} - \sqrt{G^*(\boldsymbol{q})} e^{i\boldsymbol{q}\cdot\boldsymbol{\delta}_{\mu_1}/2}\right]}_{\equiv F_{\pm,\mu_1}(\boldsymbol{q})},
\end{aligned} \tag{3.120}$$

where in the first line we used Eq. (3.117), in the last line $(\boldsymbol{r}_i + \boldsymbol{r}_j)/2$ is the position of midpoint of the corresponding edge, and $\boldsymbol{\delta}_{\mu_1}$ is the displacement vector from sublattice A to B along edge μ_1. The eigenvectors can then be written up to normalization as

$$|\boldsymbol{q}, \pm\rangle_1 \propto \mathrm{d}_0 |\boldsymbol{q}, \pm\rangle_0 = \sum_{\mu} F_{\pm,\mu_1}(\boldsymbol{q})\, |\boldsymbol{q}, \mu_1\rangle_1 . \tag{3.121}$$

Note that at $\boldsymbol{q} = \boldsymbol{0}$, $F_{+,\mu_1}(\boldsymbol{0}) = 0$, because $\mathrm{d}_0 |\boldsymbol{0}, +\rangle$ is the gradient of a constant function, meaning that the eigenvector which would correspond to the base of the dispersive band touching the flat bands in $\mathrm{d}_0\delta_1$ is "missing". This is because the eigenvectors of the band touching point is the space H^1 of harmonic 1-forms, which cannot be expressed as the gradient of a 0-form. Consider on the other hand the eigenvector $|\boldsymbol{0}, -\rangle$, which is the gradient of a scalar function alternating sign on the two diamond sublattices. It is a constant 1-form which is positive for all edges oriented from the A to B sublattice, but it is not a harmonic form since it is exact. The remaining eigenvectors of $\mathrm{d}_0\delta_1$ are the rotational modes, which span $\operatorname{im}\delta_2$, forming the two flat bands with zero eigenvalue. The harmonic modes correspond to the band touching point, which are the modes that cannot be expressed as a gradient. To summarize, the HHD is a decomposition of the spin configuration space Ω^1 into three subspaces,

$$\Omega^1 = \operatorname{im}\mathrm{d}_0 \oplus H^1 \oplus \operatorname{im}\delta_2, \tag{3.122}$$

where $\operatorname{im}\mathrm{d}_0$ is the space of irrotational modes, spanned by the dispersive band eigenvectors of $\mathcal{A}^{\mathrm{lg}}$, $\operatorname{im}\delta_2$ is the space of rotational modes, spanned by the eigenvectors of the flat bands, and H^1 is the space of harmonic modes, corresponding to the band touching point at the zone center.

3.4.2 HHD Projectors

We now turn to the explicit computation of the HHD for spin configurations. Surprisingly, the importance of the HHD for describing spin ice physics was not appreciated until quite recently, and interest in the HHD has been primarily confined to the literature surrounding spin ice fragmentation [133, 134], the phenomenon by which the irrotational modes partially condense while the rotational modes are free to continue fluctuating. In particular, Brooks-Bartlett et al. [133] demonstrated the decomposition for a single unit cell (or in their account, a single tetrahedron). A partial decomposition via an iterative algorithm was also used by [135], but until now an explicit method for the computation on the full lattice has not been demonstrated. Using the Hodge approach, we will define the HHD unambiguously for *any* spin configuration, in a way that allows for a straightforward computation.

We demonstrated the general case in Section 3.1.15, from which we concluded that the projector to the irrotational and rotatational mode subspaces are

$$\mathcal{P}^{\operatorname{im} \mathrm{d}_0} = \mathrm{d}_0(\Delta_0)^{-1}\delta_1 \qquad \text{(irrotational)}, \tag{3.123}$$

$$\mathcal{P}^{\operatorname{im} \delta_2} = \mathrm{d}_2(\Delta_2)^{-1}\mathrm{d}_1 \qquad \text{(rotational)}. \tag{3.124}$$

We also know now that the projector to the harmonic modes is just the projector to the three-dimensional band touching point eigenspace at $\boldsymbol{q} = \mathbf{0}$. This projector is easy to construct by directly diagonalizing the $\boldsymbol{q} = \mathbf{0}$ adjacency matrix, which is given by[32]

$$\langle \mathbf{0}, \mu_1 | \, \mathcal{A}^{\mathrm{lg}} \, | \mathbf{0}, \nu_1 \rangle = 2 \begin{pmatrix} 0 & 1 & 1 & 1 \\ 1 & 0 & 1 & 1 \\ 1 & 1 & 0 & 1 \\ 1 & 1 & 1 & 0 \end{pmatrix}, \tag{3.125}$$

As we discussed in the previous chapter (cf. Section 2.4.2), this matrix has one eigenvector $|\mathbf{0}, Q\rangle_1 \equiv (1/2) \sum_{\mu_1} |\mathbf{0}, \mu_1\rangle$ which corresponds to the band maximum and the Q irrep eigenmode, $Q(\mathbf{0}) \equiv \langle \mathbf{0}, Q | S \rangle$ (Eq. (2.88)). The three remaining eigenvectors in the orthogonal three dimensional subspace correspond to the three $\boldsymbol{B}$ eigenmodes (Eq. (2.89)). The three $\boldsymbol{B}$ eigenmodes are therefore the harmonic modes and, in spin ice, correspond to the three components of the bulk magnetization. On a torus, they correspond to spin configurations which are polarized around the handles of the torus. Thus we can simply write the harmonic mode projector as

$$\mathcal{P}^{H^1} = \mathbb{1}_{\mathbf{0}} - |\mathbf{0}, Q\rangle \langle \mathbf{0}, Q| \,, \quad \text{(harmonic)} \tag{3.126}$$

[32] As we mentioned for the adjacency matrix of the diamond lattice, we note that $\mathcal{A}^{\mathrm{lg}}$ is not really a linear operator on Ω^1, and this formula assumes the global A to B orientation of edges.

where $\mathbb{1}_{\mathbf{0}} = \sum_\mu |\mathbf{0}, \mu\rangle \langle \mathbf{0}, \mu|$ is the identity in the $\boldsymbol{q} = \mathbf{0}$ subspace and zero everywhere else.

We will now compute the projector in Eq. (3.123) in the Fourier basis. To begin, we consider the Fourier representation of the operators d_0 and δ_1, which can be easily obtained from Eqs. (3.117) and (3.118). Starting with d_0, we have

$$\langle \boldsymbol{q}, \mu_1 | \, \mathrm{d}_0 \, | \boldsymbol{q}, \mu_0 \rangle = \sum_{e_{ij}} \langle \boldsymbol{q}, \mu_1 | ij \rangle \left(\langle j | \boldsymbol{q}, \mu_0 \rangle - \langle i | \boldsymbol{q}, \mu_0 \rangle \right) \tag{3.127}$$

$$= \frac{1}{L^3} \sum_{e_{ij} \in \mu_1} e^{i\boldsymbol{q}\cdot(\boldsymbol{r}_i + \boldsymbol{r}_j)/2} \left(e^{-i\boldsymbol{q}\cdot\boldsymbol{r}_j} \delta_{j \in \mu_0} - e^{-i\boldsymbol{q}\cdot\boldsymbol{r}_i} \delta_{i \in \mu_0} \right), \tag{3.128}$$

where $\delta_{i\in\mu_0} = 1$ if vertex i is in sublattice μ_0 and 0 otherwise, and for the sum over edges e_{ij} is oriented from A to B, since we defined $|\boldsymbol{q}, \mu_1\rangle$ in Eq. (3.119) with all edges oriented from A to B. We then have

$$\langle \boldsymbol{q}, \mu_1 | \, \mathrm{d}_0 \, | \boldsymbol{q}, \mu_0 \rangle = \frac{1}{L^3} \sum_{e_{ij} \in \mu_1} \delta_{j \in \mu_0} e^{i\boldsymbol{q}\cdot(\boldsymbol{r}_i - \boldsymbol{r}_j)/2} - \frac{1}{L^3} \sum_{e_{ij} \in \mu_1} \delta_{i \in \mu_0} e^{i\boldsymbol{q}\cdot(\boldsymbol{r}_j - \boldsymbol{r}_i)/2}. \tag{3.129}$$

Since the sum is over edges oriented A to B in a single sublattice μ_1, we have that $\boldsymbol{r}_j - \boldsymbol{r}_i \equiv \boldsymbol{\delta}_{\mu_1}$, and we obtain

$$\langle \boldsymbol{q}, \mu_1 | \, \mathrm{d}_0 \, | \boldsymbol{q}, \mu_0 \rangle = \delta_{\mu_0 = B} \, e^{-i\boldsymbol{q}\cdot\boldsymbol{\delta}_{\mu_1}/2} - \delta_{\mu_0 = A} \, e^{i\boldsymbol{q}\cdot\boldsymbol{\delta}_{\mu_1}/2}, \tag{3.130}$$

which can be expressed in matrix form as[33]

$$\begin{pmatrix} \langle \boldsymbol{q}, 1 | \, \mathrm{d}_0 \, | \boldsymbol{q}, A \rangle & \langle \boldsymbol{q}, 1 | \, \mathrm{d}_0 \, | \boldsymbol{q}, B \rangle \\ \langle \boldsymbol{q}, 2 | \, \mathrm{d}_0 \, | \boldsymbol{q}, A \rangle & \langle \boldsymbol{q}, 2 | \, \mathrm{d}_0 \, | \boldsymbol{q}, B \rangle \\ \langle \boldsymbol{q}, 3 | \, \mathrm{d}_0 \, | \boldsymbol{q}, A \rangle & \langle \boldsymbol{q}, 3 | \, \mathrm{d}_0 \, | \boldsymbol{q}, B \rangle \\ \langle \boldsymbol{q}, 4 | \, \mathrm{d}_0 \, | \boldsymbol{q}, A \rangle & \langle \boldsymbol{q}, 4 | \, \mathrm{d}_0 \, | \boldsymbol{q}, B \rangle \end{pmatrix} = \begin{pmatrix} -e^{i\boldsymbol{q}\cdot\boldsymbol{\delta}_1/2} & e^{-i\boldsymbol{q}\cdot\boldsymbol{\delta}_1/2} \\ -e^{i\boldsymbol{q}\cdot\boldsymbol{\delta}_2/2} & e^{-i\boldsymbol{q}\cdot\boldsymbol{\delta}_2/2} \\ -e^{i\boldsymbol{q}\cdot\boldsymbol{\delta}_3/2} & e^{-i\boldsymbol{q}\cdot\boldsymbol{\delta}_3/2} \\ -e^{i\boldsymbol{q}\cdot\boldsymbol{\delta}_4/2} & e^{-i\boldsymbol{q}\cdot\boldsymbol{\delta}_4/2} \end{pmatrix}. \tag{3.131}$$

[33] We have always taken the convention that all 0-cells have the same global "in" orientation. We could have taken the orientations to be "out" on the A sublattice and "in" on the B sublattice. This would remove the negative signs in the first column of Eq. (3.131). This matrix can also be found, without the minus signs, in [57], where it is defined to be the Fourier transform of an adjacency matrix whose matrix elements are +1 for neighboring diamond-pyrochlore sites and 0 otherwise. This adjacency matrix is in fact precisely $\langle e | \, \mathrm{d}_0 \, | i \rangle$ (cf. Eq. (3.117)) if we take the convention that sublattice A 0-cells are oriented "out", B 0-cells are oriented "in", and edges are always oriented from A to B. We note also that this matrix expression first appeared (implicitly) in the work of Henley [136] in the study of the soft spin ground state correlations of spin ice, which we will discuss further in this chapter, along with the matrix form of the Fourier transformed Laplacian Eq. (3.139), but none of the connection to discrete differential operators or lattice field theory was known at the time.

Doing the same computation for δ_1, using Eq. (3.118), we obtain

$$\langle \boldsymbol{q}, \mu_0 |\, \delta_1 \,| \boldsymbol{q}, \mu_1 \rangle = \sum_j \sum_{i@j} \langle \boldsymbol{q}, \mu_0 | j \rangle \, \langle ij | \boldsymbol{q}, \mu_1 \rangle \,, \tag{3.132}$$

$$= \frac{1}{L^3} \sum_j \sum_{i@j} \left(\delta_{j \in \mu_0} \, e^{i \boldsymbol{q} \cdot \boldsymbol{r}_j} \right) \left(\delta_{e_{ij} \in \mu_1} \, \left(\delta_{j \in B} - \delta_{j \in A} \right) \, e^{-i \boldsymbol{q} \cdot (\boldsymbol{r}_i + \boldsymbol{r}_j)/2} \right), \tag{3.133}$$

where in the second line we used the fact that $|\boldsymbol{q}, \mu_1\rangle$ was defined with the convention that the sign is positive if the edge is oriented from A to B. Utilizing the Kronecker deltas, we can rewrite this as a sum over edges in sublattice μ_1 with an overall sign determined by μ_0,

$$\langle \boldsymbol{q}, \mu_0 |\, \delta_1 \,| \boldsymbol{q}, \mu_1 \rangle = (\delta_{\mu_0 = B} - \delta_{\mu_0 = A}) \frac{1}{L^3} \sum_{e_{ij} \in \mu_1, \, j \in \mu_0} e^{-i \boldsymbol{q} \cdot (\boldsymbol{r}_i - \boldsymbol{r}_j)/2} \tag{3.134}$$

$$= \delta_{\mu_0 = B} \, e^{i \boldsymbol{q} \cdot \boldsymbol{\delta}_{\mu_1}/2} - \delta_{\mu_0 = A} \, e^{-i \boldsymbol{q} \cdot \boldsymbol{\delta}_{\mu_1}/2}. \tag{3.135}$$

We conclude that

$$\begin{pmatrix} \langle \boldsymbol{q}, A |\, \delta_1 \,| \boldsymbol{q}, 1 \rangle & \langle \boldsymbol{q}, A |\, \delta_1 \,| \boldsymbol{q}, 2 \rangle & \langle \boldsymbol{q}, A |\, \delta_1 \,| \boldsymbol{q}, 3 \rangle & \langle \boldsymbol{q}, A |\, \delta_1 \,| \boldsymbol{q}, 4 \rangle \\ \langle \boldsymbol{q}, B |\, \delta_1 \,| \boldsymbol{q}, 1 \rangle & \langle \boldsymbol{q}, B |\, \delta_1 \,| \boldsymbol{q}, 2 \rangle & \langle \boldsymbol{q}, B |\, \delta_1 \,| \boldsymbol{q}, 3 \rangle & \langle \boldsymbol{q}, B |\, \delta_1 \,| \boldsymbol{q}, 4 \rangle \end{pmatrix}$$

$$=$$

$$\begin{pmatrix} -e^{-i \boldsymbol{q} \cdot \boldsymbol{\delta}_1/2} & -e^{-i \boldsymbol{q} \cdot \boldsymbol{\delta}_2/2} & -e^{-i \boldsymbol{q} \cdot \boldsymbol{\delta}_3/2} & -e^{-i \boldsymbol{q} \cdot \boldsymbol{\delta}_4/2} \\ e^{i \boldsymbol{q} \cdot \boldsymbol{\delta}_1/2} & e^{i \boldsymbol{q} \cdot \boldsymbol{\delta}_2/2} & e^{i \boldsymbol{q} \cdot \boldsymbol{\delta}_3/2} & e^{i \boldsymbol{q} \cdot \boldsymbol{\delta}_4/2} \end{pmatrix}. \tag{3.136}$$

Comparing this expression to Eq. (3.130) and Eq. (3.131), we see that the matrix elements of δ_1 are just the conjugate transpose of the matrix elements of d_0, which they had to be since these operators were defined to be adjoint with respect to the inner product (Eq. (3.30)).

Finally, we can compute the Fourier basis matrix elements of the scalar Laplacian, $\Delta_0 = \delta_1 \mathrm{d}_0$, by simply multiplying the corresponding matrices,

$$\langle \boldsymbol{q}, \mu_0 |\, \Delta_0 \,| \boldsymbol{q}, \nu_0 \rangle = \sum_{\mu_1} \langle \boldsymbol{q}, \mu_0 |\, \delta_1 \,| \boldsymbol{q}, \mu_1 \rangle \, \langle \boldsymbol{q}, \mu_1 |\, \mathrm{d}_0 \,| \boldsymbol{q}, \nu_0 \rangle \tag{3.137}$$

$$= \sum_{\mu_1} \left(\delta_{\mu_0 = B} \, e^{i \boldsymbol{q} \cdot \boldsymbol{\delta}_{\mu_1}/2} - \delta_{\mu_0 = A} \, e^{-i \boldsymbol{q} \cdot \boldsymbol{\delta}_{\mu_1}/2} \right) \left(\delta_{\nu_0 = B} \, e^{-i \boldsymbol{q} \cdot \boldsymbol{\delta}_{\mu_1}/2} - \delta_{\nu_0 = A} \, e^{i \boldsymbol{q} \cdot \boldsymbol{\delta}_{\mu_1}/2} \right) \tag{3.138}$$

$$= \begin{pmatrix} 4 & -G^*(\boldsymbol{q}) \\ -G(\boldsymbol{q}) & 4 \end{pmatrix}. \tag{3.139}$$

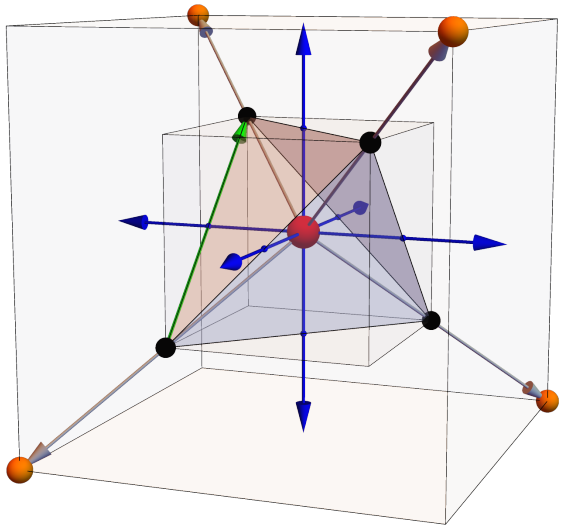


Figure 3.13: A single diamond A sublattice site (red) and its four neighboring B sublattice sites (orange). The four displacement vectors $\boldsymbol{\delta}_\mu$ are shown in gray. The four pyrochlore sites sitting at the midpoints of the diamond bonds are shown as black spheres, forming a tetrahedron surrounding the central diamond site. The vectors $(\boldsymbol{\delta}_{\mu_1} \pm \boldsymbol{\delta}_{\nu_1})/2$ which appear in the matrix elements of the irrotational projector Eq. (3.146) are shown, with the positive sign corresponding to the blue vectors pointing to the cube face centers and the negative sign corresponding to the pyrochlore separation vectors along the edges of the tetrahedron, with one example vector shown in green.

One can readily verify that the eigenvectors of this matrix agree with $|\boldsymbol{q},\pm\rangle$ defined in Eq. (3.115) and the eigenvalues are $4 \pm |G(\boldsymbol{q})|$.

Finally, we can write down the projector to the irrotational subspace, Eq. (3.123). The inverse[34] of the Laplacian is readily computed in the basis of Eq. (3.139) to be

$$\langle \boldsymbol{q},\mu_0|\,(\Delta_0)^{-1}\,|\boldsymbol{q},\nu_0\rangle = \frac{1}{16-|G(\boldsymbol{q})|^2}\begin{pmatrix}4 & G^*(\boldsymbol{q})\\ G(\boldsymbol{q}) & 4\end{pmatrix}, \tag{3.140}$$

and the matrix elements of the projector are

$$\begin{aligned}
\mathcal{P}^{\mathrm{im}\,\mathrm{d}_0}_{\mu_1\nu_1}(\boldsymbol{q}) &\equiv \langle \boldsymbol{q},\mu_1|\,\mathrm{d}_0(\Delta_0)^{-1}\delta_1\,|\boldsymbol{q},\nu_1\rangle && (3.141)\\
&= \sum_{\mu_0,\nu_0}\langle \boldsymbol{q},\mu_1|\,\mathrm{d}_0\,|\boldsymbol{q},\mu_0\rangle\,\langle \boldsymbol{q},\mu_0|\,(\Delta_0)^{-1}\,|\boldsymbol{q},\nu_0\rangle\,\langle \boldsymbol{q},\nu_0|\,\delta_1\,|\boldsymbol{q},\nu_1\rangle && (3.142)\\
&= \frac{1}{16-|G(\boldsymbol{q})|^2}\begin{pmatrix}-e^{i\boldsymbol{q}\cdot\boldsymbol{\delta}_{\mu_1}/2} & e^{-\boldsymbol{q}\cdot\boldsymbol{\delta}_{\mu_1}/2}\end{pmatrix}\begin{pmatrix}4 & G^*(\boldsymbol{q})\\ G(\boldsymbol{q}) & 4\end{pmatrix}\begin{pmatrix}-e^{-i\boldsymbol{q}\cdot\boldsymbol{\delta}_{\nu_1}/2}\\ e^{i\boldsymbol{q}\cdot\boldsymbol{\delta}_{\nu_1}/2}\end{pmatrix} && (3.143)\\
&= \frac{1}{16-|G(\boldsymbol{q})|^2}\begin{pmatrix}-e^{i\boldsymbol{q}\cdot\boldsymbol{\delta}_{\mu_1}/2} & e^{-\boldsymbol{q}\cdot\boldsymbol{\delta}_{\mu_1}/2}\end{pmatrix}\begin{pmatrix}-4e^{-i\boldsymbol{q}\cdot\boldsymbol{\delta}_{\nu_1}/2}+G^*(\boldsymbol{q})e^{i\boldsymbol{q}\cdot\boldsymbol{\delta}_{\nu_1}/2}\\ 4e^{i\boldsymbol{q}\cdot\boldsymbol{\delta}_{\nu_1}/2}-G(\boldsymbol{q})e^{-i\boldsymbol{q}\cdot\boldsymbol{\delta}_{\nu_1}/2}\end{pmatrix} && (3.144)
\end{aligned}$$

Multiplying this out, we obtain

$$(16-|G(\boldsymbol{q})|^2)P_{\mu_1\nu_1}(\boldsymbol{q}) = \left[4e^{i\boldsymbol{q}\cdot(\boldsymbol{\delta}_{\mu_1}-\boldsymbol{\delta}_{\nu_1})/2} - G^*(\boldsymbol{q})e^{i\boldsymbol{q}\cdot(\boldsymbol{\delta}_{\mu_1}+\boldsymbol{\delta}_{\nu_1})/2}\right] + \text{h.c.} \tag{3.145}$$

Expanding $G(\boldsymbol{q}) = \sum_{\rho_1}\exp\bigl(i\boldsymbol{q}\cdot\boldsymbol{\delta}_{\rho_1}\bigr)$, we obtain

$$\mathcal{P}^{\mathrm{im}\,\mathrm{d}_0}_{\mu_1\nu_1}(\boldsymbol{q}) = \frac{1}{16-|G(\boldsymbol{q})|^2}\left[8\cos\left(\frac{\boldsymbol{\delta}_{\mu_1}-\boldsymbol{\delta}_{\nu_1}}{2}\cdot\boldsymbol{q}\right) - \sum_{\rho_1}2\cos\left(\left(\boldsymbol{\delta}_{\rho_1}-\frac{\boldsymbol{\delta}_{\mu_1}+\boldsymbol{\delta}_{\nu_1}}{2}\right)\cdot\boldsymbol{q}\right)\right]. \tag{3.146}$$

Notice that this is ill-defined at $\boldsymbol{q}=\boldsymbol{0}$, where $G(\boldsymbol{q})=4$.

We can Taylor expand the matrix elements to see the long-wavelength behavior. Starting with the denominator, we have

$$16 - \left|\sum_{\rho_1}e^{i\boldsymbol{q}\cdot\boldsymbol{\delta}_{\rho_1}}\right|^2 \approx 16 - \left|4-\frac{1}{2}(\boldsymbol{\delta}_{\rho_1}\cdot\boldsymbol{q})^2\right|^2 = 4\sum_{\rho_1}(\boldsymbol{\delta}_{\rho_1}\cdot\boldsymbol{q})^2. \tag{3.147}$$

[34] Here we mean the Moore-Penrose pseudo-inverse, which means that we only invert the linear map in the subspace where it has non-zero eigenvalues, in this case for all $\boldsymbol{q}\neq\boldsymbol{0}$.

Expanding the cosines, we have

$$\mathcal{P}^{\mathrm{im\,d_0}}_{\mu_1\nu_1}(\boldsymbol{q}) \approx \frac{1}{4\sum_{\rho_1}(\boldsymbol{\delta}_{\rho_1}\cdot\boldsymbol{q})^2}\left[-4\left(\frac{\boldsymbol{\delta}_{\mu_1}-\boldsymbol{\delta}_{\nu_1}}{2}\cdot\boldsymbol{q}\right)^2+\sum_{\rho_1}\left(\left(\boldsymbol{\delta}_{\rho_1}-\frac{\boldsymbol{\delta}_{\mu_1}+\boldsymbol{\delta}_{\nu_1}}{2}\right)\cdot\boldsymbol{q}\right)^2\right] \tag{3.148}$$

$$= \frac{1}{4\sum_{\rho_1}(\boldsymbol{\delta}_{\rho_1}\cdot\boldsymbol{q})^2}\left[-4\left(\frac{\boldsymbol{\delta}_{\mu_1}-\boldsymbol{\delta}_{\nu_1}}{2}\cdot\boldsymbol{q}\right)^2+\sum_{\rho_1}(\boldsymbol{\delta}_{\rho_1}\cdot\boldsymbol{q})^2+4\left(\frac{\boldsymbol{\delta}_{\mu_1}+\boldsymbol{\delta}_{\nu_1}}{2}\cdot\boldsymbol{q}\right)^2\right] \tag{3.149}$$

$$= \frac{1}{4\sum_{\rho_1}(\boldsymbol{\delta}_{\rho_1}\cdot\boldsymbol{q})^2}\left[\sum_{\rho_1}(\boldsymbol{\delta}_{\rho_1}\cdot\boldsymbol{q})^2+4(\boldsymbol{\delta}_{\mu_1}\cdot\boldsymbol{q})(\boldsymbol{\delta}_{\nu_1}\cdot\boldsymbol{q})\right] \tag{3.150}$$

$$= \frac{1}{4}+\frac{(\boldsymbol{\delta}_{\mu_1}\cdot\boldsymbol{q})(\boldsymbol{\delta}_{\nu_1}\cdot\boldsymbol{q})}{\sum_{\rho_1}(\boldsymbol{\delta}_{\rho_1}\cdot\boldsymbol{q})^2}. \tag{3.151}$$

$$= \frac{1}{4}+\frac{(\hat{\boldsymbol{\delta}}_{\mu_1}\cdot\hat{\boldsymbol{q}})(\hat{\boldsymbol{\delta}}_{\nu_1}\cdot\hat{\boldsymbol{q}})}{\sum_{\rho_1}(\hat{\boldsymbol{\delta}}_{\rho_1}\cdot\hat{\boldsymbol{q}})^2}. \tag{3.152}$$

The denominator of the second term is actually a constant, because $\sum_{\rho_1}\hat{\boldsymbol{\delta}}_{\rho_1}\hat{\boldsymbol{\delta}}^T_{\rho_1} = (4/3)\mathbb{1}$ (the $\boldsymbol{\delta}_\mu$ form an overcomplete basis for $\mathbb{R}^3$), so we obtain

$$\mathcal{P}^{\mathrm{im\,d_0}}_{\mu_1\nu_1}(\boldsymbol{q}) \approx \frac{1}{4}+\frac{3}{4}(\hat{\boldsymbol{\delta}}_{\mu_1}\cdot\hat{\boldsymbol{q}})(\hat{\boldsymbol{\delta}}_{\nu_1}\cdot\hat{\boldsymbol{q}}). \tag{3.153}$$

We can therefore write down the long-wavelength form of projector

$$\mathcal{P}^{\mathrm{im\,d_0}}(\boldsymbol{q}) \approx \sum_{\mu_1,\nu_1}|\boldsymbol{q},\mu_1\rangle\left(\frac{1}{4}+\frac{3}{4}(\hat{\boldsymbol{\delta}}_{\mu_1}\cdot\hat{\boldsymbol{q}})(\hat{\boldsymbol{\delta}}_{\nu_1}\cdot\hat{\boldsymbol{q}})\right)\langle\boldsymbol{q},\nu_1|\,, \tag{3.154}$$

$$= \left(\frac{1}{2}\sum_{\mu_1}|\boldsymbol{q},\mu_1\rangle\right)\left(\frac{1}{2}\sum_{\nu_1}\langle\boldsymbol{q},\nu_1|\right)+\left(\frac{\sqrt{3}}{2}\sum_{\mu_1}(\hat{\boldsymbol{\delta}}_{\mu_1}\cdot\hat{\boldsymbol{q}})\,|\boldsymbol{q},\mu_1\rangle\right)\left(\frac{\sqrt{3}}{2}\sum_{\nu_1}(\hat{\boldsymbol{\delta}}_{\nu_1}\cdot\hat{\boldsymbol{q}})\,\langle\boldsymbol{q},\nu_1|\right), \tag{3.155}$$

From which we read off the long-wavelength normalized eigenvectors of the irrotational bands, which we previously obtained in the form of Eq. (3.121), as

$$|\boldsymbol{q},-\rangle_1 = \frac{1}{2}\sum_{\mu_1}|\boldsymbol{q},\mu_1\rangle\,, \tag{3.156}$$

$$|\boldsymbol{q},+\rangle_1 = \frac{\sqrt{3}}{2}\sum_{\mu_1}(\hat{\boldsymbol{\delta}}_{\mu_1}\cdot\hat{\boldsymbol{q}})\,|\boldsymbol{q},\mu_1\rangle\,, \tag{3.157}$$

which precisely correspond to the long-wavelength normal modes of the dispersive bands that we obtained in Section 2.5.2 by an analysis of the irreducible representations of the point group,

$$Q(\boldsymbol{q}) = \langle \boldsymbol{q}, -|S\rangle = \frac{1}{2}\sum_{\mu_1} S_{\mu_1}(\boldsymbol{q}), \tag{3.158}$$

$$\hat{\boldsymbol{q}} \cdot \boldsymbol{B}(\boldsymbol{q}) = \langle \boldsymbol{q}, +|S\rangle = \hat{\boldsymbol{q}} \cdot \left(\frac{\sqrt{3}}{2}\sum_{\mu_1} \hat{\boldsymbol{\delta}}_{\mu_1} S_{\mu_1}(\boldsymbol{q})\right), \tag{3.159}$$

where $\hat{\boldsymbol{\delta}}_{\mu_1}$ is the easy-axis unit vector oriented from diamond sublattice A to B. Indeed we should recover the same result from a Taylor expansion of Eq. (3.121). Notice that in the limit $\boldsymbol{q} \to \boldsymbol{0}$, Eq. (3.156) becomes $(1/2)\sum_{\mu_1} |\boldsymbol{0}, \mu_1\rangle$ which is precisely the eigenvector of the band maximum, whereas Eq. (3.157) does not have a well defined $\boldsymbol{q} = \boldsymbol{0}$ limit.

The rotational projector $\mathcal{P}^{\mathrm{im}\,\delta_2}$ matrix elements could be found by a computation of $\delta_2(\Delta_2)^{-1}\mathrm{d}_1$, but we can more easily obtain it simply as $\mathcal{P}^{\mathrm{im}\,\delta_2} = \mathbb{1} - \mathcal{P}^{\mathrm{im}\,\mathrm{d}_0} - \mathcal{P}^{H^1}$. Its matrix elements for $\boldsymbol{q} \neq \boldsymbol{0}$ are therefore given by $\mathcal{P}^{\mathrm{im}\,\delta_2}_{\mu_1\nu_1}(\boldsymbol{q}) = \delta_{\mu_1\nu_1} - \mathcal{P}^{\mathrm{im}\,\mathrm{d}_0}_{\mu_1\nu_1}(\boldsymbol{q})$, which with some massaging can be written at long wavelength as

$$\mathcal{P}^{\mathrm{im}\,\delta_2}_{\mu_1\nu_1}(\boldsymbol{q}) \approx \frac{3}{4}(\hat{\boldsymbol{\delta}}_{\mu_1} \times \hat{\boldsymbol{q}}) \cdot (\hat{\boldsymbol{\delta}}_{\nu_1} \times \hat{\boldsymbol{q}}). \tag{3.160}$$

From this we conclude that the long-wavelength eigenvectors of the two flat bands are given by the two linearly-independent components of the three vectors $(\sqrt{3}/2)\sum_{\mu_1}(\hat{\boldsymbol{\delta}}_{\mu_1} \times \hat{\boldsymbol{q}})\,|\boldsymbol{q}, \mu_1\rangle$.

3.4.3 Explicit Computation of the HHD for Spin Configurations

Now that we have the matrix elements of the irrotational projector Eq. (3.146) in the Fourier basis, we can utilize it to conveniently compute the HHD of spin configurations. Thus for example, to compute the direct-space components of the irrotational part of the HHD for a spin configuration, we Fourier transform the spin configuration, act on it with the projector, then inverse Fourier transform,

$$\left\langle e \middle| S^{\mathrm{irr}} \right\rangle = \langle e|\, \mathcal{P}^{\mathrm{im}\,\mathrm{d}_0}\, |S\rangle \tag{3.161}$$

$$= \sum_{\boldsymbol{q},\boldsymbol{q}'} \sum_{\mu_1,\nu_1} \langle e|\boldsymbol{q}, \mu_1\rangle \langle \boldsymbol{q}, \mu_1|\, \mathcal{P}^{\mathrm{im}\,\mathrm{d}_0}\, |\boldsymbol{q}', \nu_1\rangle \langle \boldsymbol{q}', \nu_1|S\rangle \tag{3.162}$$

$$\equiv \sum_{\boldsymbol{q}} \sum_{\nu_1} \frac{1}{\sqrt{L^3}} e^{-i\boldsymbol{q}\cdot\boldsymbol{r}_e} \mathcal{P}^{\mathrm{im}\,\mathrm{d}_0}_{\mu_1(e)\nu_1}(\boldsymbol{q}) S_{\nu_1}(\boldsymbol{q}), \tag{3.163}$$

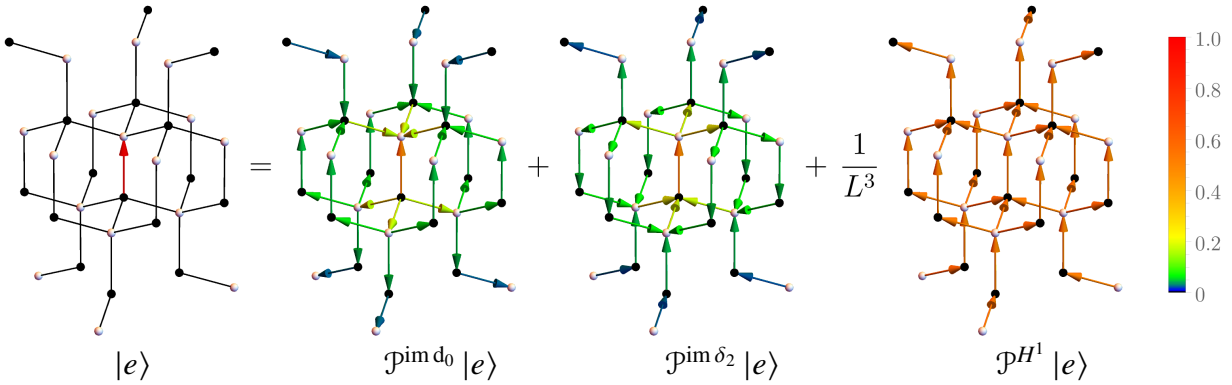


Figure 3.14: Helmholtz-Hodge decomposition for a 1-form on the diamond cell complex which is supported on a single oriented edge on which it has unit magnitude. The subextensive harmonic portion has been multiplied by L^3, in this example we took $L = 6$. Only the 1-skeleton of the complex is shown. We can visually observe that the irrotational part has non-zero divergence, and the rotational part has zero divergence and forms closed loops.

where $\boldsymbol{r}_e$ is the midpoint of edge e, $\mu_1(e)$ is the sublattice of edge e, and the we assumed e is oriented from A to B to fix the overall sign.

An example is shown in Fig. 3.14 for the HHD of a 1-cochain $|e\rangle$ supported on a single edge e of the diamond cell complex with unit magnitude, zero on all other edges. We represent the 1-cochain's value on each 1-cell by an arrow, where the direction indicates the orientation of the underlying 1-cell for which the cochain is positive and the color indicates the magnitude. We choose this example because the HHD of an arbitrary configuration can be reduced to a linear combination of the decompositions for each individual edge, since the projectors are linear operators,

$$\mathcal{P}^{\mathrm{im}\,\mathrm{d}_0}|S\rangle = \sum_e S(e)\mathcal{P}^{\mathrm{im}\,\mathrm{d}_0}|e\rangle\,. \tag{3.164}$$

Notice that the intensities (colors) on each edge for the rotational and irrotational portions are approximately equal, with the primary difference being the orientations of the arrows. Notice that the rotational component $\mathcal{P}^{\mathrm{im}\,\delta_2}|e\rangle$ has arrows forming closed head-to-tail loops. In the last term, the harmonic portion, we have multiplied the intensity by L^3, since it is subextensive. The harmonic portion is a constant 1-form, with zero divergence at each vertex and zero circulation around every plaquette.

3.4.4 Ground State Constraints and Projections

Nearest-neighbor ice models on a bipartite parent lattice are described in the SCGA approximation by the discrete 1-form Hamiltonian Eq. (3.86) on the parent cellular complex,

$$\beta H = \frac{1}{2}\langle S, (\lambda_0 + \beta J \mathrm{d}_0\delta_1)S\rangle = \frac{1}{2}\left[\lambda_0\langle S, S\rangle + \beta J\langle \delta_1 S, \delta_1 S\rangle\right]. \tag{3.165}$$

The 1-form spin correlation functions are

$$\langle S(e)S(e')\rangle \equiv \langle e|\, \mathcal{G}\, |e'\rangle\,. \tag{3.166}$$

The correlation matrix can be read off from the Hamiltonian as $\mathcal{G} = (\lambda_0 + \beta J \mathrm{d}_0\delta_1)^{-1}$. At zero temperature $\lambda_0 = (N^1_{\mathrm{sl}} - N^0_{\mathrm{sl}})/N^1_{\mathrm{sl}}$, where N^k_{sl} is the number of k-cells (sublattices) per primitive unit cell. The low temperature behavior of the correlations can be seen by a spectral decomposition of $\mathcal{G}$,

$$\mathcal{G} = \sum_{\boldsymbol{q}}\sum_{n} \frac{|\boldsymbol{q}, n\rangle\langle \boldsymbol{q}, n|}{\lambda_0 + \beta\epsilon_n(\boldsymbol{q})} \xrightarrow{\beta\to\infty} \sum_{\boldsymbol{q}} \sum_{\{n|\epsilon_n(\boldsymbol{q})=0\}} {\lambda_0}^{-1} |\boldsymbol{q}, n\rangle\langle \boldsymbol{q}, n|\,, \tag{3.167}$$

where n indexes the bands, from which we conclude that

$$\lim_{T\to 0} \mathcal{G} = \lambda_0^{-1}(\mathcal{P}^{\mathrm{im}\,\delta_2} + \mathcal{P}^{H^1}), \tag{3.168}$$

that is, the zero temperature correlation matrix is precisely λ_0^{-1} (the generalized susceptibility of the ground state modes) times the projector to the null space of $\mathrm{d}_0\delta_1$, the space of rotational and harmonic modes.

Ground states of the system are characterized by the constraint that $\delta_1 |S\rangle = 0$. In direct space it reads as one constraint per 0-cell,

$$0 = \sum_{e_{ij}} \langle k|\, \delta_1\, |ij\rangle\langle ij|S\rangle = \sum_{i@k} S_{ik}, \tag{3.169}$$

$$\equiv \sum_{e_{ij}} L^{(k)}_{e_{ij}} S(e_{ij}), \tag{3.170}$$

These equations are not all linearly independent, in particular if we sum over all k we obtain zero,

$$\sum_k \langle k|\, \delta_1\, |S\rangle = \left\langle \sum_k |k\rangle\,, \delta_1\, |S\rangle \right\rangle_1 = \left\langle \mathrm{d}_0 \sum_k |k\rangle\,, |S\rangle \right\rangle_1 = 0, \tag{3.171}$$

because $\sum_k |k\rangle$ is a constant 0-form (meaning a harmonic 0-form, an element of H^0) such that $\mathrm{d}_0 \sum_k |k\rangle = 0$, so there are only $(N_{\mathrm{sl}}^0 L^3 - 1)$ constraints. In the Fourier basis, the constraints can be expressed using Eq. (3.136) as one equation for each $\boldsymbol{q}, \mu_0$

$$0 = \langle \boldsymbol{q}, \mu_0 | \, \delta_1 \, |S\rangle \tag{3.172}$$

$$= \sum_{\mu_1} \langle \boldsymbol{q}, \mu_0 | \, \delta_1 \, |\boldsymbol{q}, \mu_1\rangle \, \langle \boldsymbol{q}, \mu_1 | S\rangle \,, \tag{3.173}$$

$$= \sum_{\mu_1} \left[\delta_{\mu_0 = B} \, e^{i\boldsymbol{q}\cdot\boldsymbol{\delta}_{\mu_1}/2} - \delta_{\mu_0 = A} \, e^{-i\boldsymbol{q}\cdot\boldsymbol{\delta}_{\mu_1}/2} \right] S_{\mu_1}(\boldsymbol{q}), \tag{3.174}$$

$$\equiv \sum_{\mu_1} L_{\mu_1}^{(\mu_0)}(\boldsymbol{q}) S_{\mu_1}(\boldsymbol{q}). \tag{3.175}$$

Note that here again we see that there are only $(L^3 - 1)$ constraints, because at $\boldsymbol{q} = \boldsymbol{0}$ the equations for $\mu_0 = A, B$ are both equivalent to $\langle \boldsymbol{0}, Q | S\rangle = 0$. This final version of the constraint Eq. (3.175) was originally obtained by Henley [136] as the statement that the projection on the dispersive bands of the adjacency matrix is zero, and so in the Hodge language it can be stated as the coordinate-free zero-divergence statement $\delta_1 \, |S\rangle = 0$, and the coefficients in the constraint equation L are the matrix elements of δ_1. The constraint can be rephrased as the statement that the projection of the spin configuration $|S\rangle$ on any basis of the irrotational modes is zero, and can be written as any of the following,

$$\left\langle \mathrm{d}_0 \, |k\rangle \, , |S\rangle \right\rangle = 0 \quad \forall \, k \tag{3.176}$$

$$\left\langle \mathrm{d}_0 \, |\boldsymbol{q}, \mu_0\rangle \, , |S\rangle \right\rangle = 0 \quad \forall \, \boldsymbol{q}, \mu_0 \tag{3.177}$$

$${}_1\langle \boldsymbol{q}, \pm | S\rangle \equiv \left\langle \, |\boldsymbol{q}, \pm\rangle_1 \, , |S\rangle \right\rangle = 0 \quad \forall \boldsymbol{q}, \pm \tag{3.178}$$

3.4.5 Topology of the Constraint

The constraint Eq. (3.175) was recently utilized as a starting point by Ref. [137] to define frustrated models with spin liquid ground states which can be classified topologically by an analysis of the functions $L_{\mu_1}^{(\mu_0)}(\boldsymbol{q})$, taking inspiration from Henley. Their method interprets these functions as forming two vector fields, one for each μ_0, on the FBZ torus. They do not present the analysis for spin ice and instead only study models with a single constraint vector field defined with a parent lattice which is Bravais (single sublattice). We will take up the analysis of spin ice here, and we present the analysis utilizing the language of fiber bundles. Our discussion will not be extremely rigorous, rather the goal is to be illustrative of the idea and to motivate the future development of a more rigorous perspective based on the mathematics of vector bundles.

The space Ω^1 is not just a vector space but a topological vector space, because it can be viewed as a discrete fiber bundle. Fiber bundles are a powerful tool which allows for the formalization of many notions from field theory and are ubiquitous objects of study in geometry and topology. Conceptually a fiber bundle B consists of a base manifold M with fibers glued to each point, where each fiber is isomorphic to the fiber space F. Restricted to a local contractible patch $U \subset M$, a fiber bundle looks like a direct product $U \times F$, but globally may not be a direct product. The basic examples are a cylinder, which is locally a product $[0,1] \times [0,1]$ and is globally a product $S^1 \times [0,1]$, and a Möbius strip, which is locally a product $[0,1] \times [0,1]$ but globally is "twisted" and not a product.

The most common type of fiber bundle is a vector bundle, for which the fiber space is a vector space. A real scalar field can thus be described as a real line bundle, where the fiber is a 1-dimensional real vector space (i.e. a real line), a complex scalar field is a complex line bundle (whose fibers are 1-dimensional complex vector spaces, i.e. a "complex line"), and scalar fields with multiple components are described by fiber bundles with multi-dimensional vector spaces for their fibers. The most common examples of vector bundles in physics are the tangent and cotangent bundles of a manifold. At each point of a manifold is attached the vector space of tangent vectors, which are directional derivatives on functions, or equivalently equivalence classes of 1-dimensional curves passing through that point. The cotangent vector space is then defined to be the dual space to tangent vector spaces (the space of linear functionals sending tangent vectors to scalars). A vector field on a manifold is a section of the tangent bundle, where a section of a bundle is a map $\sigma : M \to B$ which maps the base manifold to the bundle in such a way that it "slices" across the fibers, selecting one element of each fiber (thus one vector per tangent space), meaning that it satisfies $\pi \circ \sigma = \mathrm{id}$, where π is the projection that sends each point in the bundle to its base point in the manifold. Similarly, a section of a cotangent bundle is a cotangent vector field, a.k.a. a 1-form. k-forms are defined to be sections of the exterior bundle, which is defined by an alternating (i.e. anti-symmetric) tensor product of copies of the cotangent space. The spaces Ω^k in the continuum case can then be identified with the space of k-form sections of the exterior bundle.

In Fig. 3.9 we showed a cartoon illustration of a *discrete* line bundle, with a 1-dimensional fiber over each vertex of a cellular complex, which is the discrete version of a manifold. The space Ω^0 in the discrete case is then the space of sections of this discrete bundle. We can similarly define a discrete bundle with fibers over each edge, plaquette, etc., although these generalized discrete bundles are not precise analogues of their continuum counterparts, where the fibers are always localized at points of the manifold. Despite this, we can utilize the Fourier basis to describe Ω^k as a vector bundle over the FBZ torus, which in the thermodynamic $L \to \infty$ limit becomes a continuous manifold rather than a discrete space, where the fibers are N_{sl}^k-dimensional complex vector spaces. In the case of the diamond cell complex, Ω^1 has the structure of a fiber bundle over a 3-torus with fiber space a 4-dimensional complex vector space. The standard basis for Ω^1

is $|\boldsymbol{q}, \mu_1\rangle$, where now $\boldsymbol{q}$ is interpreted as a coordinate on the manifold which has topologically non-trivial periodic boundary conditions, $|\boldsymbol{q}, \mu_1\rangle \equiv |\boldsymbol{q} + \boldsymbol{b}_i, \mu_1\rangle$, and μ_1 indexes the four basis vectors in the fibers.

The constraint $\delta_1 |S\rangle = 0$ in the Fourier basis reads as two constraints per $\boldsymbol{q}$, Eq. (3.175), save for $\boldsymbol{q} = \boldsymbol{0}$ where the two constraints become degenerate. These two constraints can then be described by two sections of our rank-4 bundle, described in local coordinates as

$$\sigma^{(\mu_0)} = \sum_{\mu_1} L^{(\mu_0)}_{\mu_1}(\boldsymbol{q}) \, |\boldsymbol{q}, \mu_1\rangle \, . \tag{3.179}$$

These two sections span a 2-dimensional subspace at each $\boldsymbol{q}$ except $\boldsymbol{q} = \boldsymbol{0}$ where there is a singularity. This singularity is topologically non-trivial. We have the freedom to choose a basis in each fiber, which will change the functional form of the constraint equations (but not their solutions of course). Such a change of basis for a vector bundle is a gauge transform, and a choice of basis is described by a section of an associated principal $GL(\mathbb{C}, 4)$ bundle, though we will not use that construction here. We will take the basis to be given by the eigenvectors of $\mathrm{d}_0\delta_1$ such that the constraint is given by Eq. (3.178). Then we know that near $\boldsymbol{q} = \boldsymbol{0}$ the eigenvectors have the long-wavelength form of Eqs. (3.156) and (3.157). The section corresponding to $|\boldsymbol{q}, -\rangle$ does not have any singularity at $\boldsymbol{q} = \boldsymbol{0}$ and smoothly approaches the limit $(1/2)\sum_{\mu_1} |\boldsymbol{0}, \mu_1\rangle$. In a local patch U of the FBZ torus surrounding $\boldsymbol{q} = \boldsymbol{0}$, the bundle has local trivialization $U \times \mathbb{C}^4$,[35] and the fiber splits into a direct sum $\mathbb{C}^3 \oplus \mathbb{C}$, where the 1-dimensional subspace is spanned by the $|\boldsymbol{q}, -\rangle$. In the remaining rank-3 bundle surrounding the origin we have a local orthonormal basis in each fiber given by the three components of $|\boldsymbol{q}, \alpha\rangle \equiv (\sqrt{3}/2)\sum_{\mu_1} \hat{\delta}^{\alpha}_{\mu_1} |\boldsymbol{q}, \mu_1\rangle$. The constraint section is spanned by $\sum_\alpha \hat{q}^\alpha |\boldsymbol{q}, \alpha\rangle$, and has a singularity at $\boldsymbol{q} = \boldsymbol{0}$ where it is undefined. The topology of the section is then defined by deleting the point $\boldsymbol{q} = \boldsymbol{0}$ and considering a 2-cycle (a sphere) surrounding the deleted point, on which $\hat{\boldsymbol{q}}$ winds. Thus the topological defect can be viewed as a skyrmion in momentum space.[36] This is essentially the same result found by Ref. [137] for octochlore spin ice, which is the ice model defined on a lattice of corner sharing octahedra which is the line graph of the cubic lattice. It is interesting to note however that the cubic cell complex has only one vertex and three edges (therefore three spins) per unit cell, and thus only one constraint per unit cell, whereas the diamond has two vertices and four edges. Therefore on the diamond cell complex there are two constraint sections on a rank-4 bundle (vs. one constraint

[35] Every fiber bundle is locally a product of the form $U \times F$, but may differ from this local product structure globally. A fiber bundle which is globally a product is said to be trivial. A local trivialization of a bundle is the restriction to a local patch on which it decomposes into a product.

[36] We note that there should be a rigorous method to define the topology of the defect utilizing relative cohomology. The idea would be that the topology of the defect is classified by $H^2(T^3, D^3)$, the second relative cohomology of the 3-torus T^3 relative to a 3-disk D^3 surrounding the origin which is to be deleted, since the charge of the defect is a cohomology class computed from the integral of a 2-form over the surface of the ball [138].

section on a rank-3 bundle), but we find that despite this the topology of the constraint section is the same: there is only one skyrmion topological defect because one of the two sections is topologically trivial. This is a topic that warrants further development and study.

3.4.6 Correlations of Rotational and Irrotational Modes

Up until now we have considered spin-spin correlation functions, which in the Hodge language are interpreted as 1-form correlators $\langle S(e)S(e')\rangle \equiv \langle e|\,\mathcal{G}\,|e'\rangle$ for two edges e and e'. Interpreting spins as 1-forms we can define a number of natural correlation functions which target the correlations of the rotational and irrotational modes separately. First, we can define correlators for the divergence of the spin configuration, which are correlators between vertices of the cell complex, and can be computed directly from the spin-spin correlation matrix as follows,

$$\langle \delta_1 S(v)\delta_1 S(v')\rangle = \left\langle S(\partial^\dagger v)S(\partial^\dagger v')\right\rangle \tag{3.180}$$

$$= \left\langle S^{\text{irr}}(\partial^\dagger v)S^{\text{irr}}(\partial^\dagger v')\right\rangle . \tag{3.181}$$

This correlation functions is only sensitive to the irrotational component of the spin configuration (since δ_1 annihilates the rotational and harmonic components). It can be expressed in terms of the correlation matrix $\mathcal{G}$ as

$$\langle \delta_1 S(v)\delta_1 S(v')\rangle = \left\langle \partial^\dagger v\right|\mathcal{G}\left|\partial^\dagger v'\right\rangle = \langle v|\,\delta_1\mathcal{G}\mathrm{d}_0\,|v'\rangle\,. \tag{3.182}$$

We can define a corresponding structure factor for this correlator as

$$S_{\delta_1}(\boldsymbol{q}) \equiv \sum_{\mu_0,\nu_0}\langle \boldsymbol{q},\mu_0|\,\delta_1\mathcal{G}\mathrm{d}_0\,|\boldsymbol{q},\nu_0\rangle\,. \tag{3.183}$$

This structure factor is shown in Fig. 3.15. Note that, since this is a correlator of diamond lattice vertices it exhibits the symmetries of the BCC lattice formed by the diamond lattice and its dual, whose FBZ is an FCC Wigner Seitz cell.

This structure factor contains information about the correlation of the local divergence at each diamond lattice vertex. At high temperature we observe broad intensity near the zone corner H point in the cubic directions such as (002) and low intensity near the zone center. As temperature decreases, the intensity drops precipitously as the irrotational modes depopulate and the spectral weight becomes flat throughout the zone, indicating decorrelation of the local divergence, except at the zone center where it remains zero and at the P point. The zero intensity at zone center can be understood simply as a consequence of the fact that a 0-form that is a pure divergence (coexact) can never be a harmonic 0-form, which would be constant on every vertex. The space of

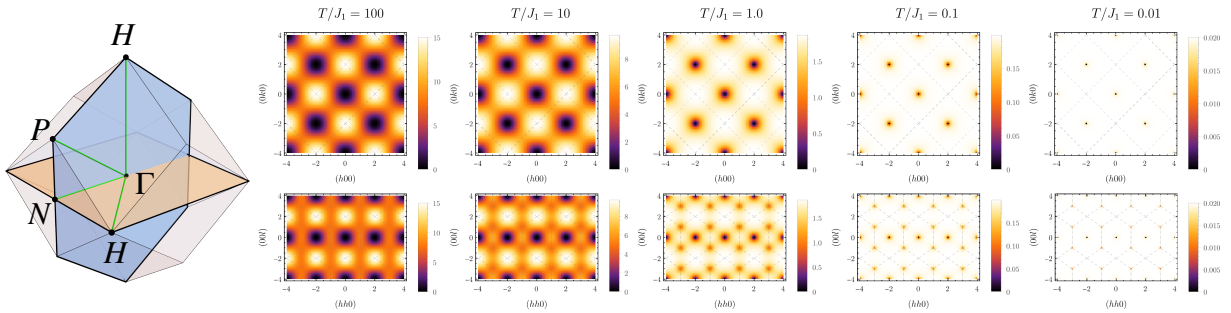


Figure 3.15: The divergence-divergence correlator structure factor defined in Eq. (3.183) computed for the nearest-neighbor spin ice model in the SCGA.

harmonic 0-forms H^0 is 1-dimensional and is spanned by the vector $|\mathbf{0}, A\rangle + |\mathbf{0}, B\rangle$. Thus the zero intensity at zone center is a consequence of the fact that we are looking at the correlator of an "emergent" 0-form degree of freedom (the local divergence) which is derived from the physical 1-form degrees of freedom. The same must hold true at any temperature, so the pattern observed in the high-temperature limit, where the underlying 1-forms are decorrelated, is indicative purely of correlation effects among the 0-form divergence caused by the fact that the space of 0-form configurations is not sampled uniformly due to the true degrees of freedom being 1-forms. The broad intensity appearing near the zone boundary is likely due to the fact that for a unit flux along a single edge the divergence at the two ends always has opposite sign, thus when the 1-form field is fully disordered its divergence appears to exhibit some short range correlations. We also see intensity at the zone corner P point which appears to be exactly half of the near uniform intensity throughout the zone at low temperatures. It is unclear what the interpretation of this zone corner intensity is, and it warrants further study, but it is almost certainly caused by the same mechanism.

Similarly, we can define curl-curl correlation functions of the form

$$\langle \mathrm{d}_1 S(p) \mathrm{d}_1 S(p') \rangle = \langle S(\partial p) S(\partial p') \rangle \tag{3.184}$$

$$\equiv \left\langle S^{\mathrm{rot}}(\partial p) S^{\mathrm{rot}}(\partial p') \right\rangle . \tag{3.185}$$

If we introduce a basis $|\boldsymbol{q}, \mu_2\rangle$ for Ω^2, defined as plane waves which are non-zero and centered on a single sublattice of plaquettes (note that the diamond lattice is self-dual, the plaquettes are in one-to-one correspondence with edges of the dual lattice, meaning that there are four sublattices for plaquettes), then we would define

$$S_{\mathrm{d}_1}(\boldsymbol{q}) \equiv \sum_{\mu_2, \nu_2} \langle \boldsymbol{q}, \mu_2 | \, \delta_2 \mathcal{G} \mathrm{d}_1 \, | \boldsymbol{q}, \nu_2 \rangle . \tag{3.186}$$

More generally, we can consider the correlations of the integral of the spin configuration 1-form around any two 1-cycles,

$$\langle S(C)S(C')\rangle \sim \left\langle \oint_C S \oint_{C'} S \right\rangle = \left\langle \oint_C S^{\text{rot}} \oint_{C'} S^{\text{rot}} \right\rangle, \tag{3.187}$$

which is only sensitive to the rotational component of the spin configuration. If the 1-cycles wind around a non-trivial cycle then there will also be a contribution of the harmonic part. We can express this correlator as

$$\langle S(C)S(C')\rangle = \langle C|\, \mathcal{G}\, |C'\rangle\,, \tag{3.188}$$

For the nearest-neighbor ice model in the soft spin approximation, we can express the correlation matrix as

$$\mathcal{G} = \underbrace{\sum_{\boldsymbol{q},\pm} \frac{|\boldsymbol{q},\pm\rangle\langle\boldsymbol{q},\pm|}{\lambda_0 + \beta\epsilon_\pm(\boldsymbol{q})}}_{\text{irrotational}} + \underbrace{\frac{1}{\lambda_0}\mathcal{P}^{\text{im}\,\delta_2}}_{\text{rotational}} + \underbrace{\frac{1}{\lambda_0}\mathcal{P}^{H^1}}_{\text{harmonic}}\,. \tag{3.189}$$

Since C and C' are closed, the correation function Eq. (3.187) is insensitive to the irrotational component, and therefore

$$\langle S(C)S(C')\rangle = \frac{1}{\lambda_0}\langle C|\,(\mathcal{P}^{\text{im}\,\delta_2} + \mathcal{P}^{H^1})\,|C'\rangle = \frac{1}{\lambda_0}\langle C|C'\rangle\,, \tag{3.190}$$

where in the last equality we utilized that $(\mathcal{P}^{\text{im}\,\delta_2} + \mathcal{P}^{H^1})\,|C'\rangle = |C'\rangle$ (since $|C'\rangle$ represents a circulation along a closed contour). If C and C' do not overlap each other, then this correlation function is zero. We will utilize this result in Chapter 4.

3.4.7 Further Neighbor Interactions: Extended Ice Models

Until now we have only considered nearest-neighbor ice models, which are minimal models which split into a sum over frustrated clusters by exploiting the connectivity of line graph lattices. In real materials we can generally expect that exchange interactions may be present between pairs of spins further than nearest neighbor, and indeed long-range dipole-dipole interactions are in principle always present.

The more further-neighbor interactions we add to the model the greater the connectivity of the network of interacting spins, and the more our system looks like a graph and less like a cellular complex. The underlying cell complex geometry is always present, but further neighbor interactions may spoil the description in terms of differentials and codifferentials of 1-form degrees

of freedom that we have discussed thus far. In a sense, the more connected the lattice becomes, the more the spins look like 0-forms rather than 1-forms. But it is possible to add further neighbor interactions in ways that exactly preserve the 1-form interpretation of the degrees of freedom, which we discuss in this section.

Consider a general spin Hamiltonian including further-neighbor interactions of the form

$$H = \frac{1}{2} \sum_{ij} S_i \mathcal{J}_{ij} S_j. \tag{3.191}$$

which respects the symmetries of the lattice. The space of possible models can be parameterized by the exchange coupling for the various neighbors. Due to symmetries, we can categorize neighbors according the their separation distance plus a symmetry class. For the pyrochlore lattice, for first and second neighbor separations there is only a single symmetry class, all such pairs of spins can be mapped to each other by a lattice symmetry, while at third neighbor there are two symmetry classes which we denote $3a$ and $3b$. In Tables B.1 to B.5 we tabulate all neighbors up to 100th, which will be necessary for our discussion of dipolar interactions. We define adjacency matrices $\mathcal{A}_{ij}^{(n,s_n)} = 1$ for neighbors at separation distance r_n of symmetry class s_n and 0 otherwise. Then the Hamiltonian can be expressed as

$$H = \frac{1}{2} \sum_{ij} S_i \left(\sum_n \sum_{s_n} J_{n,s_n} \mathcal{A}_{ij}^{(n,s_n)} \right) S_j. \tag{3.192}$$

Thus a model is specified by specifying the values of the exchange couplings J_{n,s_n}. The question we wish to answer is, what combinations of J_{n,s_n} will yield a model where the spins can be treated as 1-forms? We will not be able to give a complete answer to this question, but we can illustrate a method for constructing a large family of models which will satisfy this condition.

The basic question of whether further neighbor interactions can be added in a way that preserves the frustration for spin models on the pyrochlore lattice has been studied previously [57, 139]. In particular, it is known that adding second and third neighbors with $J_2 = J_{3a}$ preserves the flat bands in the spectrum, and that the resulting Hamiltonian can be re-written in the form

$$H = J \sum_t \left(\sum_{i \in t} S_i \right)^2 + \tilde{J} \sum_{\langle tt' \rangle} \left(\sum_{i \in t} S_i \right) \left(\sum_{j \in t'} S_j \right), \tag{3.193}$$

where t indexes tetrahedra and the second sum is over nearest-neighbor tetrahedra. For $J > 0$ and sufficiently small $\tilde{J}$, the ground states still obey the zero-net-spin constraint $\sum_{i \in t} S_i = 0$, but also the net spin of neighboring clusters is coupled. This is referred to as the extended spin ice

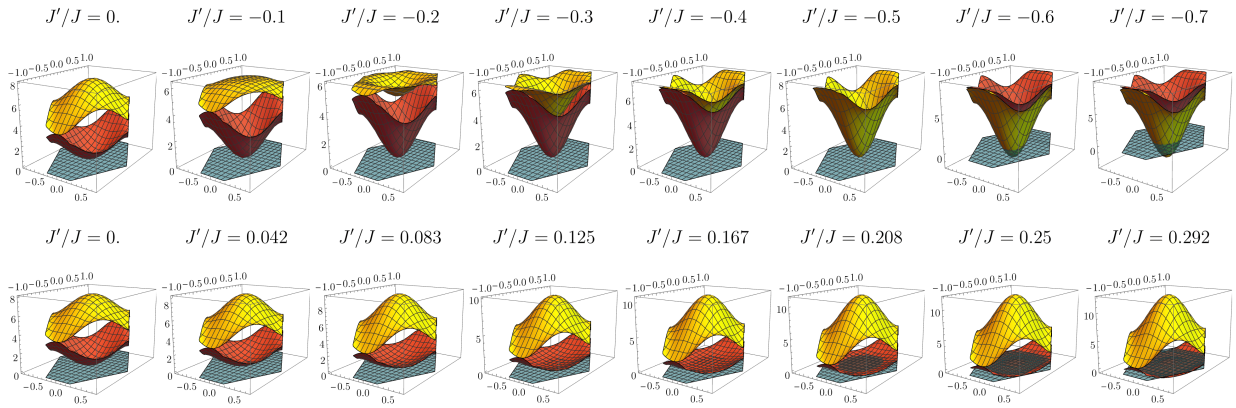


Figure 3.16: The bands of the interaction matrix of the ESI model, with $J_1 \equiv J > 0$ and $J_2 = J_{3a} = J'$. Upper row: $J' < 0$, tuning J' tunes the dispersion of the top band (yellow), until at a critical value of $J'/J = -1/2$ the upper band becomes degenerate with the lower dispersive band, after which it drops below the flat bands, leading to a condensation of the divergence at every site at sufficiently low temperature. Bottom row: $J' > 0$, tuning J' changes the dispersion of both the upper and lower irrotational bands, until at a critical value of $J'/J = 1/6$ the lower band goes soft, after which it disperses quadratically downward.

model [139]. For large $|\tilde{J}/J|$ the second term will prefer ordering of the local net spin, outweighing the first term's preference for minimizing the net spin. This model was further studied for an isotropic Heisenberg model on the pyrochlore and kagome lattices [57, 59]. It was later shown that these models can be considered part of a family of models whose interaction matrix is a polynomial of the nearest-neighbor adjacency matrix [104]. The basic idea is rather simple: the line graph adjacency matrix has flat bands with eigenvalue -2, therefore taking any polynomial of $(\mathcal{A}^{\mathrm{lg}} + 2\mathbb{1})$ will result in a matrix with flat bands with eigenvalue zero, since

$$(\mathcal{A}^{\mathrm{lg}} + 2\mathbb{1})^p = \sum_{\boldsymbol{q},n} \epsilon_n(\boldsymbol{q})^p \left|\psi_{\boldsymbol{q},n}\right\rangle \left\langle\psi_{\boldsymbol{q},n}\right| . \tag{3.194}$$

In Fig. 3.16 we show the bands of the ESI model with nearest-neighbor $J_1 \equiv J$ and second and third neighbor exchange $J_2 = J_{3a} \equiv J'$, with the flat bands set to zero energy. Tuning J' modifies the dispersion of the irrotational bands (red,yellow) but leaves the rotational bands flat. There is a range of values $-1/2 < J' < 1/6$ in which the flat bands remain below the irrotational bands and the soft spin ground state is still the Coulomb spin liquid. Beyond this range the irrotational bands drop below the flat bands, leading to a condensation of irrotational modes at sufficiently low temperatures.

With Hodge theory in hand, we can naturally recast these observations in the language of the

differential. According to our result Eq. (3.71) relating the line graph adjacency matrix to $\mathrm{d}_0\delta_1$, Eq. (3.194) is simply $(\mathrm{d}_0\delta_1)^p$. The extended spin ice soft spin Hamiltonian can be expressed as

$$H = \frac{1}{2}\left\langle S, \left(\lambda_0 + J^{(1)}\mathrm{d}_0\delta_1 + J^{(2)}(\mathrm{d}_0\delta_1)^2\right) S\right\rangle . \tag{3.195}$$

We can see naturally here that this Hamiltonian only couples the irrotational modes by utilizing the adjointness relation, writing it as

$$H = \frac{1}{2}\lambda_0\langle S, S\rangle + \frac{1}{2}J^{(1)}\,\langle \delta_1 S, \delta_1 S\rangle + \frac{1}{2}J^{(2)}\langle \delta_1 S, \underbrace{(\delta_1\mathrm{d}_0)}_{\Delta_0}\,\delta_1 S\rangle. \tag{3.196}$$

We see that the additional term provides nearest-neighbor interactions of the local divergence on the diamond vertices, coupled explicitly by the scalar Laplacian. We can therefore construct a large family of Hamiltonians which generalize extended spin ice,

$$H = \frac{1}{2}\lambda_0\langle S, S\rangle + \frac{1}{2}\sum_{n=1}^{\infty} J^{(n)}\langle S, (\mathrm{d}_0\delta_1)^n S\rangle = \frac{1}{2}\lambda_0\langle S, S\rangle + \frac{1}{2}\sum_{n=1}^{\infty} J^{(n)}\langle \delta_1 S, (\Delta_0)^{n-1}\delta_1 S\rangle. \tag{3.197}$$

The scalar Laplacian has matrix elements coupling nearest-neighbor vertices, and $(\Delta_0)^n$ couples vertices separated by up to n edges. Indeed, we can parameterize this set of Hamiltonians in terms of a potential energy for the local divergence,

$$H = \frac{1}{2}\lambda_0\langle S, S\rangle + \frac{1}{2}\sum_{v,v'}(\delta_1 S)_v V_{vv'}(\delta_1 S)_{v'}, \tag{3.198}$$

where $(\delta_1 S)_v \equiv \langle v|\,\delta_1\,|S\rangle$, and the scalar potential is

$$V_{vv'} \equiv \langle v|\left(\sum_{n=1}^{\infty} J^{(n)}(\Delta_0)^{n-1}\right)|v'\rangle\,. \tag{3.199}$$

In this sense we can interpret the divergence of S as an emergent charge. To return to the question posed at the beginning of this section, we can think of the parameters $J^{(n)}$ as forming a basis for a subspace of the full parameters space spanned by the J_{n,s_n} in which the Hamiltonian admits an exact 1-form description as provided here, since the interaction matrix $\sum_n J^{(n)}(\mathrm{d}_0\delta_1)^n$ can always be written as a linear combination of adjacency matrices. We will also consider additional terms which allow us to tune the rotational bands in a similar manner at the end of this chapter. The subspace spanned by these irrotational and rotational mode couplings is not the entire parameter space, meaning that there are combinations of parameters which do not admit an exact 1-form description. It would be worthwhile to develop a systematic understanding of how precisely the 1-form description is corrected when perturbing out of this subspace. We also note that a recent study of the extended kagome ice model finds that quantum fluctuations can stabilize a spin liquid even for small perturbations off this extended spin ice line [140].

3.4.8 Dipolar Spin Ice: Dumbbell Model and Projective Equivalence

We now turn to the long-range dipole-dipole interactions which are strong in the canonical spin ice materials $Ho_2Ti_2O_7$ and $Dy_2Ti_2O_7$ and are always present to some degree in magnetic materials. Thus far we described in detail the nearest-neighbor model and how it maps to a theory of a 1-form field on the diamond lattice cellular complex, then described how further neighbor interactions can generate a scalar potential for the emergent "charge" (i.e. the local divergence). Dipole-dipole interactions are further neighbor interactions which extend to infinite range and cannot be effectively truncated since the interaction strength falls off as distance cubed, such that the total energy is logarithmically divergent with system size.[37]

The dipole-dipole term in the spin Hamiltonian takes the form

$$H_{\text{dipole}} = \frac{1}{2}\sum_i (-\boldsymbol{\mu}_i \cdot \boldsymbol{B}(\boldsymbol{r}_i)) = -\frac{1}{2}\sum_{i\neq j} \frac{\mu_0}{4\pi} \frac{3\boldsymbol{\mu}_i \cdot \boldsymbol{\mu}_j - (\boldsymbol{\mu}_i \cdot \hat{\boldsymbol{r}}_{ij})(\boldsymbol{\mu}_j \cdot \hat{\boldsymbol{r}}_{ij})}{|\boldsymbol{r}_{ij}|^3}, \tag{3.200}$$

where $\boldsymbol{\mu}_i$ is the local magnetic moment at site i and we used the dipole field from Eq. (2.101), excluding self-interactions and dividing by two to avoid double counting. For Ising spins, we can express the magnetic moment as $\boldsymbol{\mu}_i \equiv \mu S_i \hat{\boldsymbol{\delta}}_i$,[38] and rewrite the dipolar Hamiltonian as

$$H_{\text{dipole}} = \frac{1}{2}\sum_{i\neq j} S_i \left(-\frac{\mu_0 \mu^2}{4\pi} \frac{\hat{\boldsymbol{\delta}}_i \cdot \hat{\boldsymbol{\delta}}_j - 3(\hat{\boldsymbol{\delta}}_i \cdot \hat{\boldsymbol{r}}_{ij})(\hat{\boldsymbol{\delta}}_j \cdot \hat{\boldsymbol{r}}_{ij})}{|\boldsymbol{r}_{ij}|^3} \right) S_j \equiv \frac{1}{2}\sum_{ij} S_i \mathcal{D}_{ij} S_j, \tag{3.201}$$

with $\mathcal{D}_{ii} = 0$. The dipole-dipole interaction matrix can be expressed as a linear combination of further-neighbor adjacency matrices with coefficients decaying as $1/r^3$, with couplings J_{n,s_n} rapidly fluctuating due to anisotropic nature of the matrix elements. *A priori* there is no obvious reason why adding dipole interactions should leave us in the subspace of Hamiltonians which admit a description in terms of 1-forms. But it is well known that the nearest-neighbor model does give a qualitatively correct description of spin ice physics.

There are two methods known for why the dipole interactions do not spoil the physics captured by the nearest-neighbor model, the original method due to Isakov et al. [142] was dubbed "projective equivalence", and the second method due to Castelnovo et al [119] is dubbed the "dumbbell approximation". Remarkably, Ref. [119] does not mention the connection to the work of Ref. [142], but we will demonstrate here that the two methods are essentially equivalent.

We begin with a brief description of the dumbbell approximation. The idea is quite simple and elegant: an ideal infinitesimal dipole with dipole moment μ can be approximated by a pair of

[37] More precisely, the Fourier transform of the dipole-dipolar interaction is conditionally convergent, see Ref. [141].

[38] where $\boldsymbol{\delta}_i$ is the displacement vector from the A to B diamond sublattices used in Section 3.4.2.

monopoles with charges $\pm q$ separated by a distance d, where $\hat{\boldsymbol{\mu}}$ points from the negative to the positive charge. A single charge located at $\boldsymbol{r}'$ produces a Coulomb field

$$\boldsymbol{B}(\boldsymbol{r}) = \frac{\mu_0}{4\pi} \frac{q}{|\boldsymbol{r} - \boldsymbol{r}'|}. \tag{3.202}$$

Adding together the $1/r$ Coulomb fields produced by a pair of opposite charges recovers the $1/r^3$ field of an ideal dipole at distances $r \gg d$, with $\mu = qd$, plus corrections decaying as $1/r^5$. Thus if we simply replace our ideal dipoles located at the midpoint of each diamond edge with a "dumbbell" of magnetic monopoles located at either end of the edge with charge $q = \pm\mu/a_d$ where a_d is the length of the edge, interacting via a $1/r$ Coulomb potential,

$$H = \frac{1}{2} \sum_{v \neq v'} \frac{\mu_0}{4\pi} \frac{Q_v Q_{v'}}{|\boldsymbol{r}_v - \boldsymbol{r}_{v'}|}, \tag{3.203}$$

where Q_v is the total magnetic charge at vertex v, then we will recover the problematic long-distance $1/r^3$ portion of the dipole-dipole interaction matrix. The dumbbell model Hamiltonian is then written entirely in terms of these magnetic monopole degrees of freedom,

$$H_{\text{dumbbell}} = \frac{1}{2} J_Q \sum_t Q_v^2 + \frac{1}{2} \sum_{v \neq v'} \frac{\mu_0}{4\pi} \frac{Q_v Q_{v'}}{r_{vv'}}, \tag{3.204}$$

where J_Q comes from the nearest-neighbor interaction of pyrochlore spins (which can contain both dipole and exchange contributions), serving as a chemical potential for the charges. In the dumbbell model, these "magnetic monopoles" are true emergent degrees of freedom which carry an emergent charge and interact via a long-range $1/r$ Coulomb potential [56, 58, 119].

The dumbbell Hamiltonian Eq. (3.204) looks remarkably similar to Eq. (3.198) (with the $v = v'$ term in $V_{vv'}$ separated) if we identify the total charge Q_v with $(\delta_1 S)_v$ up to a proportionality constant. The relation we need is simply

$$Q_v = \frac{\mu}{a_d} (\delta_1 S)_v, \tag{3.205}$$

which recovers the correct total charge when the 1-form field obeys the hard spin constraint $S(e) = \pm 1$. Then we identify

$$V_{vv'} = \begin{cases} \frac{\mu_0}{4\pi} \frac{\mu^2}{a_d^2} r_{vv'}^{-1} & v \neq v', \\ J_Q \frac{\mu^2}{a_d^2} & v = v'. \end{cases} \tag{3.206}$$

Thus the dumbbell approximation fits neatly into the Hodge theory, and we can see that it modifies the Hamiltonian in the irrotational sector and leaves the rotational modes degenerate. This means that the dumbbell model interaction matrix retains the perfectly flat rotational bands. In order to see this directly, however, we would have compute the Fourier transform of the dumbbell model interaction matrix $\mathcal{J}_{\mathrm{db}} = \mathrm{d}_0 V \delta_1$, with V given in direct space by Eq. (3.206). This is quite a daunting task, and leaves us with the same issue as the original dipole interaction matrix: the Fourier transform will be conditionally convergent due to the inifinite-range interaction. Furthermore, it is unclear how to systematically compute the corrections to the dumbbell approximation. These are the tasks we will now grapple with, and the primary tool will be projective equivalence.

Projective equivalence as presented by Isakov et al. [142] is the following statement: the dipole-dipole interaction matrix $\mathcal{D}_{ij}$ is proportional to the projector to the dispersive bands of the line graph adjacency matrix $\mathcal{A}^{\mathrm{lg}}$ plus corrections which decay as $1/r^5$ or faster. Intuitively, it seems that this statement must be directly related to the dumbbell approximation, which captures the long-range $1/r^3$ portion of the dipolar interactions and should have short range (integrable) corrections decaying as $1/r^5$. This connection has not been shown previously, but it follows easily from the Hodge theory. We know that the projector to the dispersive bands is given by $\mathcal{P}^{\mathrm{im}\,\mathrm{d}_0} = \mathrm{d}_0(\Delta_0)^{-1}\delta_1$. Therefore we can consider a projective equivalence (PE) Hamiltonian of the form

$$H_{\mathrm{PE}} = \frac{1}{2} J_{\mathrm{PE}} \left\langle S, \mathrm{d}_0(\Delta_0)^{-1}\delta_1 S \right\rangle = \frac{1}{2} J_{\mathrm{PE}} \left\langle \delta_1 S, (\Delta_0)^{-1}\delta_1 S \right\rangle . \tag{3.207}$$

This Hamiltonian has precisely the form of Eq. (3.198), with a scalar potential

$$V_{vv'} = J_{\mathrm{PE}} \langle v | \, (\Delta_0)^{-1} \, | v' \rangle . \tag{3.208}$$

We can compute the long-distance form of $V_{vv'}$ by utilizing the inverse Laplacian matrix elements in the Fourier basis Eq. (3.140), using $\langle v | \boldsymbol{q}, \mu_0 \rangle = L^{-3/2} \exp(-i\boldsymbol{q} \cdot \boldsymbol{r}_v)$ if v is in sublattice μ_0, to obtain

$$V_{vv'}/J_{\mathrm{PE}} = \frac{1}{L^3} \sum_{\boldsymbol{q}} \frac{1}{16 - |G(\boldsymbol{q})|^2} \begin{pmatrix} 4 & G^*(\boldsymbol{q}) \\ G(\boldsymbol{q}) & 4 \end{pmatrix}_{\mu_0(v)\nu_0(v')} e^{-i\boldsymbol{q}\cdot(\boldsymbol{r}_{v'} - \boldsymbol{r}_v)} . \tag{3.209}$$

For large separations $|\boldsymbol{r}_{v'} - \boldsymbol{r}| \gg a_d$, the sum is dominated by the long-wavelength terms, and is in fact dominated by the singularity at $\boldsymbol{q} = \boldsymbol{0}$. Thus expanding to leading order, using

$G(\boldsymbol{q}) \approx 4 - \sum_{\rho_1} (\boldsymbol{\delta}_{\rho_1} \cdot \boldsymbol{q})^2$ we obtain

$$V_{vv'}/J_{\text{PE}} \approx \frac{1}{L^3} \sum_{\boldsymbol{q}} \frac{e^{-i\boldsymbol{q}\cdot(\boldsymbol{r}_{v'} - \boldsymbol{r}_v)}}{\sum_{\rho_1} (\boldsymbol{\delta}_{\rho_1} \cdot \boldsymbol{q})^2} \tag{3.210}$$

$$= \frac{1}{L^3} \sum_{\boldsymbol{q}} \frac{e^{-i\boldsymbol{q}\cdot(\boldsymbol{r}_{v'} - \boldsymbol{r}_v)}}{(4/3)|a_d \boldsymbol{q}|^2} \tag{3.211}$$

$$\xrightarrow{L\to\infty} \frac{3}{4a_d^2} V_{\text{FCC}} \int_{\text{FBZ}} \frac{\mathrm{d}^3\boldsymbol{q}}{(2\pi)^3} \frac{e^{-i\boldsymbol{q}\cdot(\boldsymbol{r}_{v'} - \boldsymbol{r}_v)}}{|\boldsymbol{q}|^2} \tag{3.212}$$

$$= \frac{3V_{\text{FCC}}}{4a_d^2} \frac{1}{4\pi} \frac{1}{|\boldsymbol{r}_{v'} - \boldsymbol{r}_v|}. \tag{3.213}$$

where in the second line we used $\boldsymbol{\delta} = a_d \hat{\boldsymbol{\delta}}$ and the overcompleteness relation $\sum_{\rho_1} \hat{\boldsymbol{\delta}}_{\rho_1} \hat{\boldsymbol{\delta}}_{\rho_1}^T = (4/3)\mathbb{1}$. In the third line we took the continuum limit, replacing $(1/L^3)\sum_{\boldsymbol{q}} \to (1/V_{\text{FBZ}}) \int \mathrm{d}^3\boldsymbol{q}$ with the Brillouin zone volume given by $V_{\text{FBZ}} = (2\pi)^3/V_{\text{FCC}}$ where V_{FCC} is the volume a single FCC unit cell. In the final line we performed the integral by taking the integration limits to infinity, changing to spherical coordinates with axis along $(\boldsymbol{r}_{v'} - \boldsymbol{r}_v)$, and integrating the altitudinal angle, after which the integral reduces to $(1/2\pi^2) \int_0^\infty \mathrm{d}q\, \mathrm{sinc}(qr) = (1/4\pi r)$. Indeed we can recognize $1/4\pi|\boldsymbol{x} - \boldsymbol{x}'|$ as the Green's function for the continuum Laplacian, and it should not be surprising that the matrix elements of the inverse discrete scalar Laplacian Δ_0 are proportional at large separations to the inverse of the continuum Laplacian.

To determine the value of J_{PE}, we match the long-distance $1/r$ portion of the projective equivalence potential Eq. (3.213) to the dumbbell model potential Eq. (3.206), from which we obtain

$$\frac{3V_{\text{FCC}}}{4a_d^2} \frac{1}{4\pi} J_{\text{PE}} = \frac{\mu_0}{4\pi} \frac{\mu^2}{a_d^2} \longrightarrow J_{\text{PE}} = \frac{4\mu_0\mu^2}{3V_{\text{FCC}}}. \tag{3.214}$$

The projective equivalence interaction matrix $J_{\text{PE}} \langle e|\, \mathrm{d}_0 (\Delta_0)^{-1} \delta_1\, |e'\rangle$ then has the same matrix elements of the dipole-dipole interaction matrix $\mathcal{D}_{ee'}$ in the limit of large separations, provided we take the same orientations of the edges e, e' as we did for the definition of Eq. (3.201). Thus we can say that

$$J_{\text{PE}} \langle e|\, \mathrm{d}_0 (\Delta_0)^{-1} \delta_1\, |e'\rangle = \mathcal{D}_{ee'} + O\left(r_{ee'}^{-5}\right). \tag{3.215}$$

We will compute the $O\left(r_{ee'}^{-5}\right)$ correction terms in the next section.

Thus the projective equivalence Hamiltonian, where the interaction matrix is replaced with the projector to the irrotational subspace is equivalent *at large separations* to the dumbbell model. They are not exactly the same model, however: in the dumbbell model the scalar potential for the

charges is chosen to be the continuum $1/r$ at *all* distances, Eq. (3.206), whereas in the projective equivalence model it is taken to be the inverse of the discrete scalar Laplacian, which goes as $1/r$ at *large* separations. In both cases the short-range portion can be corrected to first order by modifying the nearest-neighbor coupling, which controls the monopole chemical potential (i.e. the J_Q term in the dumbbell potential). They both then reproduce the most salient feature of the dipole interactions: the $1/r$ Coulomb potential between monopoles, but they both fail to capture any corrections to the degeneracy of the rotational modes. This is where the projective equivalence approach has an advantage over the dumbbell approach. As we will see, it is relatively straightforward to compute additional corrections to the projective equivalence interaction matrix, which is not clear how to do in the dumbbell approach since it is defined in direct space and is not amenable to Fourier analysis. These corrections are weak [142] and in fact have been claimed to be entirely irrelevant to the physics of actual compounds [58]. We will demonstrate in Section 3.4.10 that the first order correction term can be obtained from the Hodge theory in an elegant manner and in Chapter 4 we will demonstrate that evidence for this correction is already present in experimental neutron scattering data.

Before we move on, it is worth noting that the projective equivalence Hamiltonian Eq. (3.207) can simply be written as

$$H_{\mathrm{PE}} = \frac{1}{2} J_{\mathrm{PE}} \langle S^{\mathrm{irr}}, S^{\mathrm{irr}} \rangle. \tag{3.216}$$

Previously we said that the most salient feature of the dipolar interactions which is captured by the projective equivalence Hamiltonian is the $1/r$ Coulomb potential between monopoles. In the form of Eq. (3.216), it appears that this is equivalent to the statement that the primary effect of dipole interactions is to *uniformly suppress* the irrotational component of the 1-form field.

3.4.9 Corrections to the Dumbbell Approximation

While computing corrections to the dumbbell model directly is a difficult task, doing so for the projective equivalence model is more straightforward. Since we already have the matrix elements for the Fourier transformed irrotational projector, we can inverse Fourier transform it to obtain its direct space components and take the difference with the direct space dipole dipole interaction matrix elements. To do so we will first normalize the dipole-dipole interaction matrix as

$$\mathcal{D}_{ij} = \frac{\mu_0 \mu^2}{4\pi a^3} \tilde{\mathcal{D}}_{ij}, \tag{3.217}$$

where a is a lengthscale and $\tilde{\mathcal{D}}_{ij}$ is now dimensionless, given by Eq. (3.201) with $|\boldsymbol{r}_{ij}|^3 \to |\boldsymbol{r}_{ij}/a|^3$. Then from Eq. (3.215) we have

$$\tilde{\mathcal{D}}_{ee'} = \frac{16\pi}{3}\frac{a^3}{V_{\mathrm{FCC}}}\mathcal{P}^{\mathrm{im\,d_0}}_{ee'} + \Delta_{ee'}. \tag{3.218}$$

where $\Delta_{ee'}$ contains the corrections decaying as $1/r^5_{ee'}$. We can compute the matrix elements of $\Delta_{ee'}$ numerically by inverse Fourier transforming the projector computed in Eq. (3.146),

$$\Delta_{ee'} = \tilde{\mathcal{D}}_{ee'} - \frac{64\pi}{3}\frac{1}{L^3}\sum_{\boldsymbol{q}} \mathcal{P}^{\mathrm{im\,d_0}}_{\mu_1(e)\nu_1(e')}(\boldsymbol{q}), \tag{3.219}$$

where we have chosen a to be the standard cubic cell edge length, such that $V_{\mathrm{FCC}}/a^3 = 1/4$.[39] The results are shown for $r_{ee'}$ up to twelve times the nearest-neighbor separation on the pyrochlore lattice, or 100th nearest-neighbors, in Fig. 3.17, computed for $L = 60$ (note that twelve times the nearest-neighbor distance is equivalent to only six FCC unit cells). The matrix elements indeed decay as $1/r^5$. We can express the corrections in terms of further neighbor interactions as

$$\Delta_{ee'} = \sum_{n=1}^{\infty}\sum_{s_n} \Delta_{n,s_n}\mathcal{A}^{(n,s_n)}_{ee'}. \tag{3.220}$$

Explicit values for $\Delta_{(n,s_n)}$ are tabulated in Tables B.6 to B.8.

Since the matrix elements of the correction terms Δ decay as an integrable $1/r^5$,[40] they can be safely truncated to obtain an accurate approximation of the matrix elements of the dipole interaction matrix. In Fig. 3.18 we show (a) the nearest-neighbor model bands of $\mathrm{d_0}\delta_1$, (b) the bands of the dipole-dipole interaction matrix, normalized by dividing by the nearest-neighbor value $\mathcal{D}_{\langle ij\rangle}$, and (c) the *nearly* flat rotational bands of the dipole-dipole interaction matrix, demonstrating a very weak dispersion on the order of one percent of the band width.

By far the most prominent difference between the nearest-neighbor model and the dipole-dipole interaction matrices is the lifting of the irrotational bands away from the rotational bands. This is most easily understood in the projective equivalence model, where the interaction matrix is simply the projector to the irrotational subspace, which lifts the irrotational bands up by J_{PE}. This is less

[39] Note that there is an apparent error in Ref. [142], they obtain a coefficient of $8\pi/3$ and claim that they have taken a to be the nearest-neighbor distance on the pyrochlore lattice, which is $1/\sqrt{8}$ times the cubic cell edge length, but this would produce a coefficient $4\pi/3\sqrt{2}$. The correctness of our formula is validated by the fact that the matrix elements of Δ indeed decay as $1/r^5$ as seen in Fig. 3.17.

[40] By integrable here we mean that the number of neighbors at distance r_n increases as r_n^2, so the total contribution to the energy from the correction terms at distance r_n decays as $1/r_n^3$.

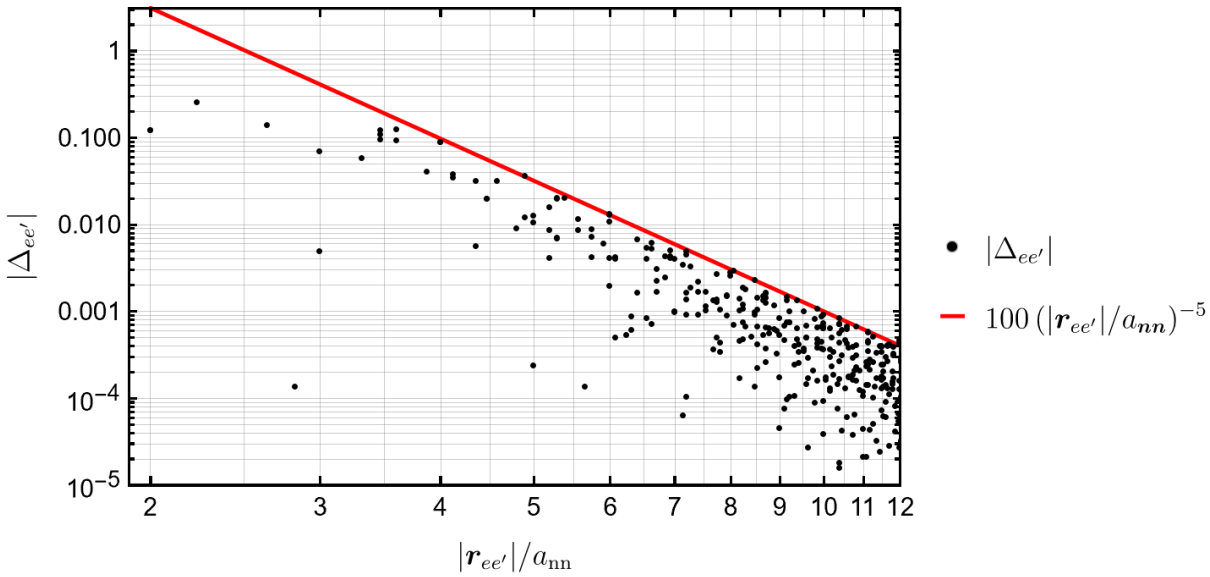


Figure 3.17: The matrix elements of $\Delta_{ee'}$ computed from Eq. (3.219) for $L = 60$ up to twelve times the nearest-neighbor pyrochlore separation or 100'th nearest neighbor, which decay as $1/r^5$.

obvious based on the direct space construction of the dumbbell model, but with the projective equivalence model we can conclude that the lifting of the irrotational bands is a direct manifestation of the long-range $1/r$ Coulomb interaction of monopoles and the singular nature of the dipolar interactions. In particular, the band touching has been lifted and the harmonic modes are not even well-defined—a thermodynamically large system with a constant uniform magnetization has infinite energy due to the non-integrability of the infinite-range dipole interactions.

Note that if the Hamiltonian were just proportional purely to the irrotational projector then the upper bands would be flat and degenerate, and just by observing Fig. 3.18(a,b) we can conclude that the leading correction to the projector comes from the nearest-neighbor terms which give the lifted irrotational bands the shape of the nearest-neighbor spectrum, which is clear from the fact that the nearest-neighbor term in the corrections Δ in Table B.6 is two orders of magnitude larger than any of the other correction terms. In addition to the lifting of the irrotational bands, their dispersion receive some level of modification compared to the nearest-neighbor bands, though it is clearly imperceptible to the eye in Fig. 3.18 owing to the further-neighbor corrections to Δ being on the order of one percent of the nearest-neighbor term, which determines the chemical potential or self-energy of the monopoles in the dumbbell model. More importantly, however, is the correction to the flat rotational bands. Since the bands are perfectly flat and degenerate in the nearest neighbor model, they are insensitive to the nearest-neighbor corrections but are highly

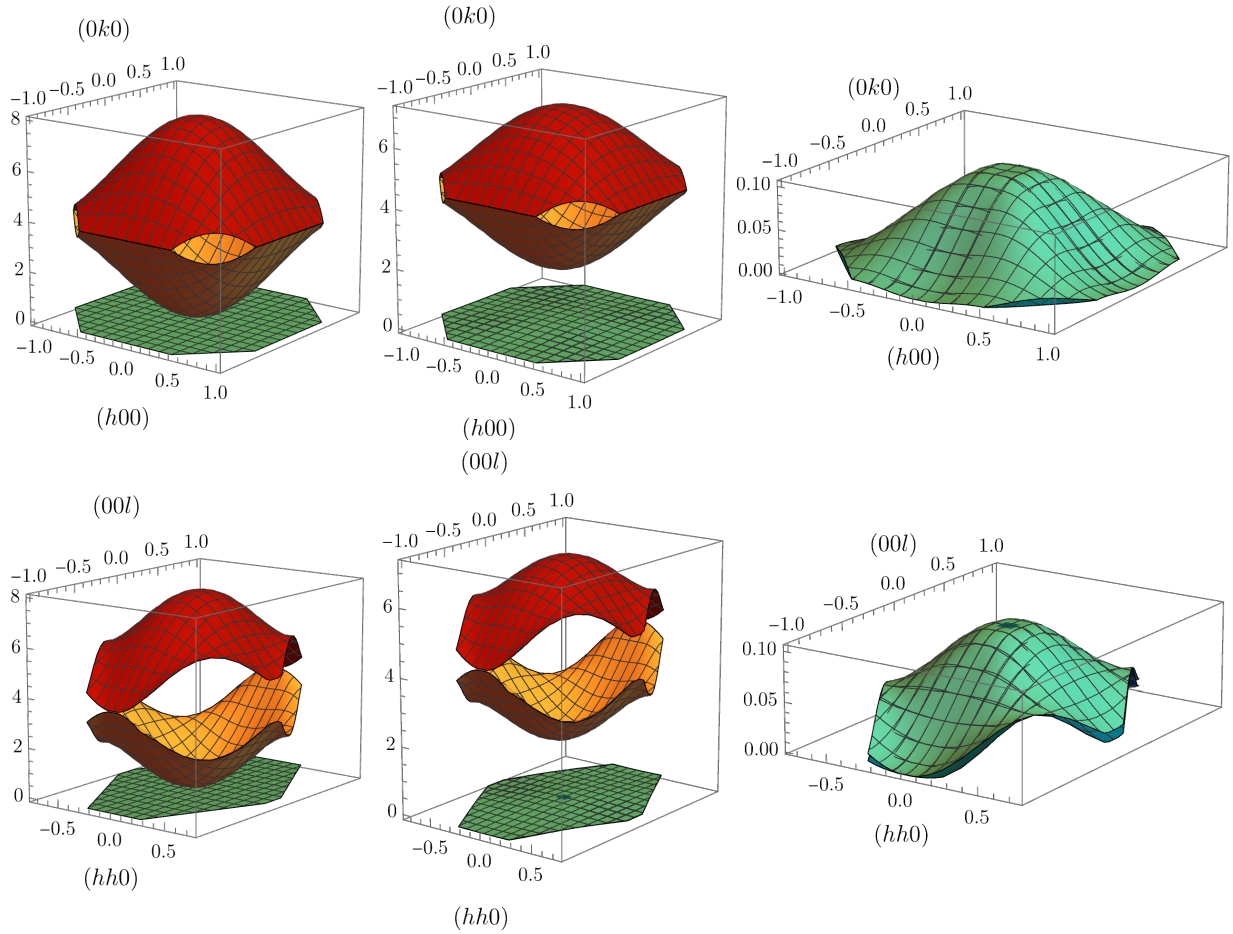


Figure 3.18: (a) The bands of the nearest-neighbor interaction matrix $d_0\delta_1$. (b) The bands of dipole-dipole interaction matrix, normalized such that the nearest-neighbor coupling is 1, computed by truncating the correction terms Δ_{n,s_n} to the projective equivalence interaction matrix at 100'th nearest-neighbor. (c) The two rotational bands of the dipole-dipole interaction matrix, showing a very weak dispersion and a slight lifting of the degeneracy of the two bands.

sensitive to the further neighbor corrections. In Fig. 3.18(c) we can clearly observe they are no longer exactly flat, having gained a dispersion downward away from $\boldsymbol{q} = \mathbf{0}$ with a band minimum the zone boundary X point. In the next section we will demonstrate that this dispersion can be captured by a simple correction of the form $\delta_2 \mathrm{d}_1$ and how it eventually should drive the system to long range order.

3.4.10 Corrections to the Flat Bands: Vorticity

It is not difficult to deduce the first order correction to the dumbbell and projective equivalence models which lifts the flatness of the rotational bands. From the Hodge theory perspective, the elephant in the room that we have until now ignored, $\delta_2 \mathrm{d}_1 \sim \mathrm{curl}\,\mathrm{curl}$, which is the complementary operator to $\mathrm{d}_0 \delta_1$ acting in the rotational band subspace. But let us first focus on the magnetic picture, and begin with a qualitative understanding of why the rotational bands should become dispersive at all.

In the presence of a magnetic field, a magnetic dipole minimizes its energy by orienting itself with the local field. The field of a dipole is fully rotational in the language of the Helmholtz decomposition, having zero divergence, and thus the field lines form closed loops. Consider a single fixed dipole, and consider placing an array of test dipoles around it. They each align along the local field lines of the fixed dipole, just as iron filings do near a bar magnet. As such the array of dipoles will prefer to form closed loops to minimize their energy. The same holds true if the array of dipoles are all mutually interacting, as in a magnetic compound. Another way to think of this is that they seek to mutually reduce the total magnetic energy $(\mu_0/2) \int \mathrm{d}^3 r \, |\boldsymbol{B}|^2$, and the best way to do this is to make the field lines form tight loops, with regions of parallel field lines costing more energy.

Based on this logic, we can conclude that the presence of dipole interactions will create an energetic preference for the magnetic moments to align in small loops. In the nearest-neighbor model, where the rotational bands are flat and the rotational modes are completely degenerate, there is no preference for longer or shorter loops, since all closed-loop configurations have the same energy, with the preference only being driven entropically by the hard spin length constraint [143]. The dispersion of the rotational bands of the dipole interaction matrix thus ought to be a consequence of an energetic preference for smaller loops, which is evidenced by the fact that they have a minimum near the zone boundary in Fig. 3.18(c).

We can naively construct an interaction term for our Hamiltonian which will accomplish just this. The energetic preference is to *maximize circulation* around the smallest loops of the lattice: the boundaries of the plaquettes, illustrated in Fig. 3.19. The term we would like to add to our

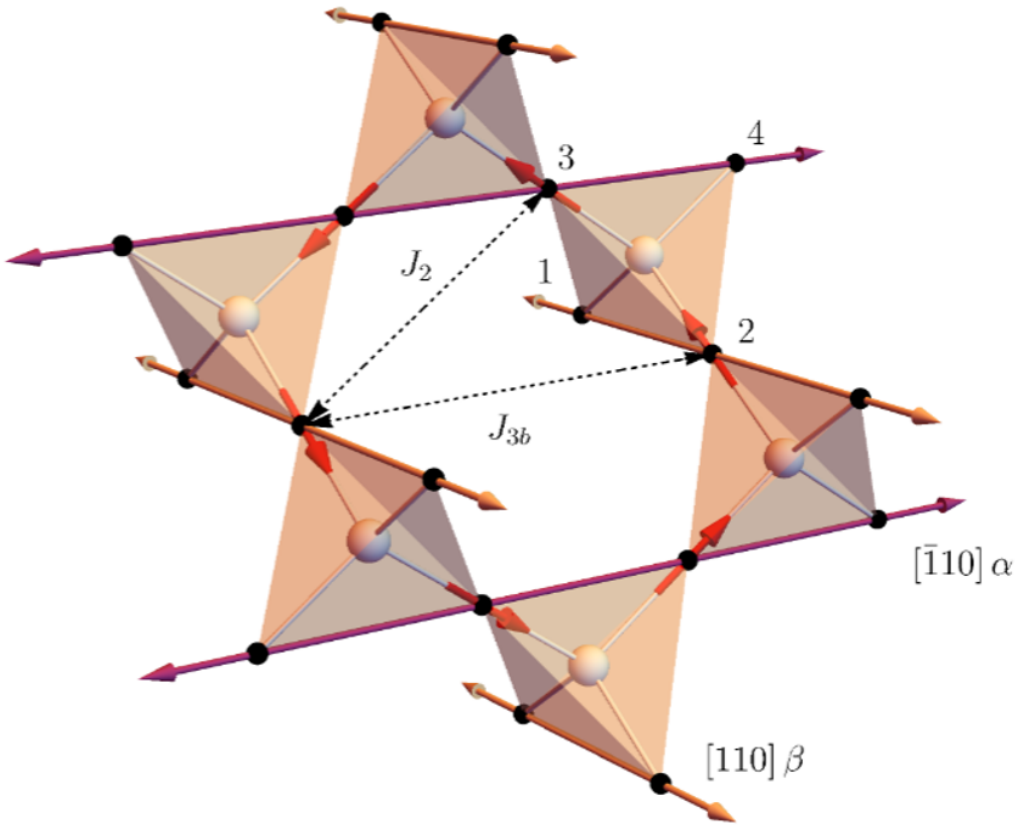


Figure 3.19: The minimal hexagonal loop motif of the pyrochlore lattice, with the diamond cell complex shown inside the tetrahedra. second and third-type-*b* neighbors are indicated. Six spins (red arrows) are illustrated forming a closed loop around the hexagon, maximizing the total circulation.

Hamiltonian should then take the form

$$H_\omega \sim -\frac{1}{2}|J_\omega| \sum_p \left(\oint_{\partial p} S \right)^2 , \tag{3.221}$$

where the sum is over all plaquettes, which creates an energetic preference for maximizing the circulation around every plaquette boundary. If we index the six sites in a hexagon sequentially 1 through 6 around the boundary and take the standard convention that edges are oriented from the A to B diamond sublattices, then we can express the above interaction term as

$$H_\omega = -\frac{1}{2}|J_\omega| \sum_p (S_1 - S_2 + S_3 - S_4 + S_5 - S_6)_p^2 , \tag{3.222}$$

where the alternating sign follows from the alternating orientations of the edges. Expanding the square, we observe that (i) every nearest-neighbor pair appears as $-S_1 S_2$, (ii) every second-neighbor pair appears as $+S_1 S_3$, (iii) every third-type-b neighbor appears as $-S_1 S_4$, cf. Fig. 3.19. Since these plaquettes share edges, summing over every hexagon results in a double counting of nearest-neighbors. Therefore we can write the above in terms of adjacency matrices as

$$H_\omega = -\frac{1}{2} \sum_{ij} S_i \underbrace{|J_\omega| \left(6\delta_{ij} - 2\mathcal{A}_{ij}^{(1)} + \mathcal{A}_{ij}^{(2)} + \mathcal{A}_{ij}^{(3b)} \right)}_{\equiv \mathcal{J}_{ij}^{(\omega)}} S_j , \tag{3.223}$$

where the factor of six comes from the fact that six hexagons meet at each edge of the diamond cell complex.

Let us now see that this interaction matrix is equivalent to the operator $\delta_2 \mathrm{d}_1$. First, the differential d_1 acting on a 1-form S produces a 2-form $\mathrm{d}_1 S$ whose value on a 2-cell is $\mathrm{d}_1 S(p) = S(\partial p) \sim \oint_{\partial p} S$ by Stokes' theorem. Therefore we can write Eq. (3.221) as

$$H_\omega = -\frac{1}{2}|J_\omega| \langle \mathrm{d}_1 S, \mathrm{d}_1 S \rangle = \frac{1}{2}|J_\omega| \langle S, \delta_2 \mathrm{d}_1 S \rangle, \tag{3.224}$$

where we used the adjointness relation. Let us define the *vorticity 2-form*,

$$\omega = \mathrm{d}_1 S, \tag{3.225}$$

which can be contrasted with the "charge" or "divergence" 0-form,[41]

$$Q = \delta_1 S. \tag{3.226}$$

[41] Until now we have mostly avoided calling the divergence of the 1-form field "charge." In the context of the

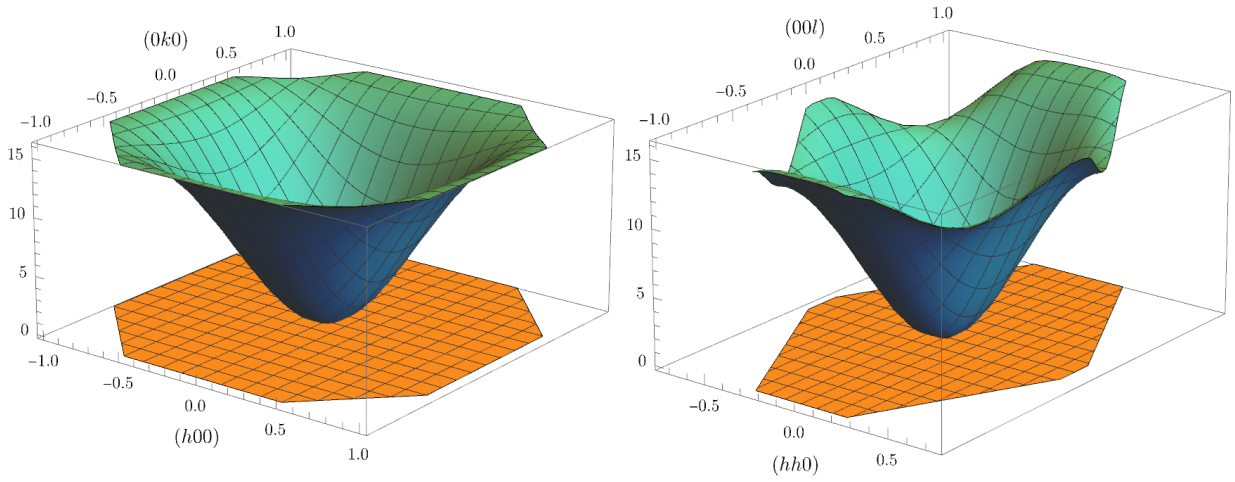


Figure 3.20: The band spectrum of $\mathcal{J}^{(\omega)}/|J_\omega|$ defined in Eq. (3.223), or equivalently the operator $\delta_2 \mathrm{d}_1$. There are two irrotational bands and two dispersive rotational bands which are completely degenerate.

Then the nearest-neighbor model has Hamiltonian $H_Q = (1/2)J\langle Q, Q\rangle_0$, while the vorticity interaction term is $H_\omega = -|J_\omega|\langle \omega, \omega\rangle_2$.

In Fig. 3.20 we show the band structure of the vorticity-vorticity interaction matrix $\mathcal{J}^{(\omega)}/|J_\omega|$, equivalently $\delta_2 \mathrm{d}_1$, where the Fourier transformed adjacency matrices for second and third neighbors are given in Appendix B. These bands should be contrasted with the bands of $\mathrm{d}_0\delta_1$ in Fig. 3.18(a). The most important feature is that there are two degenerate flat bands with zero eigenvalue and two dispersive bands, with all four bands touching at $\boldsymbol{q} = \mathbf{0}$, and the dispersive bands are degenerate in contrast to the bands of $\mathrm{d}_0\delta_1$. Here the two dispersive bands are *rotational* while the flat bands are irrotational. We will discuss these bands in more detail in Chapter 5.

In Fig. 3.21 we compare the dispersive rotational bands of the vorticity interaction term

dumbbell model, it is natural to interpret it as proportional to the charge of the monopoles, but in the pure 1-form theory the term "charge" may be somewhat misleading. It is common in the literature to refer to the divergence as a "gauge charge" and the Coulomb phase as having an "emergent gauge structure" [57, 58], but this terminology is imprecise. In the Hodge theory a different analogy would be to think of a 1-form field as analogous to the flow of a fluid. Then the curl of the velocity field of the fluid is referred to as the vorticity, but the divergence of the velocity field does not have a standard name. The Coulomb phase with zero divergence is then analogous to an incompressible fluid. Excitations in the Coulomb phase take the form of localized point defects of non-zero divergence, which indeed appear as "charges". But they are not, in the classical model, charged under the action of a symmetry group, and are somewhat analogous to a bubble in an incompressible fluid. Of course we can also view the 1-form field as being analogous to an electric field 1-form, but the important point is that there are no 0-form matter fields with internal degrees of freedom which are charged under the action of a symmetry.

$-|J_\omega|\langle\omega,\omega\rangle_2$ (green) to the rotational bands of the dipole-dipole interaction matrix (transparent, yellow), where we normalized the dipole interaction matrix at nearest-neighbor and we chose J_ω to match the bandwidths. Clearly the vorticity interaction term captures the general shape of the dispersion and accurately reproduces the band minimum at the X points (001). This confirms that the vorticity interaction term should be the leading order correction to the dumbbell model which captures the dispersion of the rotational bands. In principle further corrections could be computed for the rotational bands in the form of further-neighbor plaquette-plaquette interactions, and in principle one could formulate a rigorous and systematic expansion for the dipolar spin ice Hamiltonian starting from either the dumbbell or projective equivalence Hamiltonians and adding corrections to the irrotational and rotational modes, though this remains as future work to be done. What we can say for now is that a minimal model for dipolar spin which can serve as an alternative dumbbell approximation would be the following:

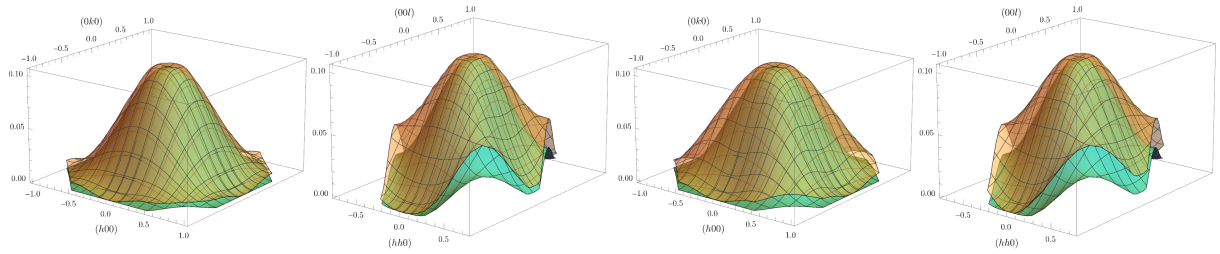
$$H_{\mathrm{DSI}} = \frac{1}{2}J_Q\langle\delta_1 S,\delta_1 S\rangle + \frac{1}{2}J_{\mathrm{PE}}\langle S^{\mathrm{irr}},S^{\mathrm{irr}}\rangle - |J_\omega|\langle \mathrm{d}_1 S,\mathrm{d}_1 S\rangle + \cdots, \tag{3.227}$$

where the last third term is our first order correction to the dumbbell approximation capturing the dispersion of the rotational bands.

Moreover, it gives us new insight into the putative ground state of dipolar spin ice. It was found via Monte Carlo simulations [144] that the dipolar spin ice model (with nearest-neighbor exchange plus dipole interactions) will have a phase transition at a sufficiently low temperature to a long range ordered ground state due to the lifting of the flat band degeneracy. Though this ground state is not observed experimentally due to exponentially growing time scales at low temperature [145, 146], it is interesting to note that it is indeed the ground state predicted by the vorticity interaction term, in which every plaquette has maximal circulation, and thus can be considered a "vortex condensate."

3.5 Conclusion and Outlook

To conclude this chapter, we have demonstrated that the physics of ice models can be naturally described in terms of a 1-form field on a discrete cellular complex. This perspective allows us to unambiguously define the Helmholtz-Hodge decomposition for spin configurations into rotational, irrotational, and harmonic components. This description is exact in the soft spin approximation, but the HHD is well defined for hard spin configurations just as well as soft spin configurations. In the nearest-neighbor model the Hamiltonian is simply proportional to $\langle\delta_1 S,\delta_1 S\rangle$. We demonstrated that the addition of further neighbor interactions can be done in a way that exactly preserves the 1-form description, and demonstrated that the celebrated dumbbell



(a) (b)

Figure 3.21: Bands of the vorticity Hamiltonian with negative coefficient (green) compared to the bands of the dipole-dipole interaction matrix computed by truncating corrections at 100'th neighbor (transparent, yellow). The dipole interaction matrix has been normalized so that the nearest-neighbor coupling is unity. The coefficient J_ω was chosen to match the bandwidths. Since the dipole interaction matrix splits the degeneracy of the two rotational bands, we show separate comparisons in (a) and (b) for the two bands, in both $(hk0)$ and (hhl).

approximation falls precisely within the 1-form description. Furthermore, we showed that the projective equivalence model is the same as the dumbbell model, where the $1/r$ Coulomb potential is replaced by the inverse of the discrete scalar Laplacian. Lastly, we demonstrated that the first order correction to the dumbbell model which captures the dispersion of the rotational modes induced by the dipole interactions is described by a term proportional to $-\langle \mathrm{d}_1 S, \mathrm{d}_1 S \rangle$.

In Chapter 4 we will demonstrate that this correction is observable in polarized neutron scattering experiments, and in Chapter 5 we will consider the complementary model to nearest neighbor spin ice whose Hamiltonian is $J_\omega \langle \mathrm{d}_1 S, \mathrm{d}_1 S \rangle$ with $J_\omega > 0$, which we will show harbors a novel spin liquid ground state which generalizes the Coulomb phase. In Chapter 6 we will utilize some of the insights gained from this chapter to define a HHD for $N = 3$ component Heisenberg spins and utilize it to study fluctuations about a long range ordered ground state.

Numerous questions and avenues for further development remain open. We have applied this method to study spin models on line graph lattices with bipartite parent lattices, but the non-bipartite case remains open to further investigation. There are also various non-linegraph frustrated lattices which may also admit a related Hodge theoretic treatment, in particular Ref. [103] also found closed loop eigenstates on the dice lattice. It is also evident that the class of lattices known as "split lattices" would be directly amenable to a straightforward generalization of our method, because they amount to a line graph lattice with extra lattice sites located at the parent lattice sites. The most well known example is the Lieb lattice, which is a decorated square lattice with extra atoms at the midpoint of each edge. A more recent and exciting example is the so-called "centered pyrochlore" lattice [147], which has extra magnetic degrees of freedom

located at the center of each tetrahedron. The degrees of freedom located on the parent lattice vertices should naturally be treated as zero-forms, with those on the edges treated as 1-forms, and they can naturally be coupled by a term of the for $J'\langle \mathrm{d}_0 S^0, S^1\rangle = J'\langle S^0, \delta_1 S^1\rangle$. In the second form in particular, we see that for $T \ll \mathrm{sgn}(J')$ we would have the constraint $-\mathrm{sgn}(J')S^0 = \delta_1 S^1$, which reads as a Gauss law for the fields. Thus these lattices would have both "matter" and "gauge" fields naturally built in, instead of the "emergent" matter fields of the ice models.

Chapter 4

Polarized Neutron Scattering as a Probe of Flat Band Physics in Spin Ice Systems

4.1 Polarized Neutron Scattering in Spin Ice: Spin Flip and Non Spin Flip Channels

As we can see from Figs. 2.16 and 2.17, structure factors reveal the three dimensional structure of correlations in a magnetic material, because they offer cross sections of the spin spin correlation functions $\langle S_i^\alpha S_j^\beta \rangle$ which is generally the underlying object of interest. This makes neutron scattering an extremely powerful probe for magnetic correlations, because it probes the magnetic susceptibility tensor χ,

$$T\chi^{\alpha\beta}(\boldsymbol{q}) = \sum_{\mu\nu} \langle m_\mu^\alpha(-\boldsymbol{q}) m_\nu^\beta(\boldsymbol{q}) \rangle, \tag{4.1}$$

where the magnetic moment is related to the pseudospin by the g-tensor, $m_i^\alpha = \sum_\beta g_i^{\alpha\beta} S_i^\beta$. We will work with the reduced susceptibility tensor $\tilde{\chi} = T\chi$. Using polarization analysis it is, at least *in principle* and in some cases, possible to reconstruct the full tensor, though such a full reconstruction is extremely involved, see for example Ref. [123].[1] Polarized neutron scattering generally gives access to two of the components, encoded in the spin flip (SF) and non-spin-flip

[1] In particular their Supplementary Information section 1.4.

(NSF) scattering cross sections, proportional to Eqs. (2.104) and (2.105) which can be written as

$$S_{\mathrm{SF}}(\boldsymbol{q}) \equiv \sum_{\alpha\beta} \hat{y}^{\alpha}_{\mathrm{sc}}\, \mathcal{X}^{\alpha\beta}(\boldsymbol{q})\, \hat{y}^{\beta}_{\mathrm{sc}}, \tag{4.2}$$

$$S_{\mathrm{NSF}}(\boldsymbol{q}) \equiv \sum_{\alpha\beta} \hat{z}^{\alpha}_{\mathrm{sc}}\, \mathcal{X}^{\alpha\beta}(\boldsymbol{q})\, \hat{z}^{\beta}_{\mathrm{sc}}, \tag{4.3}$$

where $\hat{\boldsymbol{z}}_{\mathrm{sc}}$ is the polarization axis of the incoming neutron beam and $\hat{\boldsymbol{y}}_{\mathrm{sc}} = \hat{\boldsymbol{z}}_{\mathrm{sc}} \times \hat{\boldsymbol{q}}$, see Fig. 2.15.

In the study of spin liquid physics, polarized neutron scattering is the method of choice when single crystal samples are available to extract as much correlation data as possible from a magnetic compound. In particular in the study of spin ice, polarization analysis was famously used by Fennel et al. [148] to study the compound $Ho_2Ti_2O_7$, who reported the first clear observation of pinch points in spin ice. They reported both the SF and NSF intensities in the (hhl) plane, with pinch points appearing in the SF channel while the NSF channel exhibits a broad diffusive pattern. The pinch points are obscured in the unpolarized scattering cross section, which is the sum of the SF and NSF. Since then, it has become common to find both theoretical and experimental results focusing on the SF channel in the (hhl) plane [120].

In this chapter we will demonstrate that the (hhl) NSF channel, which has received little attention, can be used to isolate the correlations of only the rotational modes, while the SF channel mixes correlations of the rotational and irrotational modes, based on the results published in Ref. [149]. Note that Ref. [149] was published prior to the detailed development of the Hodge theoretic perspective of the previous chapter, and it does not contain the language "rotational" and "irrotational", but rather uses the words "transverse" and "longitudinal". The primary result of this chapter, that the (hhl) NSF probes the rotational modes, will follow almost trivially from the Hodge perspective, so we will keep our presentation brief to highlight the most important results and applications. Our results rely on the g-tensor being Ising, $g_i = g^{\parallel} \hat{\boldsymbol{z}}_i\, \hat{\boldsymbol{z}}_i^T$, where $\hat{\boldsymbol{z}}_i$ points along the local easy axis direction and we take the usual convention that it points from diamond sublattice A to B. We will throughout set $g^{\parallel} = 1$.

4.2 The Flat NSF

In Fig. 4.1, we show the temperature evolution of the SF, NSF, and unpolarized neutron scattering cross sections in the (hhl) plane for the nearest-neighbor model of spin ice computed using the SCGA. While the SF channel shows clear pinch points, the NSF channel is *flat*, i.e. $\boldsymbol{q}$-independent, showing only a uniform intensity that rises monotonically as temperature decreases. The results from Monte Carlo simulations are qualitatively identical [149]. This "flat NSF" of the NNSI model

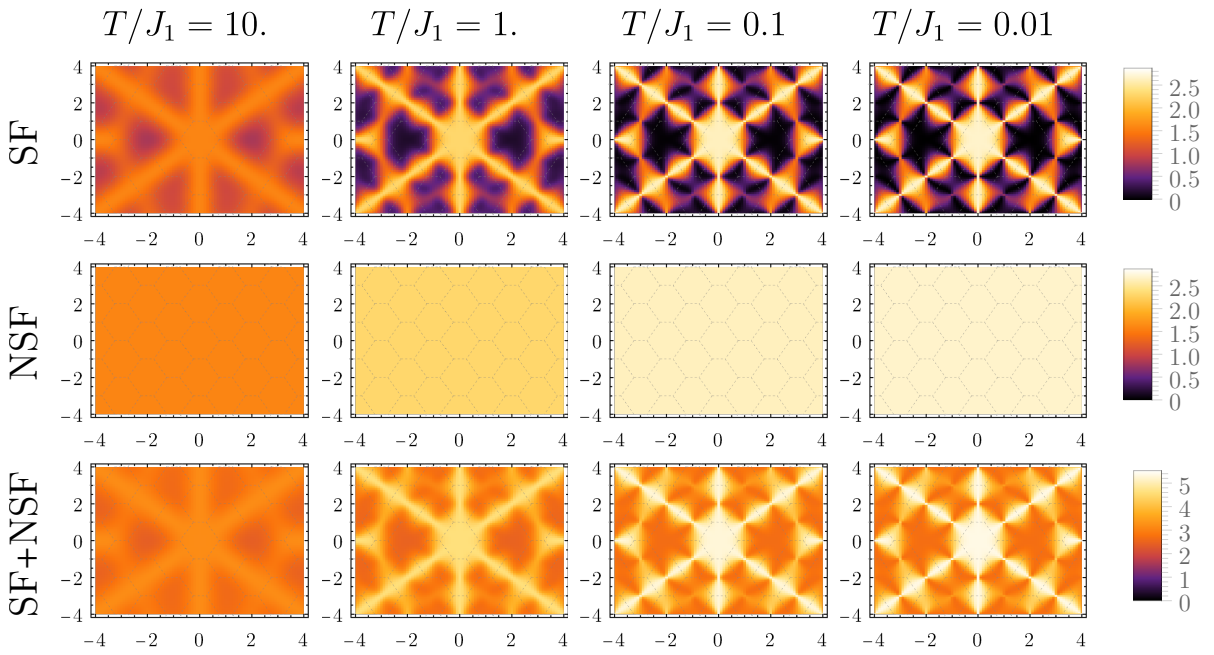


Figure 4.1: The SF, NSF, and unpolarized neutron scattering cross sections in the (hhl) plane for the nearest-neighbor spin ice model computed in the SCGA. The SF channel shows clear pinch points while the NSF is completely flat.

was noted previously in the literature [120, 148, 150–152] but, prior to the work in Ref. [149], had received essentially no scrutiny.

Our first questions is, why is the (hhl) NSF flat in the nearest-neighbor spin ice model? Based on the understanding developed in the previous chapter, we know that spin configurations have a natural Helmholtz-Hodge decomposition into rotational, irrotational, and harmonic components. We know that in the nearest-neighbor model the rotational modes are completely degenerate, corresponding to the two flat bands in the spectrum of the adjacency matrix, and the irrotational modes are energetically suppressed by the fact that the Hamiltonian amounts to $(\operatorname{div} S)^2$, corresponding to the two dispersive bands in the spectrum. At first blush, it is natural to ask: is the dispersionless NSF somehow a reflection of the dispersionless flat bands of the adjacency matrix and the corresponding degeneracy of the rotational modes (cf. Section 3.3.3)? And if so, how can we explicitly see the mechanism that underpins their relation?

To understand this better, consider a fully disordered paramagnet, i.e. the infinite temperature limit. In this limit the spin-spin correlation matrix becomes $\mathcal{G} \to \lambda_0 \mathbb{1}$, where $\lambda_0 \to 1$ at infinite temperature ($\beta \to 0$), i.e. $\langle S_i S_j \rangle = \delta_{ij}$. Correspondingly, we see in Fig. 4.1 that all of the structure

factors become uniform and $\boldsymbol{q}$-independent at high temperature. On the other hand, at zero temperature we have $\mathcal{G} \to 2(\mathcal{P}^{\mathrm{im}\,\delta_2} + \mathcal{P}^{H^1})$, where $\mathcal{P}^{\mathrm{im}\,\delta_2}$ is the identity in the rotational sector. This means that the rotational modes effectively *remain paramagnetic* at all temperatures.

Now, let us apply some heuristic reasoning based on our intuition built up in Chapter 3, which we will later verify to be correct. Consider the spin-flip channel at the lowest temperature ($T/J_1 = 0.010$) in Fig. 4.1, which exhibits the classic pinch point features. Is some regions in the scattering plane we observe zero intensities while in others we observe strong scattering. As we saw in Chapter 3, in the limit of zero temperature the correlation matrix is proportional to the projector to the rotational mode subspace. We can therefore infer that in the regions with zero intensity the SF channel structure factor must be probing the fluctuations of the irrotational modes, which have been thermally depopulated. Conversely, in the regions of maximum intensity, it must probe the correlations of the rotational modes, and in the regions of intermediate intensity it probes the correlations of *both*. If that is the case, then the dispersion in the SF channel at zero temperature can be understood not as reflecting the dispersion of the correlation matrix, which is just a projector and therefore is dispersionless, but rather is a consequence of the fact that at different $\boldsymbol{q}$, this structure factor probes a different linear combination of rotational and irrotational modes. On the other hand the flat NSF surely seems to probe the correlations of the paramagnetic rotational modes at all temperatures and for all $\boldsymbol{q}$, and directly reflects the underlying flatness of the rotational mode correlations. In the remainder of this chapter we will make these arguments precise.

4.3 NSF as Rotational Mode Correlator

Let us now consider what exactly these structure factors are probing in the vector language. From Eqs. (2.104) and (2.105), taking an Ising g-tensor with $g_\parallel = 1$ so that $\boldsymbol{m}_i = S_i \hat{\boldsymbol{z}}$, we can write the SF as

$$S_{\mathrm{SF}}(\boldsymbol{q}) = \left\langle \left| \sum_\mu S_\mu(\boldsymbol{q}) \hat{\boldsymbol{z}}_\mu \cdot \hat{\boldsymbol{y}}_{\mathrm{sc}} \right|^2 \right\rangle, \tag{4.4}$$

$$\equiv \left\langle \left| \sum_\mu \langle \boldsymbol{q}, \mu | S \rangle \, \hat{\boldsymbol{z}}_\mu \cdot \hat{\boldsymbol{y}}_{\mathrm{sc}} \right|^2 \right\rangle, \tag{4.5}$$

$$= \left\langle \left| \left(\sum_\mu \langle \boldsymbol{q}, \mu | \, \hat{\boldsymbol{z}}_\mu \cdot \hat{\boldsymbol{y}}_{\mathrm{sc}} \right) | S \rangle \right|^2 \right\rangle, \tag{4.6}$$

and similarly for the NSF. Thus let us define the vectors (one for each $\boldsymbol{q}$)

$$|\boldsymbol{q}, \mathrm{SF}\rangle \equiv \sum_{\mu} \left(\hat{\boldsymbol{z}}_{\mu} \cdot \hat{\boldsymbol{y}}_{\mathrm{sc}}\right) |\boldsymbol{q}, \mu\rangle \,, \tag{4.7}$$

$$|\boldsymbol{q}, \mathrm{NSF}\rangle \equiv \sum_{\mu} \left(\hat{\boldsymbol{z}}_{\mu} \cdot \hat{\boldsymbol{z}}_{\mathrm{sc}}\right) |\boldsymbol{q}, \mu\rangle \,. \tag{4.8}$$

Note that $\hat{\boldsymbol{y}}_{\mathrm{sc}}$ is $\boldsymbol{q}$-dependent while $\hat{\boldsymbol{z}}_{\mathrm{sc}}$ is not. We can then rewrite the structure factors as measuring the fluctuations of the corresponding modes,

$$S_{\mathrm{SF}}(\boldsymbol{q}) \equiv \langle \boldsymbol{q}, \mathrm{SF}| \, \mathcal{G} \, |\boldsymbol{q}, \mathrm{SF}\rangle \,, \tag{4.9}$$

$$S_{\mathrm{NSF}}(\boldsymbol{q}) \equiv \langle \boldsymbol{q}, \mathrm{NSF}| \, \mathcal{G} \, |\boldsymbol{q}, \mathrm{NSF}\rangle \,. \tag{4.10}$$

At zero temperature, the SF and NSF structure factors are then simply[2]

$$S_{\mathrm{SF}}(\boldsymbol{q}) \to 2 \, \langle \boldsymbol{q}, \mathrm{SF}| \, (\mathcal{P}^{\mathrm{im}\,\delta_2} + \mathcal{P}^{H^1}) \, |\boldsymbol{q}, \mathrm{SF}\rangle \,, \tag{4.11}$$

$$S_{\mathrm{NSF}}(\boldsymbol{q}) \to 2 \, \langle \boldsymbol{q}, \mathrm{NSF}| \, (\mathcal{P}^{\mathrm{im}\,\delta_2} + \mathcal{P}^{H^1}) \, |\boldsymbol{q}, \mathrm{NSF}\rangle \,. \tag{4.12}$$

The harmonic portion only contributes to the intensity at the $\boldsymbol{q} = \boldsymbol{0}$. For each $\boldsymbol{q} \neq \boldsymbol{0}$ we can decompose these vectors as $|\boldsymbol{q}, \mathrm{SF}\rangle = \mathcal{P}^{\mathrm{im}\,\mathrm{d}_0} \, |\boldsymbol{q}, \mathrm{SF}\rangle + \mathcal{P}^{\mathrm{im}\,\delta_2} \, |\boldsymbol{q}, \mathrm{NSF}\rangle$, and we can say that

$$S_{\mathrm{SF}}(\boldsymbol{q}) \to 2 \left| \mathcal{P}^{\mathrm{im}\,\delta_2} \, |\boldsymbol{q}, \mathrm{SF}\rangle \right|^2 , \tag{4.13}$$

$$S_{\mathrm{NSF}}(\boldsymbol{q}) \to 2 \left| \mathcal{P}^{\mathrm{im}\,\delta_2} \, |\boldsymbol{q}, \mathrm{NSF}\rangle \right|^2 , \tag{4.14}$$

i.e these structure factors are simply twice the norm squared of the rotational component of the corresponding vector (harmonic component at $\boldsymbol{q} = \boldsymbol{0}$). Since $\hat{\boldsymbol{y}}_{\mathrm{sc}}$ is $\boldsymbol{q}$-dependent, $|\boldsymbol{q}, \mathrm{SF}\rangle$ rotates at different $\boldsymbol{q}$, and for some $\boldsymbol{q}$ is mostly in the rotatinal sector, where we see maximum intensity in the low temperature SF, and at other $\boldsymbol{q}$ is mostly in the irrotational sector, where we see minimal intensity, while at a general $\boldsymbol{q}$ it will be some linear combination of rotational and irrotational components. The zero-temperature intensity variation of the SF channel therefore is simply twice the magnitude squared of the rotational component of $|\boldsymbol{q}, \mathrm{SF}\rangle$.

On the other hand, the NSF intensity is the same for all $\boldsymbol{q} \in (hhl)$, meaning that the rotational component of $|\boldsymbol{q}, \mathrm{NSF}\rangle$ has a constant magnitude, and indeed implies that $|\boldsymbol{q}, \mathrm{NSF}\rangle$ is entirely contained in the rotational sector since (i) $\hat{\boldsymbol{z}}_{\mathrm{sc}}$ is $\boldsymbol{q}$-independent and (ii) the NSF remains flat at all

[2]To remind the reader: $\mathcal{P}^{\mathrm{im}\,\delta_2}$ is the projector to the rotational mode subspace, and $\mathcal{P}^{H^1}$ is the projector to the harmonic mode subspace.

temperatures, even when the dispersive irrotational modes are populated. We can also see this by computing the magnitude of $|\boldsymbol{q}, \mathrm{NSF}\rangle$,

$$\langle \boldsymbol{q}, \mathrm{NSF}|\boldsymbol{q}, \mathrm{NSF}\rangle = \sum_{\mu} |\hat{\boldsymbol{z}}_{\mu} \cdot \hat{\boldsymbol{z}}_{\mathrm{sc}}|^2 = \hat{\boldsymbol{z}}_{\mathrm{sc}} \cdot \left(\sum_{\mu} \hat{\boldsymbol{z}}_{\mu} \hat{\boldsymbol{z}}_{\mu}^T \right) \cdot \hat{\boldsymbol{z}}_{\mathrm{sc}} = \frac{4}{3}, \tag{4.15}$$

where we used the overcompleteness relation for the $\hat{\boldsymbol{z}}_{\mu}$ used in the previous chapter. Similarly $\langle \boldsymbol{q}, \mathrm{SF}|\boldsymbol{q}, \mathrm{SF}\rangle = 4/3$ for the same reason. Therefore we conclude from Eq. (4.14) that if either $|\boldsymbol{q}, \mathrm{SF}\rangle$ or $|\boldsymbol{q}, \mathrm{NSF}\rangle$ is fully rotational, the zero-temperature intensity is exactly $8/3 \approx 2.67$. Comparing to the intensity scale of Fig. 4.1 we conclude that (i) the (hhl) NSF channel probes *only* rotational modes ($|\boldsymbol{q}, \mathrm{NSF}\rangle$ is entirely contained in the rotational sector away from $\boldsymbol{q} = \boldsymbol{0}$) and (ii) the SF channel probes only rotational modes in the regions where it has maximum intensity at zero temperature, in particular near $\boldsymbol{q} = \boldsymbol{0}$ and along the high-symmetry (hhh) and $(00l)$ lines. The same logic can be applied to all of the low-temperature structure factors shown in Fig. 2.16: their high-intensity regions are probing rotational modes and their low-intensity regions are probing irrotational modes. Note in particular that the NSF in the $(hk0)$ plane is *not* $\boldsymbol{q}$-independent, and therefore must probe a mix of rotational and irrotational modes. Thus the statement that the (hhl) NSF probes *only* rotational modes is *only true in the* (hhl) *plane*. We will see that it is closely tied to the geometry of the lattice.

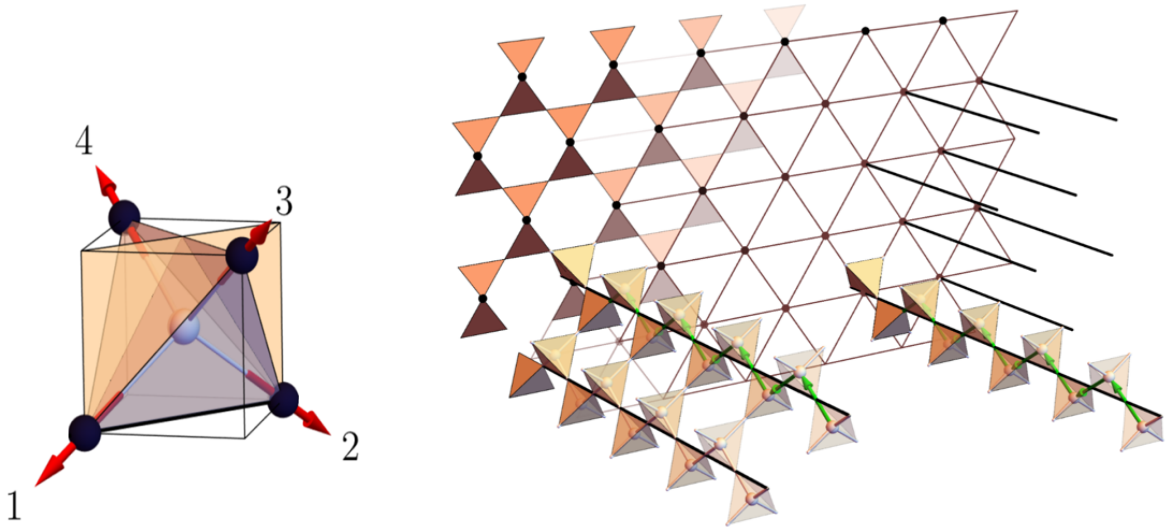
The projection of the SF and NSF vectors to the rotational and irrotational subspace are independent of temperature, and indeed independent of the Hamiltonian. The (hhl) NSF probes purely rotational modes at all temperatures. Since $\mathcal{G}$ is simply λ_0^{-1} times the identity in the rotational sector, we can conclude from Eq. (4.10) that

$$S_{\mathrm{NSF}}(\boldsymbol{q}) = \lambda_0^{-1} \langle \boldsymbol{q}, \mathrm{NSF}|\boldsymbol{q}, \mathrm{NSF}\rangle = \frac{4}{3\lambda_0}. \tag{4.16}$$

Therefore the uniform intensity interpolates from $8/3$ at low temperature ($\lambda_0 = 1/2$) to $4/3$ at high temperature ($\lambda_0 = 1$), as seen in Fig. 4.1.

4.4 NSF as Chain-Chain Correlator

We will now demonstrate why the vectors $|\boldsymbol{q}, \mathrm{NSF}\rangle$ for the (hhl) NSF are purely rotational by considering what they correspond to in direct space. First, let us take a look at the coefficients $\left(\hat{\boldsymbol{z}}_{\mu} \cdot \hat{\boldsymbol{z}}_{\mathrm{sc}}\right)$ in Eq. (4.8). For (hhl) NSF scattering the polarization axis $\hat{\boldsymbol{z}}_{\mathrm{sc}}$ is orthogonal to the (hhl) plane, along $[\bar{1}10]$, cf. Fig. 2.15. In Fig. 4.2(a) we show a single tetrahedron of the pyrochlore lattice, and have labeled the four corners with $\mu = 1, 2, 3, 4$. The cubic $[100]$ and equivalent



(a) (b)

Figure 4.2: (a) A single tetrahedron of the pyrochlore lattice, with the direct-space $[hhl]$ plane shown and sublattices labeled. The local $\hat{z}_\mu$ easy axis vectors are shown in red, where the central diamond lattice site (in the center of the tetrahedron) is an A site. (b) Depiction of the alpha chains of the pyrochlore lattice, which correspond to 1-dimensional cycles of the diamond lattice cell complex wrapping around the periodic boundaries. When projected to the direct-space $[hhl]$ plane they form a triangular lattice.

directions are along the cube edges, and without loss of generality we can take the polarization axis $\hat{z}_{\text{sc}} \parallel [\bar{1}10]$ to point parallel to the edge between corners 3 and 4, pointing from 3 to 4. The corresponding (hhl) plane orthogonal to $[\bar{1}10]$ is then shown as the yellow rectangle bisecting the 3-4 bond and cutting the tetrahedron in half. The red vectors indicate the local easy axis directions $\hat{z}_\mu$ with the convention that they all point from the diamond sublattice A to B.

Notice $\hat{z}_1$ and $\hat{z}_2$ both lie in the direct-space $[hhl]$ plane, orthogonal to the polarization vector $\hat{z}_{\text{sc}} \parallel [\bar{1}10]$. Meanwhile sublattices 3 and 4 have equal and opposite projection onto $\hat{z}_{\text{sc}}$, which we are taking to point from corner 3 to 4. In Cartesian coordinates we can write

$$\hat{z}_3 = (\hat{\boldsymbol{x}} - \hat{\boldsymbol{y}} + \hat{\boldsymbol{z}})/\sqrt{3}, \qquad \hat{z}_4 = (-\hat{\boldsymbol{x}} + \hat{\boldsymbol{y}} + \hat{\boldsymbol{z}})/\sqrt{3}, \qquad \hat{z}_{\text{sc}} = (-\hat{\boldsymbol{x}} + \hat{\boldsymbol{y}})/\sqrt{2}. \tag{4.17}$$

Then we can write Eq. (4.8) as

$$|\boldsymbol{q}, \text{NSF}\rangle = \sqrt{\frac{2}{3}}(|\boldsymbol{q}, 4\rangle - |\boldsymbol{q}, 3\rangle), \tag{4.18}$$

where sublattices 1 and 2 dropped out entirely due to the neutron polarization being orthogonal to their easy axes. Note that this was a consequence of the Ising g-tensor, if $g_\perp \neq 0$ then there would be an additional coupling from sublattices 1 and 2 proportional to $g_\perp$.

Now let us consider the real-space components of the (hhl) NSF vector, by expanding it as

$$\sqrt{\frac{3}{2}}\,|\boldsymbol{q}, \text{NSF}\rangle = \frac{1}{\sqrt{L^3}} \sum_{\boldsymbol{r}} e^{-i\boldsymbol{q}\cdot\boldsymbol{r}} \left(e^{-i\boldsymbol{q}\cdot\boldsymbol{c}_4}\,|\boldsymbol{r}, 4\rangle - e^{-i\boldsymbol{q}\cdot\boldsymbol{c}_3}\,|\boldsymbol{r}, 3\rangle\right) \tag{4.19}$$

$$= e^{-i\boldsymbol{q}\cdot\boldsymbol{c}_3} \frac{1}{\sqrt{L^3}} \sum_{\boldsymbol{r}} e^{-i\boldsymbol{q}\cdot\boldsymbol{r}} \left(|\boldsymbol{r}, 4\rangle - |\boldsymbol{r}, 3\rangle\right), \tag{4.20}$$

where the sum is over unit cells and in the second line we used the fact that $\boldsymbol{q} \in (hhl)$ so that $\boldsymbol{q} \cdot (\boldsymbol{c}_4 - \boldsymbol{c}_3) = 0$. Now we will further utilize the fact that $\boldsymbol{q} \in (hhl)$ to note that any two unit cells with $\boldsymbol{r} - \boldsymbol{r}' \perp (hhl)$ have the same phase factor. In particular, in Fig. 4.2(b) we show that the pyrochlore lattice can be decomposed into "alpha chains" consisting of up and down tetrahedra (three chains are shown) running along $[\bar{1}10]$ [153–155]. There are exactly L^2 of these chains in the lattice. If we write the unit cell positions as linear combinations of FCC primitive vectors $\boldsymbol{r} = \sum_i n_i \boldsymbol{a}_i$, with $n_i \in \{0, \ldots, L-1\}$, then we can taken one of the FCC vectors along the chain and index chains according to the remaining two, and write $\boldsymbol{r} = n_1 \boldsymbol{r}_1 + \boldsymbol{r}_\alpha$, so there are L unit cells in a single chain and L^2 chains with positions given by $\boldsymbol{r}_\alpha = n_2 \boldsymbol{a}_2 + n_3 \boldsymbol{a}_3$. Therefore we can

rewrite the above as

$$e^{i\boldsymbol{q}\cdot\boldsymbol{3}}\sqrt{\frac{3}{2}}\,|\boldsymbol{q},\mathrm{NSF}\rangle = \frac{1}{\sqrt{L^3}}\sum_{\alpha} e^{-i\boldsymbol{q}\cdot\boldsymbol{r}_\alpha}\sum_{\boldsymbol{r}\in\alpha}\left(|\boldsymbol{r},4\rangle - |\boldsymbol{r},3\rangle\right), \tag{4.21}$$

$$= \frac{1}{\sqrt{L^2}}\sum_{\alpha} e^{-i\boldsymbol{q}\cdot\boldsymbol{r}_\alpha}\left(\frac{1}{\sqrt{L}}\sum_{\boldsymbol{r}\in\alpha}\left(|\boldsymbol{r},4\rangle - |\boldsymbol{r},3\rangle\right)\right). \tag{4.22}$$

Note that in $\boldsymbol{q}\cdot\boldsymbol{r}_\alpha$ the only portion of the position $\boldsymbol{r}_\alpha$ that matters is the part in the direct-space $[hhl]$ plane (orthogonal to $[\bar{1}10]$), and these orthogonal displacement vectors form a triangular lattice shown in Fig. 4.2(b) [154–156].

Now the remaining question is what the vector in parentheses in Eq. (4.22) means. As a cochain, $|\boldsymbol{r},\mu\rangle$ represents a unit flux along edge μ in unit cell $\boldsymbol{r}$ with orientation from diamond sublattice A to B, and $-|\boldsymbol{r},\mu\rangle$ represents a flux in the opposite direction. Therefore this vector represents flux along the entire α-chain running in a single direction with uniform intensity on every edge in the chain, which we now recognize as a 1-dimensional contour in the diamond cell complex, indicated by the green arrows in Fig. 4.2(b). Define the normalized vector

$$|\alpha\rangle \equiv \frac{1}{\sqrt{2L}}\sum_{e\in\alpha}|e\rangle, \tag{4.23}$$

where the edges are all taken to be oriented along the 1-dimensional contour and where we have normalized by the number of edges in the chain. Thus we can express the (hhl) NSF vector as

$$|\boldsymbol{q},\mathrm{NSF}\rangle = \sqrt{\frac{4}{3}}\frac{1}{\sqrt{L^2}}e^{-i\boldsymbol{q}\cdot\boldsymbol{r}_\alpha}\sum_{\alpha}|\alpha\rangle, \tag{4.24}$$

where we dropped the overall phase factor as it is a gauge choice (in other words, without loss of generality we can take the origin to be located on corner 3 so that $\boldsymbol{c}_3 = \boldsymbol{0}$ in Eq. (4.20)).

The (hhl) NSF structure factor can now be written as

$$S_{\mathrm{NSF}}(\boldsymbol{q}) = \frac{4}{3}\frac{1}{L^2}\sum_{\alpha,\alpha'} e^{-i\boldsymbol{q}\cdot(\boldsymbol{r}'_\alpha - \boldsymbol{r}_\alpha)}\,\langle\alpha|\,\mathcal{G}\,|\alpha'\rangle, \tag{4.25}$$

which is simply $4/3$ times the Fourier transform of the chain-chain correlation function

$$\langle\alpha|\,\mathcal{G}\,|\alpha'\rangle = \left\langle(\langle\alpha|S\rangle)(\langle\alpha'|S\rangle^*)\right\rangle \sim \left\langle\oint_\alpha S\oint_{\alpha'} S\right\rangle. \tag{4.26}$$

We can now clearly understand why the (hhl) NSF probes only the rotational and harmonic modes. Since the alpha chains are closed 1-cycles, $\langle \alpha | S^{\text{irr}} \rangle = 0$, meaning that the (hhl) NSF is *entirely insensitive to the irrotational sector*. Since $\mathcal{G}$ is just $\lambda_0^{-1}(\mathcal{P}^{\text{im}\,\delta_2} + \mathcal{P}^{H^1})$ in the rotational and harmonic sectors, we have

$$\langle \alpha | \mathcal{G} | \alpha' \rangle = \lambda_0^{-1} \langle \alpha | \alpha' \rangle = \lambda_0^{-1} \delta_{\alpha\alpha'}, \tag{4.27}$$

i.e. the circulation around different alpha chains is completely uncorrelated in the nearest-neighbor model, since all of the rotational and harmonic modes are completely degenerate. This direct-space decorrelation translates to the flatness of the NSF response in reciprocal space,

$$S_{\text{NSF}}(\boldsymbol{q}) = \frac{4}{3}\frac{1}{L^2}\sum_{\alpha,\alpha'} e^{-i\boldsymbol{q}\cdot(\boldsymbol{r}'_\alpha - \boldsymbol{r}_\alpha)} \lambda_0^{-1} \delta_{\alpha\alpha} = \frac{4}{3\lambda_0}. \tag{4.28}$$

The uniform intensity thus reflects the correlations of spins only the intra-chain correlator of spins, $\langle \alpha | \mathcal{G} | \alpha \rangle$.

4.5 The Non-Flat NSF of Dipolar Spin Ice

We have now established that the (hhl) NSF probes only the rotational and harmonic sectors and is insensitive to the irrotational sector because it is a circulation-circulation correlation function as expressed by Eq. (4.26). In the nearest-neighbor model where the rotational modes are perfectly degenerate, the NSF is flat throughout (hhl) since the circulation around different alpha chains is uncorrelated (indeed we showed in Section 3.4.6 the more general result that all non-overlapping cycles are decorrelated).

As we discussed in Section 3.4.10, dipolar interactions lift the flat band degeneracy, preferring for the minimal hexagonal loops to maximize their circulation. This induces correlations in the rotational sector reflecting this preference for short loops over long ones. This causes the alpha chains to become correlated and induce a non-flat NSF. Looking at Fig. 4.2(b), or Fig. 3.19, we can see that polarizing a single hexagon will cause the spins in neighboring alpha chains to become anti-ferromagnetically correlated, since they sit on opposite sides of the hexagon. Thus we should expect that dipole interactions will induce an increase in intensity near the zone boundary in the (hhl) NSF.

In Fig. 4.3, we show the SF, NSF, and unpolarized neutron scattering structure factors for a pure dipole model computed in the SCGA, which shows that the NSF indeed becomes non-flat at low temperature, with extra scattering near the zone boundary. The pattern observed in the

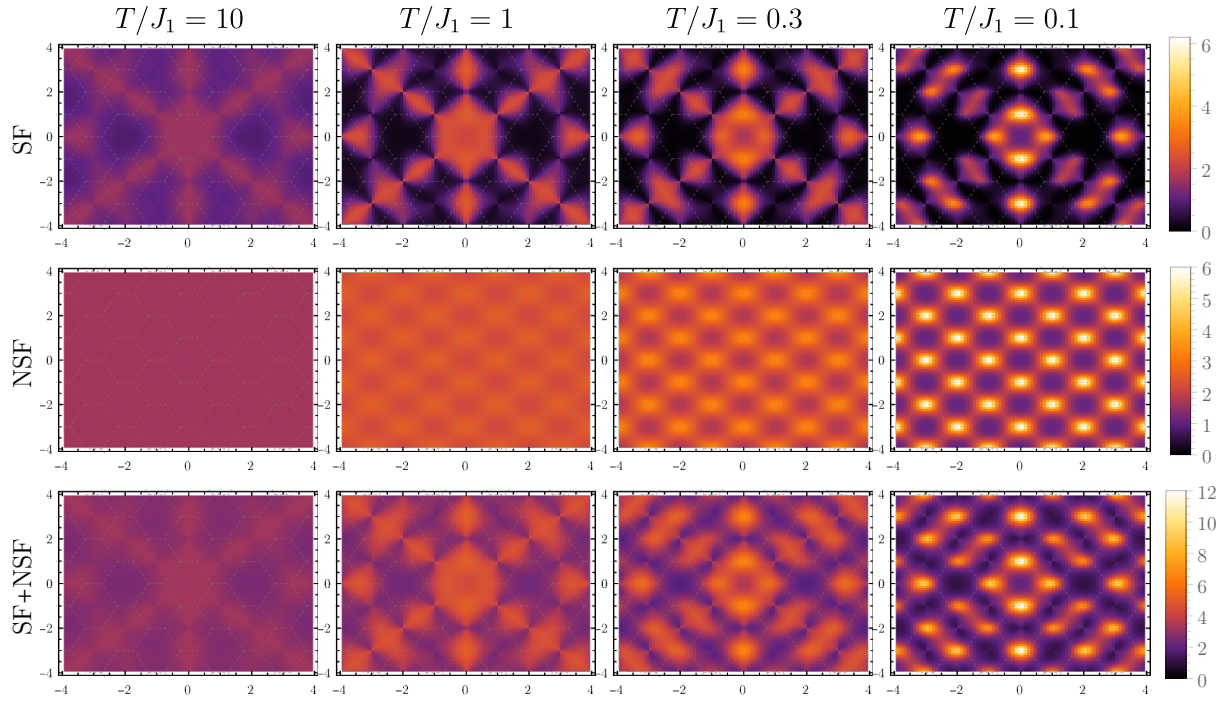
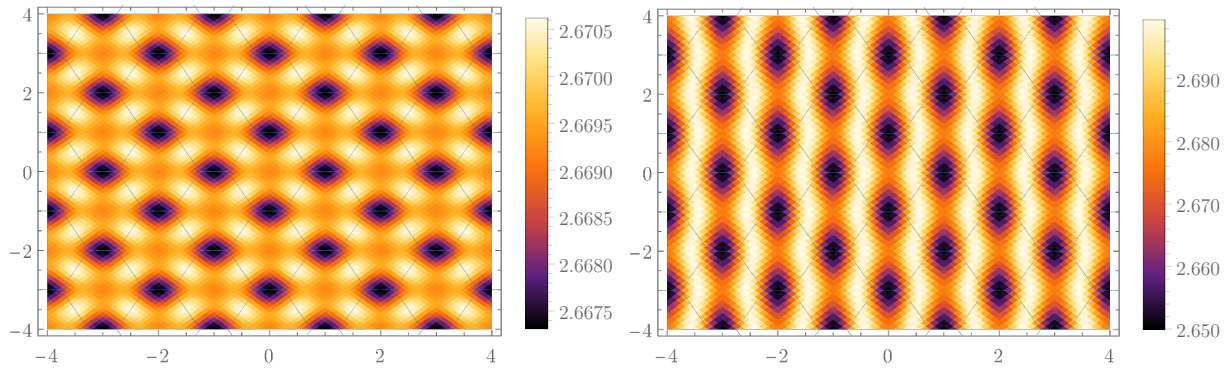


Figure 4.3: The spin flip, non-spin-flip, and unpolarized neutron scattering cross sections computed for the dipolar spin ice model in the SCGA. The non-flat NSF reflects the corrections to the flat bands which lift the rotational mode degeneracy discussed in Section 3.4.10. Notice also the presence of pinch points in the SF even at high temperatures: in the high-temperature series expansion at lowest order one would already obtain pinch points from the fact that the dipolar interaction matrix itself is singular.



(a) (b)

Figure 4.4: (a) The (hhl) NSF of NNSI measured in Monte Carlo at $T/J = 0.01$ for $L = 16$ with 10^{10} samples. (b) $1/N$ corrected (hhl) NSF at $T/J = 0.01$ computed for $L = 12$.

NSF is qualitatively the same as that observed the data presented by Fennel et al. [148], with minimum intensity near the zone center and maximum intensity near the zone boundary forming rounded rectangular features around the X point (001). We also observe the same basic effect in the SF channel in the first Brillouin zone around $\boldsymbol{q} = \mathbf{0}$, which we argued is mostly sensitive to the rotational modes based on the low temperature SCGA behavior of the SF in Fig. 4.1. In particular this means that the reduction in intensity in the SF channel near the zone center and the increased intensity at the X point is caused by the same mechanism as the dispersion in the NSF.

4.6 Hard Spin Corrections to the Flat NSF

Before concluding this chapter, we briefly comment on the consequences of the hard spin constraint. Up until now, we have utilized the SCGA extensively, and we know that the SCGA provides a generally excellent approximation for the physics of spin ice, with the primary exception that the minimal excitations at low energy in an Ising model are single spin flips, which cost some finite minimal energy ϵ and are therefore exponentially suppressed at low temperatures by a Boltzmann weight $\exp(-\beta\epsilon)$. This simply means that the Ising system reaches its "ground state" at a higher temperature than the soft spin system, in which arbitrarily small excitations can be made.

The flatness of the NSF, however, provides a highly sensitive probe that we can use to study the (very small) corrections to the SCGA coming from the hard spin constraint. In Fig. 4.4(a), we show the NSF for the nearest-neighbor model measured to extremely high precision in Monte Carlo simulations using the cluster algorithm of Ref. [128] for an $L = 16$ system, sampling 10^{10} samples

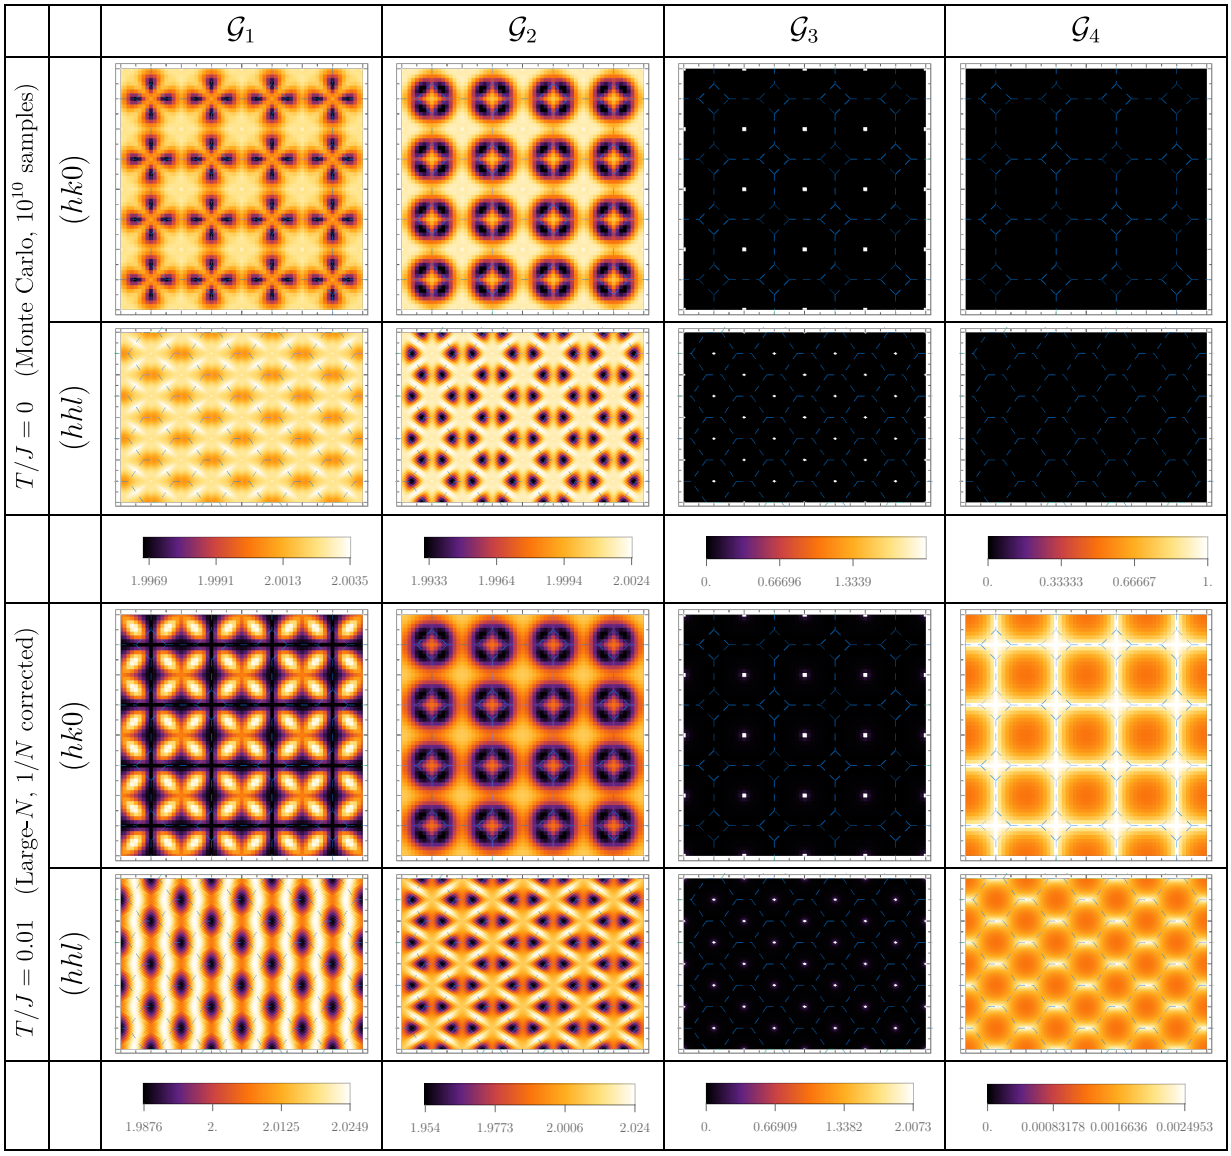


Figure 4.5: Eigenvalues of the spin-spin correlation matrix measured in Monte Carlo for NNSI at at zero temperature with $L = 16$ averaged over 10^{10} samples, compared with the Large-N, $1/N$ corrected eigenvalues at $T/J = 0.01$ with the same system size. Eigenvalues at each $\boldsymbol{q}$ are sorted by magnitude. The first and second column show weak deviations from the SCGA expectation of flat bands with eigenvalue $\lambda_0^{-1} \to 2$ at zero temperature, with qualitative agreement in the second column.

at $T/J = 0.01$. This is effectively equivalent to zero temperature, since the minimal excitation—a pair of charges—costs energy $8J$, and therefore occurs with probability $\exp(-800)$ at $T/J = 0.01$. This means that we are *only* exploring the rotational sector, having suppressed all irrotational excitations completely. We see a clear pattern emerging, with a slight reduction of intensity near the X point. In Fig. 4.4(b), we show the $1/N$ corrected NSF intensity at $T/J_1 = 0.01$, which shows a qualitatively similar pattern, at least regarding a lower intensity near the X point. This implies that the hard spin constraint induces a weak entropic correction to the complete degeneracy of the rotational modes. We much reiterate that the correction is *extremely* small, the deviation from the SCGA value of $8/3 \approx 2.667$ is on the order of 0.1% in Monte Carlo, or one part in 10^3.

To investigate this further, we can directly compute the eigenvalues of the correlation matrix $\mathcal{G}_{ij} = \langle S_i S_j \rangle$. In Fig. 4.5(a), we show the eigenvalues as measured in Monte Carlo for the same system size and number of samples, in both the $(hk0)$ and (hhl) planes. There are four eigenvalues per $\boldsymbol{q}$. In the SCGA, the zero temperature correlation matrix is exactly $2(\mathcal{P}^{\mathrm{im}\,\delta_2} + \mathcal{P}^{H^1})$, so the eigenvalues are 2 in the rotational and harmonic sectors and 0 in the irrotational sector. We can see the two of the eigenvalues are indeed zero everywhere except at $\boldsymbol{q} = \boldsymbol{0}$, corresponding to the suppressed irrotational modes, and the other two eigenvalues are *approximately* 2. We observe very weak but clear corrections to the rotational mode eigenvalues. In Fig. 4.5(b), we show the eigenvalues computed at $1/N$ at $T/J_1 = 0.01$. Due the finite temperature we observe some weak dispersion in the irrotational sector eigenvalues. More importantly, we see that the $1/N$ correction qualitatively reproduces at least some of behavior of the eigenvalues in the rotational sector, especially in the second eigenvalue (eigenvalues are sorted from largest to smallest).

We conclude that (i) the hard spin constraint does indeed modify the correlations of rotational modes, and this can be seen in the NSF as a weak reduction of intensity near the X point (001), and (ii) these corrections can be qualitatively captured by the $1/N$ expansion. In particular, we note that the reduction in intensity in the NSF near the X point is opposite of the behavior in dipolar spin ice, and implies a weak preference for neighboring alpha chains to align rather than anti-align. This is rather interesting, because it is known that entropic effects induce a weak $1/r$ potential for monopoles [157], which is the same as the effect of the dipole interaction in the irrotational sector, but it appears that the entropic effects in the rotational sector have the opposite effect to dipole interactions.

4.7 Conclusion and Outlook

In this chapter we showed that the (hhl) NSF channel is insensitive to the irrotational sector, because it probes a circulation-circulation correlator. The analogous statement in vector calculus

would be to say that

$$\oint_C \boldsymbol{E} \cdot \mathrm{d}\boldsymbol{\ell} = \oint_C (\nabla\phi + \nabla \times \boldsymbol{A} + \boldsymbol{E}_0) \cdot \mathrm{d}\boldsymbol{\ell} = \oint_C (\nabla \times \boldsymbol{A} + \boldsymbol{E}_0) \cdot \mathrm{d}\boldsymbol{\ell}, \tag{4.29}$$

i.e. a line integral around a closed loop is insensitive to the irrotational component of the field. Since in the nearest-neighbor model the rotational sector is degenerate, with the correlations of a paramagnet at all temperatures, the (hhl) NSF appears completely featureless and flat. Using this, we showed that the non-flat NSF of dipolar spin ice is sensitive only to the portion of the dipole interaction which lifts the flat band energy, and reflects the energetic preference for spins to be polarized into small loops.

These results apply to any spin model on the pyrochlore lattice, so long as the g-tensor is Ising, so that we can polarize the neutrons orthogonal to the moments of two of the sublattices, resulting in the isolation of the alpha chains. Given our understanding of the classical model, it would be worthwhile to perform a detailed study of how quantum effects can correlate the rotational modes and the resulting effect in the NSF channel. In particular, in quantum spin ice the rotational modes will become dispersive due to ring exchange, and the linearly dispersing photon (which is a collective excitation of rotational modes) should cause some resulting dispersion in the NSF. One Quantum Monte Carlo study of quantum spin ice did find dispersion in the NSF [150], which they find agrees with an analytical calculation of the correlations in the ring exchange model due to Benton [60].

In particular, our results are applicable to the dipolar-octupolar spin ice compounds which necessarily have Ising g-tensors [36, 158], and in particular to the quantum spin ice candidate $Ce_2Zr_2O_7$ [159–162]. Recent polarized neutron scattering measurements show a dispersive NSF whose correlations appear qualitatively similar to those of $Ho_2Ti_2O_7$ [162], despite Ce^{3+} having a much smaller magnetic moment. Recent work also report via exact diagonalization an increased intensity at the X point in neutron scattering structure factors for dipolar-octupolar models in the π-flux octupolar quantum spin liquid phase [163]. A more systematic study of the effects of quantum correlations in the rotational sector and the corresponding effects in neutron scattering, especially in the (hhl) NSF channel where one can isolate rotational sector correlations, is certainly warranted, and could serve as a sharp diagonistic for quantum spin ice physics.

It would also be of interest to extend our results to study the effects of quenched random disorder in spin ice within the context of the Hodge theoretic treatment, by exploiting the sensitivity of the NSF channel to the effects of disorder in the rotational bands [164–166]. Furthermore, the possibility for high-resolution energy-resolved polarized scattering experiments is on the horizon [167], which raises the possibility that one could perform a targeted study of the dynamics of the rotational modes.

Chapter 5

The Vorticity Model

At the end of Chapter 3, in Section 3.4.10, we introduced the vorticity 2-form $\omega = \mathrm{d}_1 S$ and "charge" 0-form $Q = \delta_1 S$. We showed that the nearest-neighbor spin ice model has Hamiltonian given by $(J/2)\langle Q, Q\rangle$, and that therefore the zero-temperature Coulomb phase consists of purely rotational and harmonic modes, which are exactly degenerate, and its excitations are "monopoles" [119], localized 0-dimensional irrotational defects, described by the 0-form "matter field" Q. In this chapter we introduce a new Ising model, whose Hamiltonian is $(J/2)\langle \omega, \omega\rangle$, which is the natural complement of the nearest-neighbor ice model. We dub it the vorticity model for obvious reasons. We shall find that the ground state of this model is a spin liquid phase similar to the Coulomb phase, but now with the *irrotational* modes fully degenerate. We will find that this phase appears to be a novel type of classical spin liquid hosting 1-dimensional string-like excitations rather than 0-dimensional point-like excitations.

Before we discuss the vorticity model in detail, it will be helpful to elaborate on the physics of spin ice in a way that we have not discussed yet, but which will offer a helpful juxtaposition for our discussion of the vorticity model. The ground states of spin ice obey the "ice rules", with two moments pointing in and two out of each tetrahedron of the pyrochlore lattice. In the 1-form description, this is stated as the zero-divergence statement $\delta_1 S = 0$, which is equivalent in terms of the Hodge decomposition to $S^{\mathrm{irr}} = 0$. Now, any purely rotational configuration can be written as a linear combination of the form

$$\left|S^{\mathrm{rot}}\right\rangle = \sum_p c_p \delta_2 \left|p\right\rangle, \tag{5.1}$$

for some coefficients c_p, where p indexes the plaquettes of the lattice. The 1-form $\delta_2 \left|p\right\rangle \equiv \left|\partial p\right\rangle$ represents a circulation around the boundary of the plaquette with each edge carrying unit

magnitude, and zero on all other edges of the lattice. Alternatively, we could just as well use larger loops. Any configuration of the form

$$|S\rangle = \sum_C c_C |C\rangle , \tag{5.2}$$

where C are some 1-dimensional cycles in the cellular complex, is a rotational configuration, plus a harmonic part if some of the cycles wind around the holes of the complex. This fact is exploited in the loop algorithm for spin ice [124, 125]: taking a hard spin configuration $|S\rangle$, where each $S(e) = \pm 1$, and identifying any oriented contour C of head-to-tail spins (i.e. $S(e) = +1$ along each edge in the oriented contour) then the spin configuration $|S\rangle - 2|C\rangle$ is another hard spin configuration, where all spins are reversed along the contour, which has the same energy so long as the rotational modes are degenerate. Thus the 2-in/2-out ground states of spin ice can be decomposed non-uniquely into a superposition of closed, oriented strings [143]. The Coulomb phase can therefore be described heuristically as a "soup of closed strings".

On the other hand, excited irrotational modes are linear combinations of $\left|\partial^\dagger v\right\rangle$. But since the 0-dimensional excitations are emergent from the underlying 1-form field, it is not possible to have an isolated point defect, it is always attached to another point defect by a 1-dimensional string (see Fig. 1.1(a), which shows that a single spin flip creates two irrotational defects), i.e. excited configurations can be written as linear combinations of open strings $\sum_L c_L |L\rangle$ which have non-zero divergence only at their endpoints. The minimum excitation for a hard spin ground state $|S\rangle = \sum_e S(e)|e\rangle \in \ker \delta_1$ is to flip a single spin $S(e) \to -S(e)$, which is equivalent to adding to the configuration $-2|e\rangle$, which by itself is an open string along edge e. We can also take any open string of head-to-tail spins $|L\rangle$ and flip it, $|S\rangle \to |S\rangle - 2|L\rangle$, which produces two defects, one at either end. To sum up, the Coulomb phase is a soup of closed strings and its excitations are created by breaking a string open, producing two point defects on the boundary of the string with opposite divergence. The strings do not carry any energy, i.e. their string tension is zero, meaning that the 0-dimensional excitations can be separated at no energy cost by making the string longer and are *deconfined* (can be separated to infinite distance at finite energy cost). We will contrast this with the vorticity model, which can be described as a "soup of closed membranes", whose excitations are 1-dimensional strings living on the boundary of an open membrane.

5.1 The Vorticity Model

The vorticity model is defined by the Hamiltonian

$$H_\omega = \frac{J_\omega}{2}\langle S, \delta_2 \mathrm{d}_1 S\rangle = \frac{J_\omega}{2}\langle \omega, \omega\rangle_2 = \frac{J_\omega}{2}\sum_p \left(\oint_{\partial p} s\right)^2 , \tag{5.3}$$

with $J_\omega > 0$ so that the ground states satisfy the zero-circulation constraint $\oint_{\partial p} S = 0$ on every plaquette. As such, this model is a natural complement of NNSI, and its low temperature phase is therefore (at least naïvely) a complement to the Coulomb phase.[1] As a spin model it can be written as

$$H_\omega = \frac{1}{2} \sum_{ij} S_i \mathcal{J}_{ij}^{(\omega)} S_j, \tag{5.4}$$

where

$$\mathcal{J}_{ij}^{(\omega)} = J_\omega \left[6\delta_{ij} - 2\mathcal{A}_{ij}^{(1)} + \mathcal{A}_{ij}^{(2)} - \mathcal{A}_{ij}^{(3b)} \right], \tag{5.5}$$

where $\mathcal{A}_{ij}^{(n,s_n)}$ are further-neighbor adjacency matrices. The second and third-neighbors-type-b are depicted in Fig. 3.19. Thus the ground states of H_ω have zero circulation around every plaquette. The space of ground states is therefore $\operatorname{im} \mathrm{d}_0 \oplus H^1$, the space of irrotational and harmonic configurations, i.e. ground states can be expressed as the gradient of a scalar function plus a harmonic part, or in other words their HHD has zero rotational component. Our task for this chapter is to investigate the thermodynamic properties of this model and to characterize this exotic complementary Coulomb phase.

5.2 Band Structure, Long Wavelength Expansion, and Gaussian Correlations

The bands of the interaction matrix $\delta_2 \mathrm{d}_1$ entering Eq. (5.3) are shown in Fig. 3.20, with two dispersive rotational bands and two flat irrotational bands. Unlike the nearest-neighbor spin ice model, the two dispersive bands are degenerate and the band touching point is *four*fold degenerate. This key distinction means that the all-in/all-out mode $\sum_{\mu_1} |\mathbf{0}, \mu_1\rangle$, which is the band maximum of $\mathrm{d}_0\delta_1$ and thus the "most irrotational" mode, is a ground state mode in addition to the three harmonic modes, which correspond to the bulk magnetization.

Let us consider the long-wavelength expansion of the Hamiltonian. We can follow the procedure detailed in Section 2.5.2, but we do not need to do such detailed analysis since $\delta_2 \mathrm{d}_1$ shares eigenvectors with $\mathrm{d}_0\delta_1$, and in particular with the rotational and irrotational projectors which we already obtained in Section 3.4.2. In particular, we know that the eigenvectors of the two rotational bands at long wavelength are given by the two components of the curl $\sum_{\mu_1} (\hat{\boldsymbol{q}} \times \hat{\boldsymbol{\delta}}_{\mu_1}) |\boldsymbol{q}, \mu_1\rangle$. Therefore, we can immediately deduce the long-wavelength form of the two

[1] To wit: the Coulomb phase of NNSI is characterized by zero divergence, while the zero-temperature phase of the vorticity model is characterized by zero curl.

quadratically dispersing rotational bands,

$$H \propto \frac{J_\omega}{2} \sum_{\boldsymbol{q}} |\boldsymbol{q} \times \boldsymbol{B}(\boldsymbol{q})|^2, \tag{5.6}$$

where $\boldsymbol{B}(\boldsymbol{q}) = \sum_{\mu_1} \hat{\boldsymbol{\delta}}_{\mu_1} \langle \boldsymbol{q}, \mu_1 | S \rangle$, which should be contrasted with the expansion for nearest-neighbor spin ice containing $|\boldsymbol{q} \cdot \boldsymbol{B}|^2$. If we perform the same long-wavelength expansion procedure laid out in Section 2.5.2, this corresponds to an action

$$\beta H \sim \frac{1}{2} \int \mathrm{d}^3 \boldsymbol{r} \left[\lambda_0 (Q^2 + |\boldsymbol{B}|^2) + \xi^2 |\nabla \times \boldsymbol{B}|^2 \right], \tag{5.7}$$

where $\xi = 8a\sqrt{\beta J}$, where a is the pyrochlore nearest-neighbor distance.[2] Notice that there is a key distinction here between this expression and Eq. (2.91): the presence of Q^2. In spin ice there is only one quadratically dispersing band from the harmonic threefold band touching point, corresponding to the divergence of $\boldsymbol{B}$, and the upper irrotational band corresponding to Q can be integrated out at low energies. Here the two physical components of $\nabla \times \boldsymbol{B}$ form two degenerate quadratically dispersing bands from a fourfold band touching point, and Q and $\nabla \cdot \boldsymbol{B}$ are the coarse-grained eigenmodes of the two flat irrotational bands; there is no upper band to be integrated out.

In the SCGA, the zero-temperature correlation matrix simply becomes proportional to the projector to the irrotational bands, $\mathcal{G} \to \lambda_0^{-1}(\mathcal{P}^{\mathrm{im}\, \mathrm{d}_0} + \mathcal{P}^{H^1})$. In Fig. 5.1 we show the structure factors computed in the SCGA for the vorticity model, which should be directly compared to those in Fig. 2.16 for NNSI. At low temperature, the two are essentially just inverted, since the correlation matrix for NNSI at zero temperature is $\lambda_0^{-1}(\mathcal{P}^{\mathrm{im}\, \delta_2} + \mathcal{P}^{H^1})$. We observe that the pinch points now run in the *opposite* direction, i.e. these are "longitudinal" pinch points, being pinched along the *transverse directions*. In Fig. 5.2 we show a 3 dimensional visualization of the pinch point at (220), which can be contrasted with Fig. 2.17(b).

The other difference is the presence of peaks in the "hollowed out" regions, such as those clearly visible in the $(hk0)$ SF structure factor and especially those in the (hhl) NSF. At zero temperature these peaks reflect the $\boldsymbol{q} = \boldsymbol{0}$ part of the correlation function, and appear at *every* zone center. In spin ice, these peaks would sit directly at the center of each pinch point, and some of the peaks would be suppressed by the removal of the zero-momentum AIAO mode which corresponds to the band maximum in the interaction matrix. In this model, all four bands are degenerate at $\boldsymbol{q} = \boldsymbol{0}$, including the AIAO mode, and therefore we get a peak at every zone center, including the ones where they would be absent in spin ice. Note in particular the (hhl) NSF channel. We know from the results of Chapter 4 that this is only sensitive to the rotational modes and two of the three

[2] Note this is eight times the lengthscale obtained for NNSI.

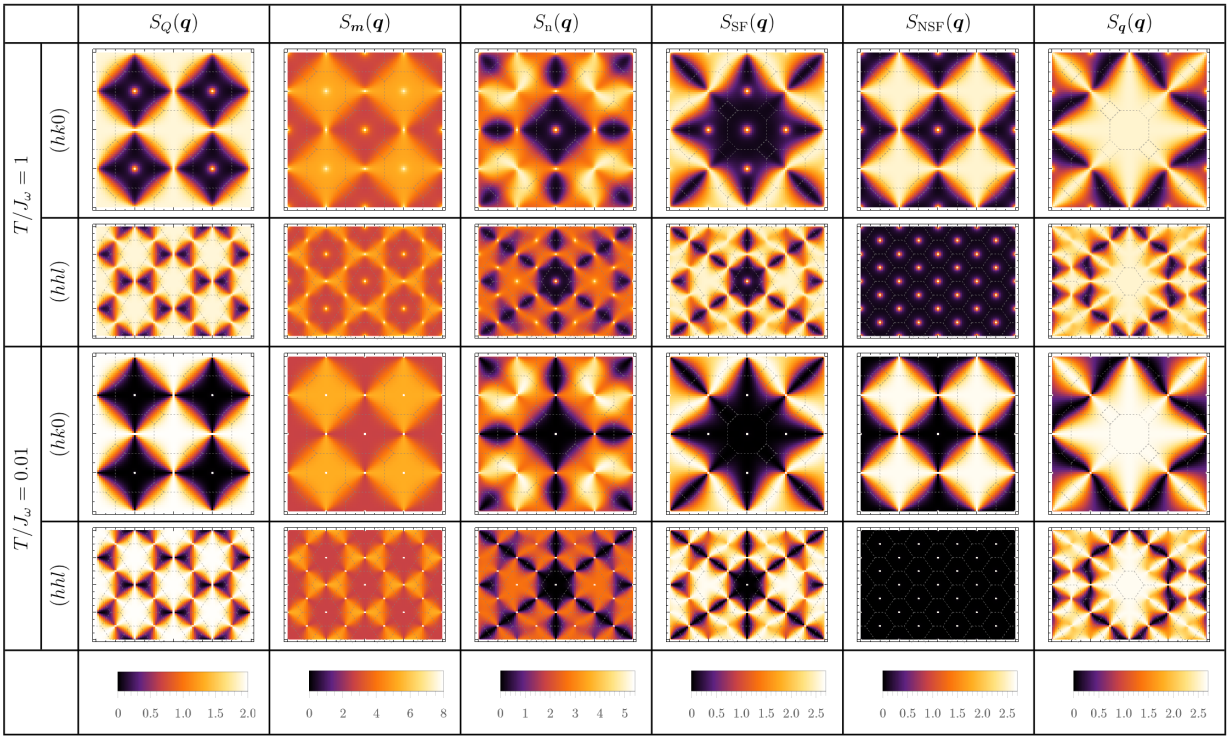


Figure 5.1: Structure factors for the vorticity model computed in the SCGA, which should be contrasted with those of the nearest neighbor spin ice model in Fig. 2.16.

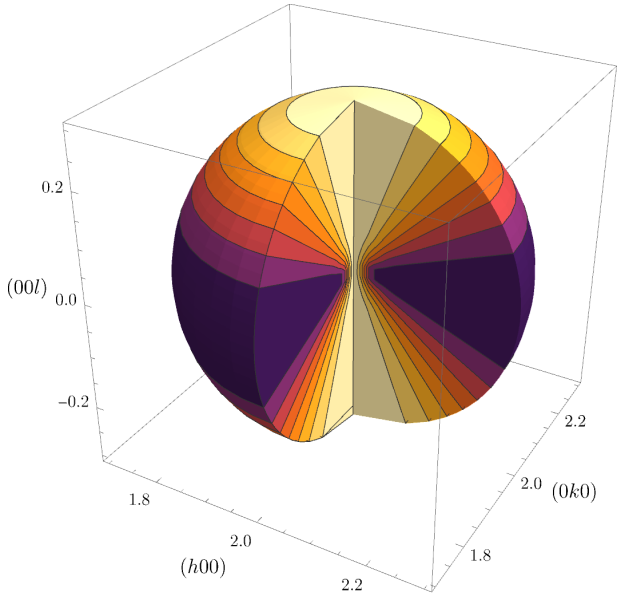


Figure 5.2: The (220) pinch point in $S_Q(\boldsymbol{q})$ for the vorticity model at $T/J_\omega = 0.01$, demonstrating that the pinch is along two directions, corresponding to the suppression of the two rotational modes, and intensity remains along the longitudinal direction. Compare to Fig. 2.17(b)

harmonic modes away from $\boldsymbol{q} = \boldsymbol{0}$. The intensity at $\boldsymbol{q} = \boldsymbol{0}$ approaches $8/3$ at zero temperature just like in spin ice, but away from zone center is completely suppressed as the rotational modes thermally depopulate. At higher temperature, where the lower energy long wavelength rotational modes begin to populate we can observe that these peaks begin to broaden. The peaks in the other structure factors broaden at higher temperatures for the same reason.

Clearly, within the SCGA, the low temperature phase is akin to a Coulomb phase, with the roles of the rotational and irrotational modes swapped, though there is the important difference that all four bands contribute in the low-energy long-wavelength limit, including the AIAO mode. This classical spin liquid phase appears to be quite exotic when contrasted with the usual Coulomb phase, which is analogous to an incompressible fluid, where spins form closed loops with zero-dimensional excitations acting as sources of local divergence. In this phase instead, the divergence is entirely unconstrained and free to fluctuate, and the "field lines" never form closed loops. Indeed we can easily compute the correlations of the local divergence, since $\mathcal{G}$ is simply

$\lambda_0^{-1}\mathcal{P}^{\text{im}\,\mathrm{d}_0}$ in the irrotational sector

$$\langle((\delta_1 S)(v)\delta_1 S(v')\rangle = \left\langle \partial^\dagger v \right| \mathcal{G} \left| \partial^\dagger v' \right\rangle \tag{5.8}$$

$$= \lambda_0^{-1} \left\langle \partial^\dagger v \right| \mathrm{d}_0 (\Delta_0)^{-1} \delta_1 \left| \partial^\dagger v' \right\rangle \tag{5.9}$$

$$= \lambda_0^{-1} \langle v | \, (\delta_1 \mathrm{d}_0)(\delta_1 \mathrm{d}_0)^{-1} (\delta_1 \mathrm{d}_0) \, | v' \rangle \tag{5.10}$$

$$= \lambda_0^{-1} \langle v | \, (\Delta_0) \, | v' \rangle \,, \tag{5.11}$$

so the divergence-divergence correlations are just given by the matrix elements of the scalar Laplacian. The corresponding divergence-divergence correlation structure factor that we defined in Eq. (3.183) looks just like like the high-temperature limit of Fig. 3.15 times λ_0^{-1} (since in the high-temperature limit the correlation matrix becomes the identity). Thus, the divergence appears to be maximally disordered at all temperatures (despite the long range correlations of the spins themselves). That is to say, the irrotational modes remain paramagnetic at all temperatures, just as the rotational modes of spin ice remain paramagnetic at all temperatures.

5.3 The 2-Coulomb Phase: Membranes and Vortex Strings

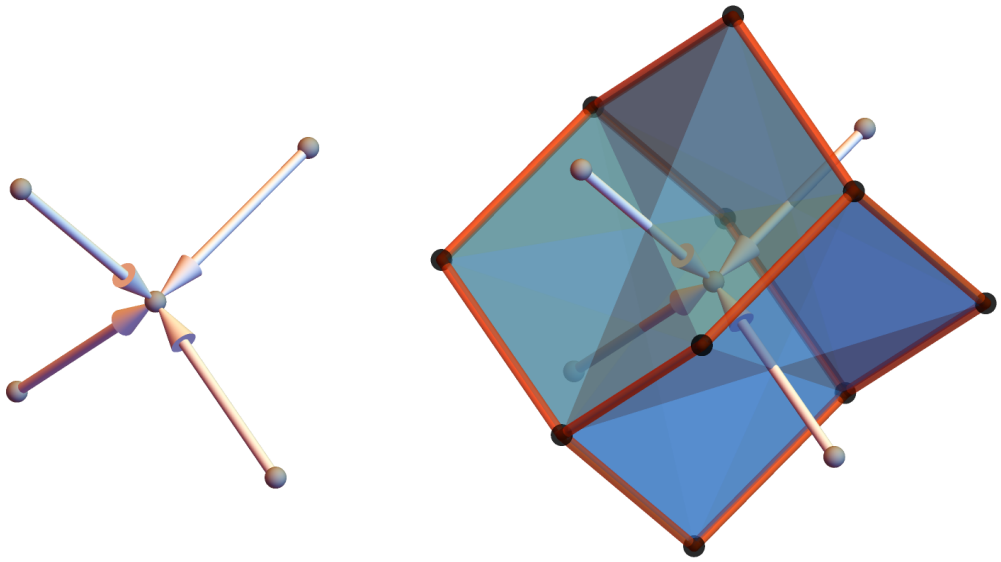
One might be tempted to describe this phase as a "charge soup",[3] where the local divergence fluctuates randomly. In this section we will demonstrate that in fact there is a better interpretation as a "membrane soup", and that the excitations of this phase are 1-dimensional vortex strings. To obtain this description we will need to invoke Hodge duality, the duality between k-forms and $(d-k)$-forms, and make full use of the dual cell complex construction described in Section 3.1.16.

5.3.1 Hodge Duality: Spins as 2-Forms

As we discussed at the beginning of this chapter, purely rotational configurations which form the ground state of the nearest-neighbor spin ice model can be expressed as linear combinations of circulations around minimal plaquettes, i.e. we can express Eq. (5.1) as

$$\left| S^{\text{rot}} \right\rangle = \sum_p c_p \left| \partial p \right\rangle . \tag{5.12}$$

[3] We will not use term "monopole" here, which we reserve to describe the excitations in spin ice. If we were to be extremely precise, it would be appropriate to reserve the term "monopole" only to describe the (in principle) physically observable, localized and isolated excitations in spin ice, which are renormalized by hard spin constraint and interaction with the bath of fluctuating rotational modes. The word "charge" is only meant to evoke the notion of the local divergence defined in the Hodge theory.



(a) (b)

Figure 5.3: (a) A "star" of the diamond lattice cell complex, the arrows represent the 1-form $\mathrm{d}_0 |v\rangle$, where each edge carries unit magnitude. (b) Every vertex (0-cell) sits at the center of a volume (3-cell) of the dual cell complex, and each edge (1-cell) passes through the center of plaquette (2-cell) of the dual complex.

In continuum vector calculus this would be equivalent to saying that a rotational vector field can be written locally as the curl of vector potential, and in the language of differential forms the statement that a rotational 1-form can be written as the coexterior derivative of a 2-form. We need not only use minimal plaquettes in the discrete case however—any rotational configuration can be expressed as linear combinations of circulations around the boundary of an *open membrane* m,

$$\left|S^{\mathrm{rot}}\right\rangle = \sum_m c_m \left|\partial m\right\rangle \equiv \sum_m c_m \delta_2 \left|m\right\rangle , \tag{5.13}$$

where m is an oriented surface. $|m\rangle$ can be be written as a linear combination of the plaquettes $|p\rangle$ making up the surface, $|m\rangle = \sum_{p\in m} |p\rangle$. Then in $\delta_2 |m\rangle = |\partial m\rangle = \sum_{p\in m} |\partial p\rangle$, the circulation around the boundary of each interior plaquette p cancels with its neighbors, leaving only the circulation on the boundary.

Let us consider the equivalent statement for irrotational modes. An irrotational configuration can be expressed as the coboundary of vertices,

$$\left|S^{\mathrm{irr}}\right\rangle = \sum_v c_v \left|\partial^\dagger v\right\rangle \equiv \sum_v c_v \mathrm{d}_0 \left|v\right\rangle . \tag{5.14}$$

We will refer to the coboundary of a single vertex as a "star," as depicted in Fig. 5.3(a). In vector calculus this would be the statement that an irrotational vector field can be expressed as the gradient of a scalar function. Then what is the analog of Eq. (5.13) for irrotational modes? Intuitively we would need to "glue" stars together along their edges to form a larger network, in which all of the "interior" edges cancel out. But this clearly does not have a continuum analog and it is not obvious how to interpret the remaining non-cancelling edges.

A more intuitive formulation can be constructed by invoking the dual cell complex, associating each k-cell of the complex with a $(d-k)$-cell of the dual complex. Then the coboundary of the original complex becomes the boundary on the dual complex. In the case of the diamond lattice, its dual is another diamond lattice, illustrated in Fig. 5.3(b) for a single star. The vertices of the original lattice map to 3-cells of the dual complex, and the edges of the original lattice map to plaquettes of the dual lattice. The star of the original complex then becomes a *closed, oriented surface* in the dual complex, the boundary of a dual 3-cell.

Let us therefore switch our perspective and think of our fundamental degrees of freedom not as a 1-form field on the original cell complex, but rather a 2-form field on the dual cell complex. The divergence δ_1 then becomes $\tilde{\mathrm{d}}_2$ on the dual lattice, where we use a tilde to denote the chains, cochains, and linear operators of the dual lattice. We can map back and forth between the two descriptions by orienting each dual 2-cell extrinsically in the same way as the 1-cell passing through it, so that the dual 2-form $\tilde{S}$ has value $\tilde{S}(\tilde{p}) = S(\star^{-1}\tilde{p})$, where $\star$ is a Hodge star operator which maps each oriented e in the primary complex the oriented plaquette $\tilde{p}$ in the dual complex.[4]

Utilizing the Hodge duality, we can now construct the irrotational analogue of Eq. (5.13),

$$\left|S^{\mathrm{irr}}\right\rangle = \sum_v c_v \left|\star^{-1}(\partial(\star v))\right\rangle . \tag{5.15}$$

To read this, $\star v$ is a dual 3-cell (volume) surrounding vertex v (Fig. 5.3(b)), $\partial(\star v)$ is the oriented surface of this 3-cell in the dual lattice, amd $\star^{-1}(\partial(\star v))$ is then the "star" surrouding vertex v (Fig. 5.3(a)). Written in terms of the dual 2-form field $\tilde{S}$, we have

$$\left|\tilde{S}^{\mathrm{irr}}\right\rangle = \sum_{\tilde{\tau}} c_{\star^{-1}\tilde{\tau}} \left|\tilde{\partial}\tilde{\tau}\right\rangle , \tag{5.16}$$

where the sum is over the 3-cells of the dual complex, $\partial\tilde{\tau}$ is the boundary of the 3-cell, and $\star^{-1}\tilde{\tau}$ is the corresponding vertex in the primary complex. That is, if we interpret our spin degrees of

[4] As we used it here, we defined $\star$ to act on chains. We could equivalently define a $\star$ acting on cochains and write $\tilde{S} \equiv \star S$, which would be more akin to the Hodge star appearing in differential geometry and Hodge theory, but we will use it this way to emphasize the relationship between the k-cells of the primary complex the $(d-k)$-cells of the dual complex

freedom as a 2-form on the dual complex, then the irrotational configurations can be written as a linear combination of *closed membranes* which are the boundaries of contiguous *volumes* in the dual lattice.

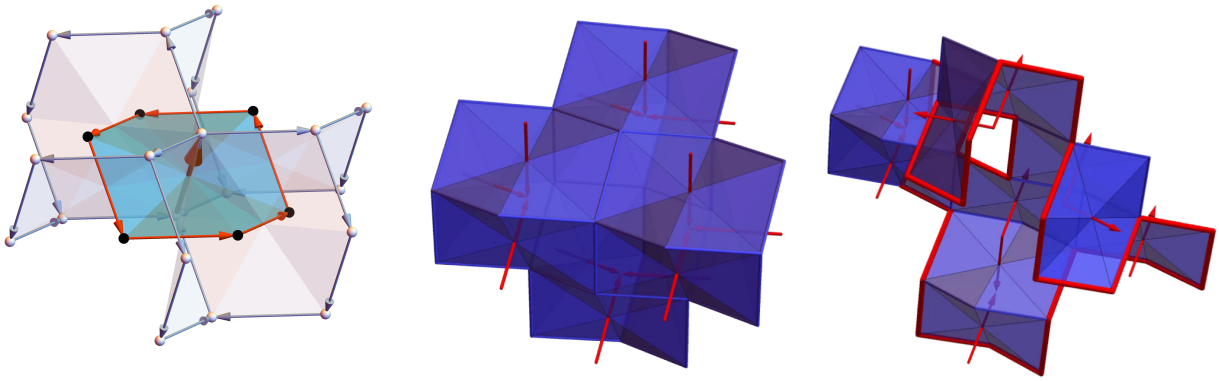
What about the harmonic components? Harmonic 1-forms satisfies the Laplace equation $\Delta_1 \gamma = 0$ and cannot be expressed as being supported on the boundary of an open surface such as Eq. (5.13), since they integrate to non-zero values on a non-trivial 1-cycle which is not a boundary of any surface. Similarly, a harmonic 2-form is an element of $\ker \Delta_2$ and integrates to a non-zero value on non-trivial 2-cycles, surfaces winding around the 2-dimensional holes of the system. Therefore the harmonic portion of the spin configuration can be expressed as a linear combination of closed membranes wrapping around the holes of the complex.

5.3.2 The 2-Coulomb Phase: A Soup of Closed Membranes

With the dual complex and the Hodge duality in hand, we can now characterize the spin liquid ground state of the vorticity model, which we will call the 2-Coulomb phase. In the regular Coulomb phase of spin ice, spins are treated as 1-forms and ground states obey the zero-divergence condition $\delta_1 S = 0$. The ground state subspace is $\ker \delta_1 = \operatorname{im} \delta_2 \oplus H^1$, the space of rotational and harmonic modes, and spin configurations can be decomposed non-uniquely into a superposition of closed strings. In the 2-Coulomb phase, we can also interpret spins as a 1-form field obeying the zero-curl constraint $\mathrm{d}_1 S = 0$, which states that S has zero circulation around every contractible 1-cycle of the lattice, and the ground state subspace is $\operatorname{im} \mathrm{d}_0 \oplus H^1$, the space of irrotational and harmonic modes. Alternatively, we can interpret spins as a 2-form field on the dual complex, obeying the dual constraint $\tilde{\delta}_2 \tilde{S} = 0$, in which case the ground state subspace is $\ker \tilde{\delta}_2 = \operatorname{im} \tilde{\delta}_3 \oplus \tilde{H}^2$, and spin configurations can be decomposed non-uniquely into a *superposition of closed membranes* on the dual complex.

Note that it is quite useful to be able to pass back and forth between treating spins as 1-forms or 2-forms. In the 2-form picture it is clear that ground states are superpositions of closed membranes, which is unclear in the 1-form picture, but the constraint is easier to read in the 1-form picture. In particular the constraint $\mathrm{d}_1 S = 0$ immediately tells us that $\mathrm{d}_1 S(p) = S(\partial p) = 0$, the circulation around any plaquette is zero (and more generally around any 1-cycle). The equivalent statement in the 2-form picture is $\tilde{\delta}_2 \tilde{S}(\tilde{e}) = \tilde{S}(\tilde{\partial}^\dagger \tilde{e}) = 0$, which has a less intuitive (though equivalent) interpretation. In fact the 2-form picture is perfectly applicable to spin ice as well as the vorticity model: the zero-divergence constraint $\delta_1 S$ is equivalent to $\tilde{\mathrm{d}}_2 \tilde{S} = 0$, or $\tilde{S}(\tilde{\partial} \tilde{\tau}) = 0$, which says that the 2-form integrates to zero on closed 2-cycles, i.e. the flux through any closed surface in the dual lattice is zero![5] There is only one example in the literature that we are aware of where such a

[5]This also serves as a very simple proof of an observation originally made by Henley [136] that the lattice can be



(a) (b) (c)

Figure 5.4: (a) The minimal excitations of the vorticity models are vortex tube. The vorticity 2-form $\omega = \mathrm{d}_1 S$ measures the circulation around each plaquette of the direct lattice. Its Hodge dual $\tilde{\omega}$ is a 1-form on the dual lattice. Flipping a single spin (center, red arrow) in the 2-form description creates a flux through a single plaquette of the dual complex (blue hexagon), and $\tilde{\omega} = \tilde{\delta}_2 \tilde{S}$ encodes the corresponding circulation around this dual plaquette (red arrows). (b) A closed membrane consisting of twelve dual plaquettes, where all twelve spins have the same orientation relative to the membrane surface. Flipping these spins will result in a new configuration with the same energy. (c) An open oriented membrane. Flipping this membrane in the ground state results in an closed string excitation localized along the boundary with energy proportional to its length. In the direct lattice this excitation is a vortex tube.

perspective was taken (implicitly), by Ref. [168], who treats spins to be associated to a plaquette of the dual lattice, defines a vector potential on the edges of the dual lattice, and matter fields on the vertices of the dual lattice.

5.3.3 Open Membranes and Vortex Tubes

Let us now consider the excitations. In spin ice, the excitations are described by the 0-form $Q = \delta_1 S$, the divergence of S. The minimal excitation is a single spin flip, which in the Coulomb phase creates a pair of opposite charges at the ends of the corresponding bond, or alternatively flipping an *open string* of head-to-tail spins, with opposite charges localized at the boundary of

partitioned into parallel planes such that the sum of spins in each plane alternates. We simply observe that any such pair of planes is the boundary of a slab in the dual complex, and therefore the flux through the planes must be equal and opposite.

the string. It is never possible to create just a single excitation, they are always created pairwise at the end of a string, and one end of the string must run off to infinity to leave behind a singe deconfined excitation.

In the vorticity model the excitations are plaquettes with non-zero vorticity, but such a plaquette cannot be isolated. If we start from a ground state with zero circulation on every plaquette and flip a single spin, we will create a non-zero circulation on *six* plaquettes containing the edge that spin is on. Specifically, these six plaquettes make up the coboundary of the edge, $\partial^\dagger e$. This is depicted in Fig. 5.4. The structure of these excitations is much clearer in the dual picture, where the vorticity 2-form ω becomes a dual 1-form $\tilde{\omega} = \tilde{\delta}_2 \tilde{S}$. Flipping a spin corresponds to adding to the 2-form spin configuration $-2\,|\tilde{p}\rangle$, and the induced vorticity is the circulation around the boundary of $\tilde{p}$, depicted in Fig. 5.4. We can also flip a larger open membrane where all spins are oriented in the same direction, the resulting excitation is then an oriented closed string localized along the boundary of the membrane along the edges of the dual lattice. In the original lattice this string pierces a sequence of hexagons, each of which has non-zero vorticity, so the excitation is a *vortex tube*, analogous to a vortex ring in a liquid. Since the Hamiltonian is $(J_\omega/2)\sum_p \omega(p)^2$, the energy of this excitation is proportional to its length, with every hexagon with non-zero ω contributing.

To summarize, the 2-Coulomb phase is most clearly understood by treating spins as a 2-form field in the dual lattice. Ground state configurations decompose non-uniquely into a superposition of closed membranes, and thus the ground state manifold can be described as a "membrane soup". Excitations are created by flipping spins forming an oriented open membrane, and the physical excitation is a vortex tube localized along the boundary of this surface, whose energy is proportional to its length. This contrasts to the standard Coulomb phase (or "1-Coulomb phase") where spins can be treated as a 1-form field, the ground state manifold is a "soup of closed strings", and excitations are created by flipping oriented open strings, with energy proportional the size of their boundary, which is independent of the length of the string. The vortex tubes of the 2-Coulomb phase are thus confined, and can not be made arbitrarily large for finite energy cost.

5.4 Thermodynamics with Hard Spins

We now turn to studying the hard spin Ising vorticity model. In the SCGA the ground state of the vorticity phase is clearly complementary to the Coulomb phase, and the correlators are trivially inverted, swapping irrotational and rotational modes. In spin ice and the traditional Coulomb phase, the SCGA provides a very good description of the hard spin model, with extremely small corrections captured by a $1/N$ correction. In contrast, we will find that the physics of the hard spin

vorticity model is remarkably modified, and in fact exhibits a highly non-trivial phase transition at finite temperature, despite having a spin liquid ground state.

5.4.1 Monte Carlo Cluster Updates: Membrane Moves

We will utilize Monte Carlo as our primary tool to investigate the hard spin model. As with spin ice, single spin flips are sufficient at high temperature, but the autocorrelation time grows rapidly at low temperature as we approach the highly correlated spin liquid phase, and cluster updates are necessary to maintain ergodicity on a reasonable timescale. The standard method for studying the Coulomb phase of spin ice is the loop move, flipping a closed loop of head-to-tail spins, which only modifies the spin configuration in the rotational sector. [125] In the vorticity model the ground states can be decomposed into a superposition of closed surfaces in the dual lattice, where every spin has the same orientation relative to the surface (i.e. the surface is fully polarized, just as an oriented loop is fully polarized). Flipping all of the spins on such as surface modifies the spin configuration only in the irrotational sector, and thus does not change the energy of the system. Therefore ideally we would like to have a "membrane algorithm" which finds such closed surfaces and flips them. Unfortunately this is remarkably more difficult than the loop algorithm.

The loop algorithm amounts to a constrained random walk on the 1-skeleton of the cell complex. Starting from a single vertex, one makes a random choice of edge to walk along, then continues this process, always choosing an exit edge with the same relative spin orientation as the spin on the incoming edge, so that the walk traces an oriented contour of spins. The walk ends when we either re-intersect the path, in which case we discard the dangling portion of the path before the intersection point, or we continue until we arrive back at the starting point [125, 143].

Attempting to grow a two-dimensional surface by this method is far more complicated. Working with the dual plaquettes, for each of the six edges of the plaquette there are five other plaquettes sharing that edge, each of which is a possible continuation of the surface, and we must choose a random one with the same orientation as the starting plaquette. We would then continue this way along the boundary of the surface constructed thus far, but since the surface is not one dimensional there are many directions to grow in. Finding a closed surface by this method is extremely unlikely, and it is possible to eventually end up in a position where one edge has no valid extensions to grow the surface on and this "membrane growing" algorithm fails. A typical resulting surface is depicted in Fig. 5.4(c). The largest closed surface that we find is shown in Fig. 5.4(b), which consists of twelve plaquettes, but even finding these surfaces is very rare and the vast majority of the time the algorithm fails to find a closed surface.

The ideal cluster algorithm that we would like to have would be the generalization of the one introduced in Ref. [128] for nearest-neighbor spin ice, which works at all temperatures and

achieves the decomposition of spin configurations into a superposition of strings, both open and closed, in a stochastic manner. The basic idea of the method is to exploit the fact that the Hamiltonian splits into a sum over tetrahedra in order to assign a decomposition of the four spins on each tetrahedron into clusters with temperature-dependent weights for each possible cluster. If we could generalize the algorithm to the vorticity model, it would decompose spin configurations into a superposition of open and closed oriented membranes. Unfortunately, despite the fact that the vorticity Hamiltonian splits into a sum over plaquettes, we find that the straightforward generalization of this method fails.[6]

Fortunately, not all is lost. We already know what the minimal membranes are: they consist of four dual plaquettes as shown in Fig. 5.3(b), corresponding in the direct lattice to a star. In spin ice, such a fully-polarized membrane corresponds to an all-in or all-out tetrahedron, i.e. a double charge. Thus finding any vertex which is fully polarized, we can flip the four spins from all-out to all-in or vice versa to obtain another configuration with the same energy. We refer to this as a *star move*. Thus, for our simulations of the vorticity model, we will utilize both single spin flips plus zero-energy star moves.

Note that this means that there will be one primary limitation on the conclusions we can draw from our Monte Carlo simulations: we cannot update the harmonic sector, which would require that we flip a non-trivial surface winding around the periodic boundaries of the system, which would change the bulk magnetization. Since flipping a single spin changes the rotational, irrotational, and harmonic components of the spin configuration (cf. Fig. 3.14), the the algorithm is in principle still ergodic, but single spin flips will be exponentially suppressed at low temperature. A single spin flip excitation from the ground state creates a loop excitation consisting of six hexagons with vorticity $\omega(p) = \pm 2$, and therefore the energy of this minimal excitation is $(J_\omega/2) \cdot 6 \cdot (\pm 2)^2 = 12 J_\omega$, so at low temperatures the probability of acceptance of a single spin flip will be $\exp(-12 J_\omega / T)$. Furthermore, in order to create a harmonic update, we would need to create a loop excitation bounding a plaquette of the dual lattice, and then grow it to the system size so that it wraps around the periodic boundaries and closes again. Therefore, at sufficiently low temperature, and in particular at zero temperature, single spin flips and star moves are not ergodic and will never achieve harmonic updates.

5.4.2 Specific Heat: An Unexpected Phase Transition

The first thermodynamic measurement we make is the specific heat. In spin ice the specific heat exhibits a broad hump without any sharp peak, and the Coulomb phase is continuously connected to the paramagnetic phase without a phase transition between them. We show the

[6] Specific, the equations for the weights of the various "breakups" do not admit solutions.

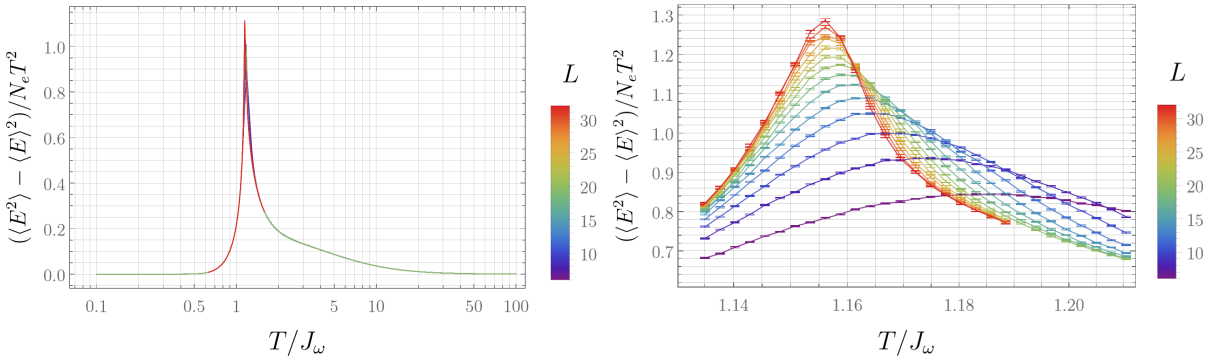


Figure 5.5: Specific heat per spin of the vorticity Ising model as a function of temperature in units of J_ω, measured in Monte Carlo simulations at various systems sizes varying from $L = 6$ to $L = 32$. (a) Broad temperature range with 10^5 sampled spin configurations per temperature, showing a sharp peak which grows with system size, indicating a phase transition. (b) Detailed structure of the peak at various system sizes with 10^7 samples per temperature.

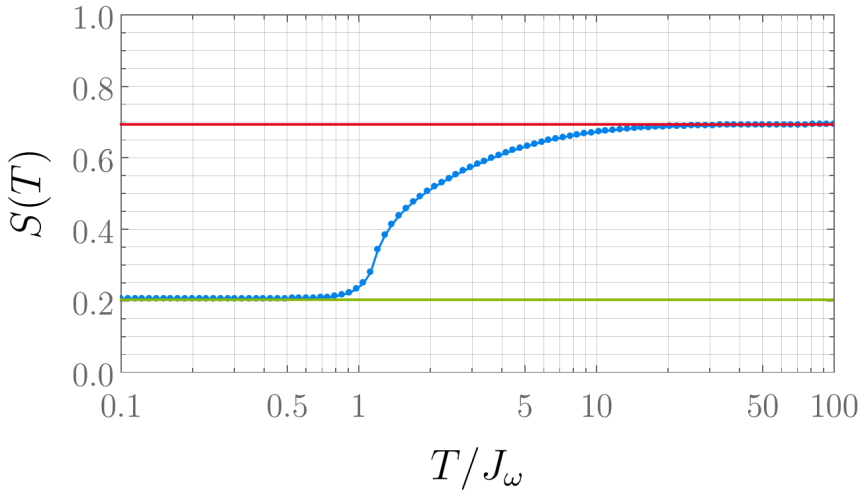


Figure 5.6: Entropy per spin computed from the Monte Carlo specific heat data, using $L = 32$ data for temperatures near the specific heat peak and $L = 20$ data away from the peak. The red line shows $S(\infty) = \ln(2)$. For $T/J_\omega > 100$ the entropy is estimated by fitting the high-temperature tail of the specific heat, yielding a very small contribution of $S(\infty) - S(100J_\omega) \approx (2.7 \pm 0.1) \times 10^{-4}$. The green line shows Pauling's estimate of $(1/2)\ln(3/2) \approx 0.202733$, while we obtain a slightly larger value of $S(0.1J_\omega) \approx 0.2061 \pm 0.0001$.

measured specific heat per spin for the vorticity model in Fig. 5.5(a) over a broad temperature range for system sizes ranging from $L = 6$ to $L = 32$, and immediately notice that there is a very prominent and sharp peak at a temperature $T \sim J_\omega$. In Fig. 5.5(b), we show the details of the peak, which grows with system size, which seems to imply a phase transition at a temperature $T_c/J_\omega \sim 1.155$. Thus clearly the hard spin physics of this model is more complicated than that of NNSI, which exhibits only a smooth crossover from the high-temperature paramagnetic phase to the low-temperature Coulomb phase.

Before we look at correlations to investigate this phase transition, we consider briefly the ground state entropy to check whether we indeed have a spin liquid phase which we spent so much time discussing, despite this apparent transition. In Fig. 5.6 we show the entropy per spin computed from the specific heat data measured in Monte Carlo, using

$$S(T) = \underbrace{S(\infty)}_{\ln(2)} + \underbrace{\int_{\infty}^{100J_\omega} \mathrm{d}T\, \frac{c(T)}{T}}_{-(2.7\pm0.1)\times10^{-4}} + \int_{100J_\omega}^{T} \mathrm{d}T\, \frac{c(T)}{T}, \tag{5.17}$$

where the change in entropy from $T = 100J_\omega$ (the highest temperature measured in Monte Carlo) to infinite temperature is estimated by fitting the high-temperature tail of the specific heat $c(T)$ and yields a negligible contribution. The entropy per spin at $T = 0.1J_\omega$, the lowest temperature measured in Monte Carlo, is found to be 0.2061 ± 0.0001, which is remarkably close to Pauling's estimate of $(1/2)\ln(3/2) \approx 0.202733$ for water ice [169]. The reason for this is not immediately clear, since Pauling's estimate can be computed by a simple counting argument based on corner sharing tetrahedra, for each of which $6/16$ Ising configurations are ground states, while the vorticity model is a model of *edge*-sharing hexagons, for which $20/64$ spin configurations are ground states (6 choose 3, since three spins must be opposite the remaining three to have zero circulation). Standard methods for computing the entropy for spin ice rely on graph theoretic methods [170] which do not work for the vorticity model due to the edge-sharing nature of the plaquettes. A more detailed analysis of the ground state entropy and its closeness to Pauling's estimate is certainly warranted here.

5.4.3 Structure Factors Above the Transition

We now turn to an analysis of the spin-spin correlations, which we will compare to the naive calculations from the SCGA shown in Fig. 5.1. In Fig. 5.7 we show the same structure factors as measured in Monte Carlo for a system of size $L = 16$ averaged over 10^5 samples for three temperatures above the putative critical temperature indicated by the specific heat peak. Note these temperatures are significantly lower than the $12J_\omega$ energy cost of the minimal vorticity excitations,

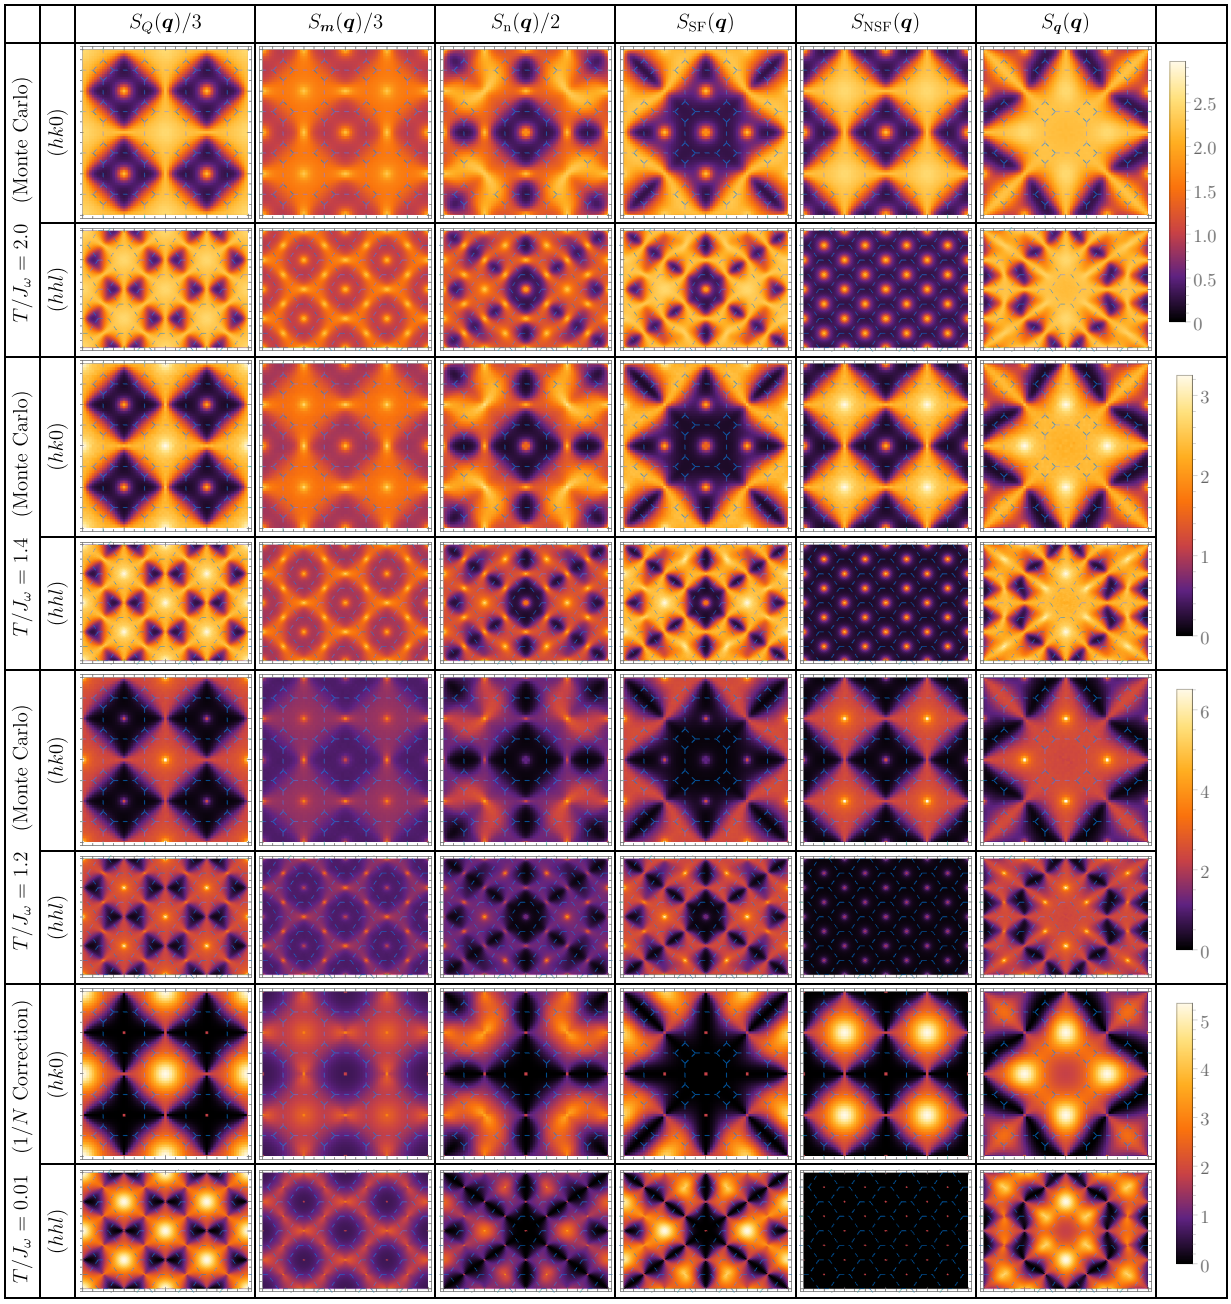


Figure 5.7: Structure factors computed at three temperatures above $T_c/J_\omega \sim 1.15$ in Monte Carlo for $L = 16$ with 10^5 samples, and at $T/J_\omega = 0.01$ in the Large-N with the $1/N$ correction.

meaning that at this temperature the rotational modes should already be strongly suppressed. At $T/J_\omega = 2$ the correlations are qualitatively similar to those of the SCGA seen in Fig. 5.1. At $T/J_\omega = 1.4$ we can observe some spectral weight has shifted towards zone centers, for example at (000) in $S_Q(\boldsymbol{q})$, (220) in the (hhl) SF channel, (220) in the $(hk0)$ NSF channel, and perhaps most prominently at (200), (020), and (002) in the longitudinal channel. These regions of higher intensity are Lorentzian peaks forming above the phase transition. At $T/J_\omega = 1.2$, close to the transition, the peak intensity increases rapidly as these Lorentzians begin to sharpen into Bragg peaks.

At the bottom of Fig. 5.7 we also show the order $1/N$ corrected Large-N structure factors. In constrast to NNSI, the $1/N$ correction exhibits a significant increase in intensity near the same positions as observed in Monte Carlo, though with much broader widths. It compares fairly well with the intermediate temperature $T/J_\omega = 1.4$ structure factors, before these peaks sharpen nearer to the transition. Thus we can at least conclude that the deviation of the Monte Carlo results from the SCGA is likely driven by the spin length constraint, and is at least to some extent qualitatively reproduced in a $1/N$ expansion. It would interesting to see if the $1/N^2$ corrections can produce an improved agreement. The $1/N$ correction to the correlation function appears as a radiative correction (see Appendix C for details), and such radiative corrections are famously known to drive symmetry breaking in field theories, called the Coleman-Weinberg mechanism [171].

5.4.4 AIAO Order by Disorder: Symmetry Breaking Coexisting with a Spin Liquid

In Fig. 5.8 we again show the same structure factors but now just below T_c at $T/J_\omega = 1.0$. We can see clear sharp Bragg peaks at zone centers indicative of long-range order. However, we can also observe intensity in the background behind the peaks (note the color scale has been adjusted to highlight the lower-intensity background). This background intensity is qualitatively the same as the SCGA correlations and the intensity above T_c, indicative of a spin liquid *coexisting* with long range order, which is consistent with our finding of a non-zero ground state entropy. But what is the nature of the long range order indicated by these peaks? On the bottom of Fig. 5.8 we show the same structure factors computed for the AIAO ordered state, divided by L^3. They appear in the same locations and with the *same relative intensities* as those observed in the hard spin vorticity model below T_c. This very strongly suggests that the symmetry breaking that we observe in the vorticity model corresponds to AIAO order.

We define the AIAO order parameter as

$$\phi = \frac{1}{4N_\nu} \sum_\nu (-1)^\nu \langle \nu | \, \delta_1 \, | S \rangle \,, \tag{5.18}$$

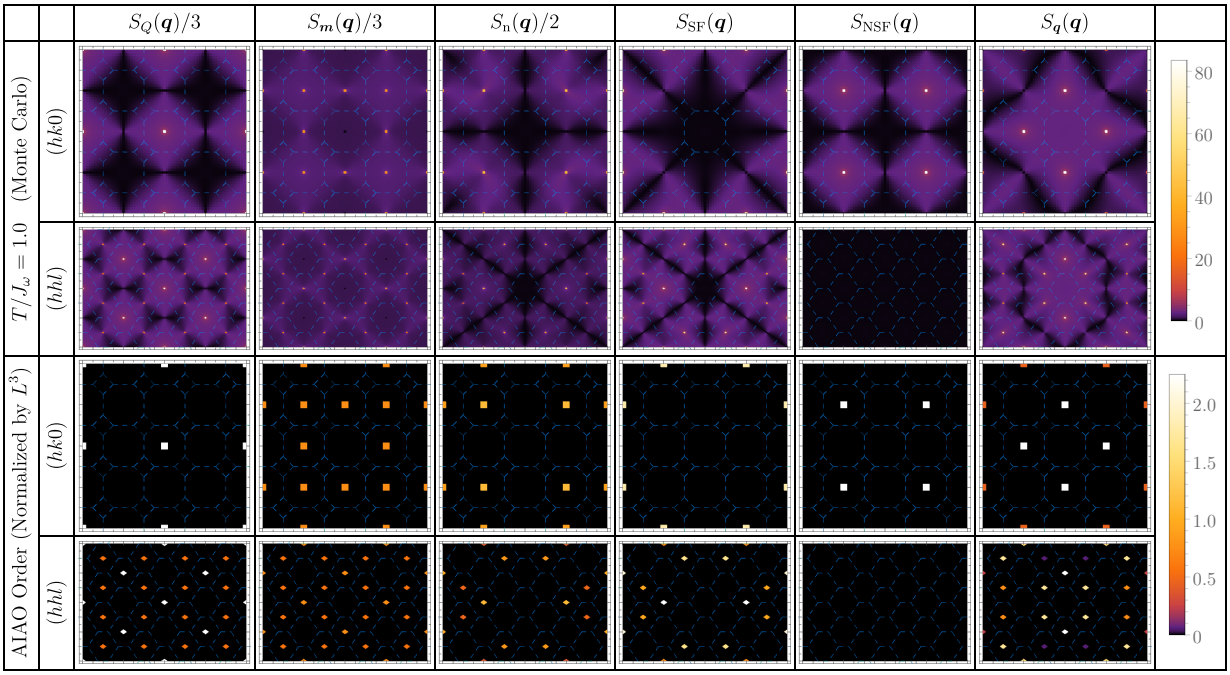
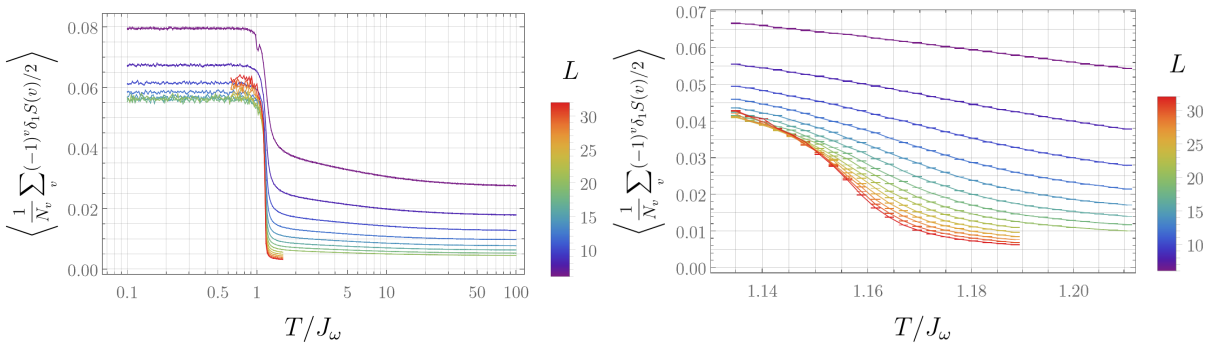


Figure 5.8: Top: Structure factors for the vorticity model at $T/J_\omega = 1.0$, just below the transition at $T_c/J_\omega \sim 1.15$, for $L = 16$ with 10^5 samples. Bottom: the corresponding structure factors for the fully ordered AIAO state, with intensity divided by L^3. Bragg peaks for the vorticity model below T_c appear at the positions and with the relative intensities expected for AIAO order, but their intensity contains only a fraction of the total spectral weight, and we can still see the spin liquid correlations in the background. Note that we have scaled the color bar intensity logarithmically to highlight the low-intensity background in the top panels.



(a) (b)

Figure 5.9: All-in-all-out order parameter as a function of temperature for system sizes varying from $L = 6$ to 32. (a) Broad temperature range with 10^5 sampled spin configurations per temperature, showing that the order parameter saturates , indicating a phase transition. (b) Detailed structure of the peak at various system sizes with 10^7 samples per temperature.

where we take all vertices given the same orientation and $(-1)^v$ has opposite sign on the A and B diamond sublattices (alternatively, orient all vertices on ones sublattice "in" and all vertices on the other sublattice "out"). This is $\pm$ for the all-in or all-out configurations. If we take the standard convention that edges are oriented from diamond sublattice A to B then we can write the order parameter as

$$\phi = \frac{1}{N_e}\sum_e S(e) \quad (e : A \to B), \tag{5.19}$$

which is +1 if all spins are oriented $A \to B$ and -1 if all spins are oriented $B \to A$. In Fig. 5.9 we show the absolute value of the AIAO order parameter measured for system sizes varying from $L = 6$ to 32 as a function of temperature. In Fig. 5.9(a), we can see that it does indeed go to a non-zero value at low temperatures, but we notice two things: first, it does not go to +1 at low temperature, and for larger systems saturates around 0.06, indicating that the system is only partially ordered. We also note that the low-temperature value is not always consistent across different runs. This is likely due to ergodicity breaking: the AIAO configuration is in the zero magnetization sector, but the system may get stuck in a sector of non-zero magnetization when cooling, and thus the order parameter may saturate to different values depending on the run. In Fig. 5.9(b), we show the behavior of the order parameter near the transition, where it appears to approach zero above T_c in the thermodynamic limit. The results are indistinguishable whether we warm the system (starting from the AIAO configuration) or if we cool from the paramagnetic phase.

The existence of all-in-all-out symmetry breaking is quite surprising in this system, and the reason for it is not at all obvious. The AIAO state is the ground state of the pyrochlore Ising *ferromagnet*,[7] but the physics of the vorticity AIAO transition is not clearly connected to the 3D Ising ferromagnet transition. Consider an Ising ferromagnet in its ground state, where every spin is parallel and every nearest-neighbor pairwise interaction $-S_1 S_2$ is minimized. Flipping a single spin costs energy for each bond which is no longer satisfied, and flipping a cluster of neighboring spins creates a defect all along the boundary of the flipped cluster. Starting from one symmetry broken ground state we must nucleate a defect (a "bubble") and grow it to be as large as the system in order to flip the system from all-up to all-down. But growing the defect costs energy proportional to the area of its boundary, which grows with the size of the defect. Note that the 3D Ising ferromagnet can therefore be described as a 3-form field, where spins are associated to the d-dimensional volumes of a cell complex,[8] then the excitations are described by a 2-form $\delta_2 S$. The Ising ferromagnet transition can thus be understood as the vanishing of a surface tension.

This is very different from the vorticity model description in terms of a 2-form field with 1-dimensional excitations, though in both cases the excitation energies scales with their size. It seems plausible that the vorticity phase transition is also linked to the confinement of these 1-dimensional excitations, and perhaps the transition can be understood as the vanishing of a string tension [172]. But it seems that despite the phase transition having the same order parameter as the 3D Ising ferromagnet, their universality classes are not guaranteed to be the same, and indeed on the low temperature side of the vorticity transition we still have a spin liquid with extensive entropy.

Even more importantly, unlike the ferromagnet, the two AIAO symmetry-broken ground states are connected by a sequence of *local zero energy* fluctuations. Specifically, if we start from the all-in configuration and apply a star move on every star centered on the A (or B) sublattice of the diamond lattice, then we arrive at the all-out configuration, and at every step in between the system is in a ground state. This is remarkably different from the 3D ferromagnetic transition! But it is somewhat reminiscent of the phenomenon of order by disorder [173–175], which is usually studied in the context of a systems with a continuously degenerate ground state. In the standard order by disorder story, at zero temperature the system has a continuous manifold of ground states to explore which are energetically degenerate, but at finite temperature the system can maximize its entropy by selecting a subset of sates in the ground state manifold about which there are more low-energy states available.

[7] By this we mean the Ising model on the pyrochlore lattice with $H_{\mathrm{NNFM}} = -J \sum_{\langle ij \rangle} S_i^z S_j^z$, with quantization axes taken with the standard $A \to B$ orientation and $J > 0$ so that the ground states are the two AIAO ordered configuration.

[8] Of course we can also treat it as a 0-form field by Hodge duality, with spins located at vertices, but the excitations are better understood if we treat the spins as a top-dimensional form.

Indeed we can observe that the AIAO states are the *maximally flippable* states for star moves, since *every* star can be flipped. But as soon as we have flipped one star, the four surrounding vertices are no longer fully polarized and can no longer be flipped. Thus, while there is a specific sequence of zero energy star moves we can apply to get from one symmetry broken ground state to the other, applying a star move to every star in a single diamond sublattice, this "bridge" between the two AIAO configurations becomes obstructed if we start flipping stars in the other sublattice. Thus we can heuristically say that there is a *subextensive bottleneck* in the ground state manifold connecting the two AIAO ground states, and while in principle it is possible to get from one to the other by applying only local zero energy updates to the system, it is practically forbidden in a thermodynamic system by the intermediate states in the "bottleneck" forming a subextensive set of states contributing zero entropy, such that the probability of the system ever transitioning from one AIAO ground state to the other goes to zero in the thermodynamic limit. This is only a heuristic conjecture, at present we don't have a precise definition of this bottleneck or how to measure its size, though we have some evidence to support it.

As a partial validation for the above conjecture, in Fig. 5.10 we show the structure factors again but this time computed at zero temperature by starting from one symmetry broken configuration ("all-in") and applying *only* star moves at random to flippable stars. On the left we recover the same Bragg peaks that were observed below T_c in Fig. 5.8 corresponding to the AIAO symmetry breaking, indicating that indeed only applying star moves is unable to take us from one AIAO ground state to the other. On the right we show the intensity beneath the Bragg peaks, which is qualitatively the same as the correlations of the 2-Coulomb phase obtained from the SCGA, including the longitudinal pinch points indicative of quasi-long-range correlations, but with some spectral weight shifted into what appear to be Lorentzians underneath the Bragg peaks. Despite the symmetry breaking the majority of the spectral weight remains in the spin liquid correlations, with the Bragg peaks containing less than 10% of the total intensity.

5.5 Interpolation to the 3D Ising Ferromagnet: Phase Diagram

To investigate the phase transition further, we will utilize the fact that the vorticity model is adiabatically connected to the nearest-neighbor ferromagnet, for which the all-in/all-out phase transition falls clearly in the 3D Ising universality class, because the vorticity Hamiltonian contains ferromagnetic nearest-neighbor interactions. The nearest-neighbor ferromagnet Hamiltonian is proportional to $-\mathrm{d}_0\delta_1$, and we will parameterize the path such that the interpolated Hamiltonian can be written

$$H(x) = \frac{J}{2}\langle S, [x(8 - \mathrm{d}_0\delta_1) + (1-x)\delta_2\mathrm{d}_1]S\rangle, \tag{5.20}$$

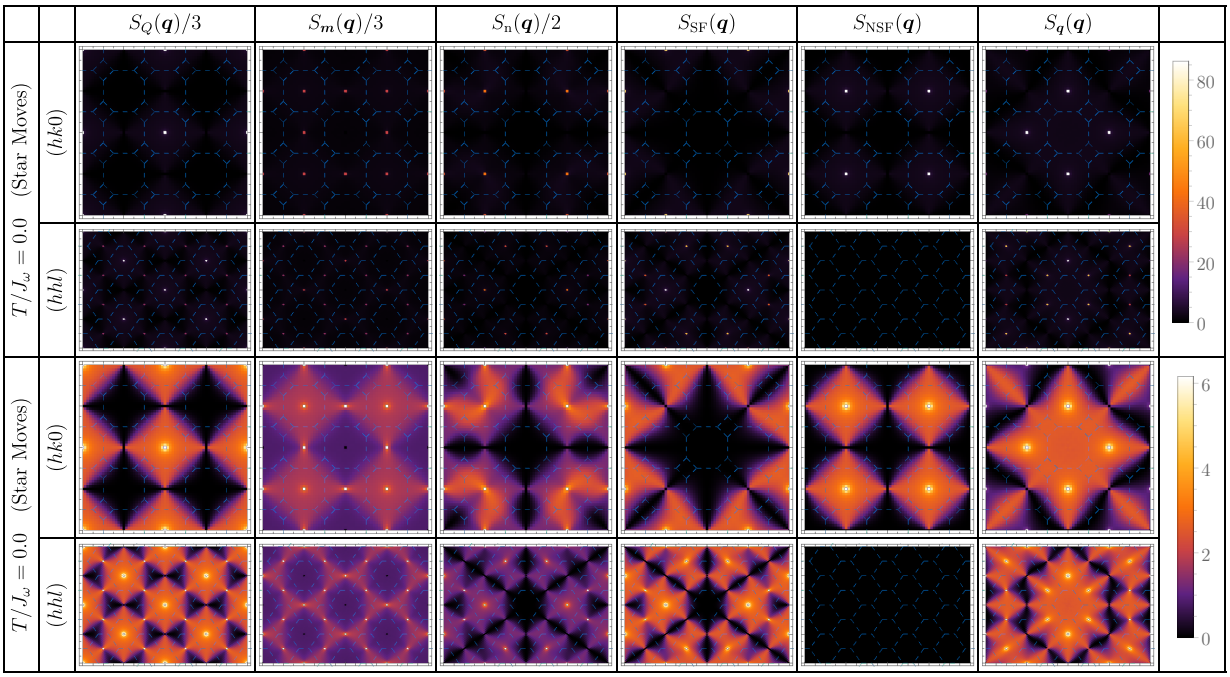


Figure 5.10: Monte Carlo structure factors at $T = 0$ obtained by starting from one AIAO configuration and only applying star moves randomly to flippable stars. Top: We observe the same Bragg peaks as seen in Fig. 5.8, which serves as a partial validation for the theory that the symmetry breaking is due to a subextensive bottleneck in the ground state manifold. Bottom: Adjusting the color scale to highlight the low-intensity background, beneath the Bragg peaks we observe the correlations of the 2-Coulomb phase, which contains the large majority of the spectral weight. Notice that the pinch points intensity zone centers is zero, because we are always in the zero-magnetization sector and do not have harmonic mode fluctuations.

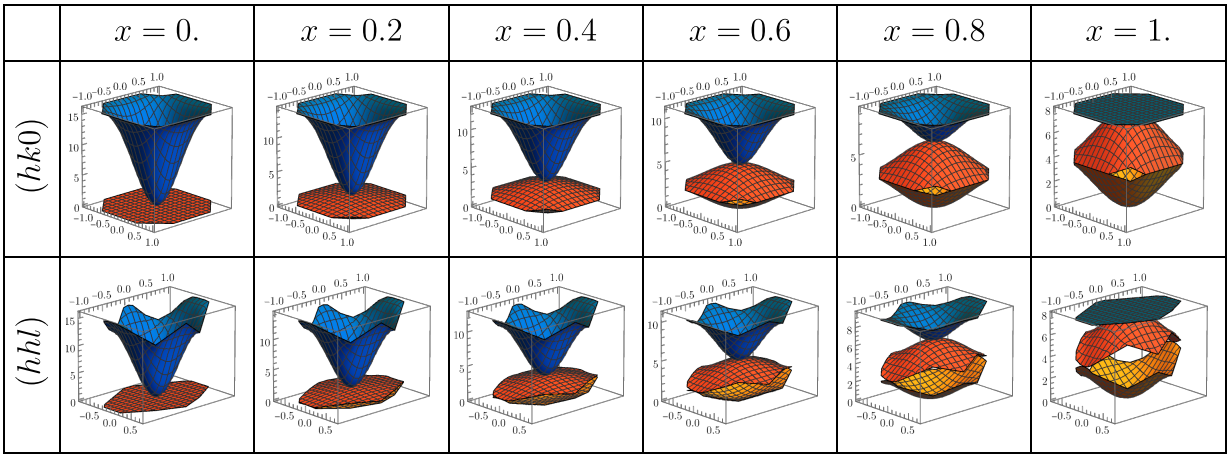


Figure 5.11: The band structure of the interpolated interaction matrix when tuning from the vorticity model at the left end to the nearest-neighbor Ising ferromagnet at the right end. Top row: (hhl), bottom row: $(hk0)$. Due to the choice of Hamiltonian the rotational (green) and irrotational (red,orange) bands remain distinct, since we tune the interaction in the sectors separately. At the left end there are two flat irrotational bands and two dispersive rotational bands, while at the right end there are two flat rotational bands and two dispersive irrotational bands Note that the degeneracy of the threefold band touching point is preserved throughout, and furthermore the two vorticity bands remain degenerate throughout.

where J is an arbitrary energy scale, so that when $x = 0$ we have the vorticity model and when $x = 1$ we have the nearest-neighbor ferromagnet. Note that the constant shift in $8 - \mathrm{d}_0\delta_1$ ensures that the minimum eigenvalue is zero since $\mathrm{d}_0\delta_1$ has bandwidth of 8 and is positive definite, so the ground state has zero energy independent of x. In terms of adjacency matrices, the interaction matrix can be written

$$\mathcal{J}_{ij}(x)/(J/2) = x(6\mathbb{1} - \mathcal{A}^{(1)}) + (1-x)(6\mathbb{1} - 2\mathcal{A}^{(1)} + \mathcal{A}^{(2)} - \mathcal{A}^{(3b)}) \tag{5.21}$$

$$= 6\mathbb{1} - (2-x)\mathcal{A}^{(1)} + (1-x)(\mathcal{A}^{(2)} - \mathcal{A}^{(3b)}). \tag{5.22}$$

We will use single spin flips supplemented by star moves again, but for $x \neq 1$ the star moves will have an energy cost to account for.

5.5.1 Interpolated Bands

Since the two terms in our Hamiltonian Eq. (5.20) act in the rotational and irrotational sector separately, they do not mix the rotational and irrotational bands in the spectrum. In Fig. 5.11 we show the interpolated spectrum, where at the left end we start with two irrational flat bands (red, orange) and two rotational dispersive bands (green), and at the right end of the spectrum we have two flat rotational bands and two dispersive irrotational bands. As soon as $x > 0$, the AIAO configurations, corresponding to the bottom of the red band, have minimum energy, and the system should in principle have a *true* AIAO phase transition, where the ground state is non-degenerate and every vertex has maximum divergence.

5.5.2 Star Moves

When performing a star move (a minimal closed-membrane move) the circulation around any plaquette is unchanged, but the divergence $\delta_1 S$ at the four sites surrounding the star changes, and the divergence at the center of the star flips sign. Let $S^{(1)}$ denote the original configuration and $S^{(2)}$ denote the updated configuration with four spins flipped from all-out to all-in or vice versa. Both of these configurations have the same vorticity, so there is no energy change from the rotational sector. On the other hand, the change of the divergence on the four vertices surrounding the star changes the energy in the irrotational sector,

$$\langle S^{(2)}, (-\mathrm{d}_0\delta_1)S^{(2)}\rangle - \langle S^{(1)}, (-\mathrm{d}_0\delta_1)S^{(1)}\rangle = -\sum_v \left[|\delta_1 S^{(2)}(v)|^2 - |\delta_1 S^{(1)}(v)|^2\right] \tag{5.23}$$

Therefore the change in energy is

$$\Delta E_{\text{star}}/J = -\frac{x}{2}\sum_t \left[\left(\sum_{e\in\partial^\dagger v} S^{(2)}(e)\right)^2 - \left(\sum_{e\in\partial^\dagger v} S^{(1)}(e)\right)^2\right] \tag{5.24}$$

Let $Q = \delta_1 S^{(1)}$. For each of the four vertices one spin flips, and correspondingly $Q(v) \to Q(v) \pm 2$, with a + sign if the spin was originally oriented "out" from that vertex.[9] Note that $\delta_1 S(v) \sim -(\operatorname{div} S)_v$ is *positive* if the net "flow" is *inward* towards v. Therefore, flipping an "out" spin *increases* $Q(v)$, because it increases the inward "flow".

Let v^* denote the central vertex. If the spins surrounding the central vertex are all "in", such that $Q(v^*) > 0$, then they are all "out" of the neighboring for four vertices. Thus we can conclude that for the four vertices surrounding v^*, $Q(v) \to Q(v) + 2\operatorname{sgn}(Q(v^*))$. The change in energy of flipping a star is therefore

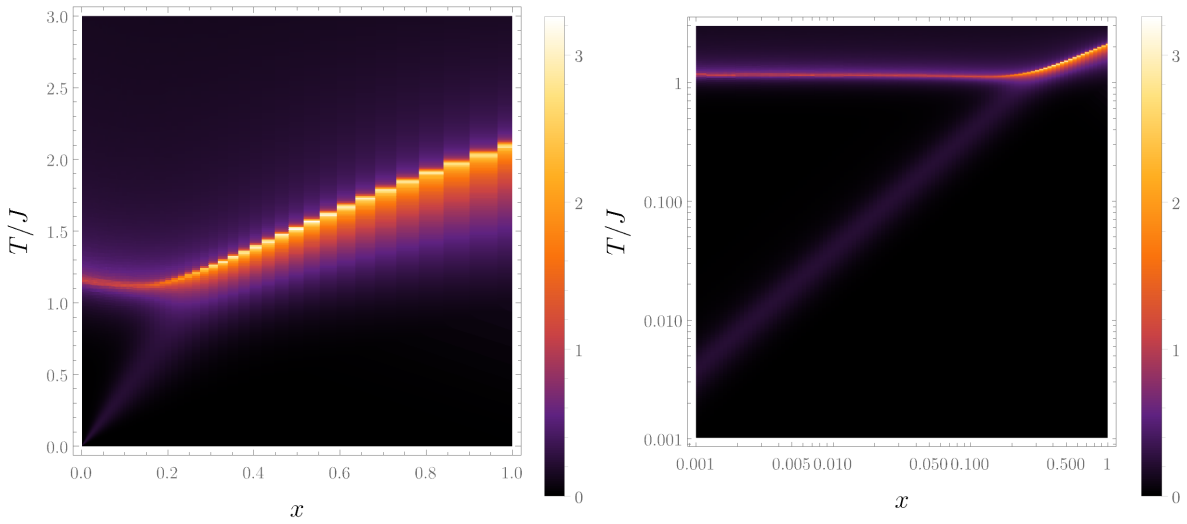
$$\begin{aligned}\Delta E_{\text{star}} &= -\frac{x}{2}\sum_{v@v^*}\left[(Q_v + 2\operatorname{sgn}(Q(v^*)))^2 - Q_v^2\right] \\ &= -\frac{x}{2}\sum_{v@v^*}\left[4 + 4\operatorname{sgn}(Q(v^*))Q_v\right] \\ &= -\frac{x}{2}\left(16 + 4\operatorname{sgn}(Q(v^*))\sum_{v@v^*} Q_v\right).\end{aligned}$$

Using this, we can continue to utilize star moves for $x \neq 0$, which are accepted according to the Metropolis condition with probability $\min(1, \exp(-\beta\Delta E_{\text{star}}))$.

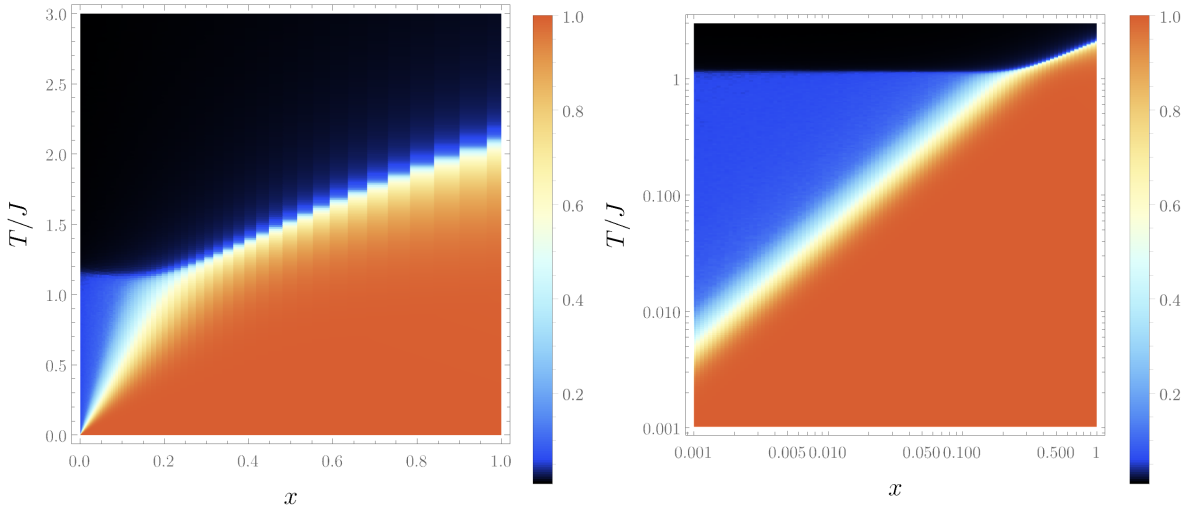
5.5.3 Phase Diagram

In Fig. 5.12(a) we show the the specific heat, and in Fig. 5.12(b) the AIAO order parameter, as a function of T/J and x for a system size $L = 16$ with 10^5 sampled configurations per temperature and x (left half linear scales and right half logarithmic scales). At $x = 0$ we have the vorticity model, and at $x = 1$ the AIAO ferromagnet. In the specific heat we can clearly see the line of phase transitions separating the high temperature paramagetic phase from the low temperature phase. In the AIAO order parameter the fully ordered AIAO symmetry broken is visible in orange, while the partially-ordered 2-Coulomb phase is visible as the blue region extending from $x = 0$.

[9] Note that in our Monte Carlo algorithm we fix a reference orientation for every edge, then store the values of $S(e_o) = \pm 1$ with respect to the chosen orientation. We take the standard $A \to B$ orientation, so that an AIAO configuration with spins pointing from sublattice A to sublattice B has $S(e_o) = +1$ for every edge.

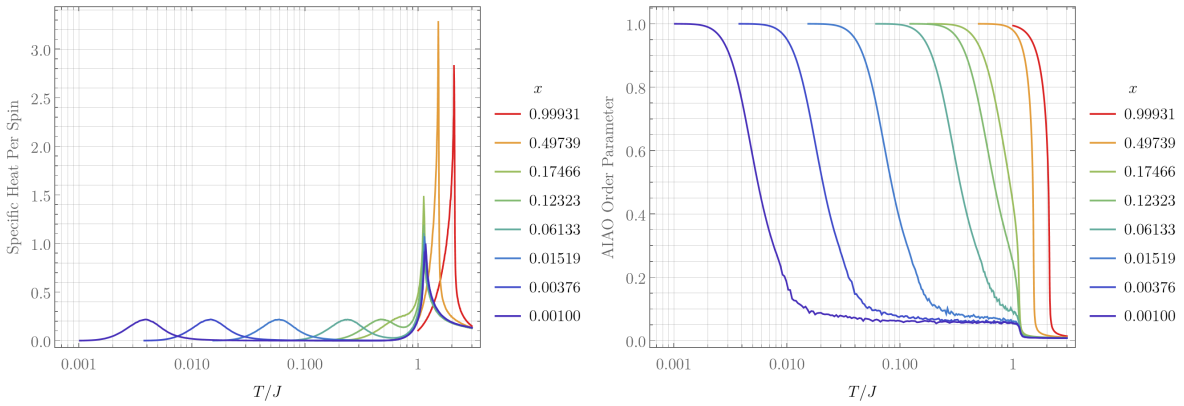


(a) Specific Heat



(b) AIAO Order Parameter

Figure 5.12: (a) Specific heat for the Hamiltonian Eq. (5.20) as a function of temperature T/J and the interpolation parameter x, measured in Monte Carlo for $L = 16$ with 10^5 sampled configurations per temperature and x. At the left end ($x = 0$) is the vorticity model and at the right end ($x = 1$) is the AIAO ferromagnet. Left: linear-linear scale. Right: log-log scale. (b) The AIAO order parameter as a function of T/J and x. The partially ordered 2-Coulomb phase is visible as the light blue region.



(a) (b)

Figure 5.13: Cuts at various x of the specific heat and AIAO order parameter with $L = 16$ and 10^5 samples per temperature.

Our initial observation is that the partial AIAO symmetry breaking extends to finite x at finite temperature. It is tempting to say that the blue region in Fig. 5.12(b) corresponds to the 2-Coulomb phase, but the AIAO order parameter only accounts for the weak partial symmetry breaking, which is not necessarily itself an intrinsic feature of the 2-Coulomb phase, which does not have an order parameter.[10] We can, however, say that in the blue region the vorticity string excitations will be heavily suppressed and should be very sparse in the system. In the orange region the 2-Coulomb phase has been suppressed and its macroscopic entropy removed. For all $x > 0$ the system eventually reaches the fully ordered AIAO ground state via a *smooth crossover*, which is visible as the blurred line in the specific heat and the white region in the AIAO order parameter.

More importantly, we see that the smooth crossover extends to higher and higher temperatures as x increases, until it eventually merges with the line of phase transitions around $x \sim 0.16$. We also see that this rising of the crossover temperature coincides with a slight *decrease* in the critical temperature, and after they merge the critical temperature rises. From the AIAO order parameter it seems clear that after this line of transitions merges with the critical temperature the phase transition is adiabatically connected to the $x = 1$ ferromagnetic transition. It appears plausible that a qualitative change in the transition may occur near $x^* \sim 0.16$.

[10] Similarly we can say that the Coulomb phase of NNSI is only sharply defined at zero temperature. A more useful characterization is perhaps to say that the spin liquid phase is well-formed once the density of defects is sufficiently small. In this case, we can estimate that the density of defects is already quite small even above T_c, judging by the width of the pinch points in Fig. 5.7.

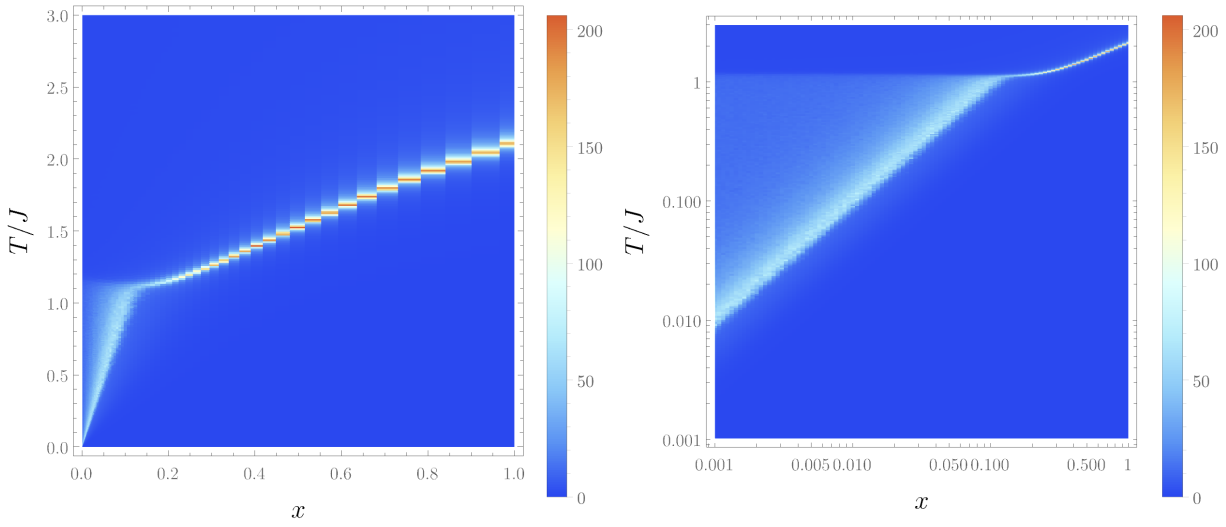
In Fig. 5.13 we show cuts of the specific heat and AIAO for a few selected values of x below and above x^*. In the specific heat we see the broad bump where the 2-Coulomb phase crosses over to AIAO order for $x < x^*$. Correspondingly in the AIAO order parameter we see the phase transition at $T_c/J \approx 1.16$ where the order parameter takes a value of around 0.06, followed by a continuous rise starting roughly around $T/J \sim 10x$ to the fully saturated ground state value as temperature decreases. At $x \sim x^*$ we can see that the bump merges with the peak at T_c, which slightly draws down the critical temperature where the peak attains its maximum, and correspondingly in the AIAO order parameter the plateau disappears. For $x > x^*$ the critical temperature rises, and the strength of the specific heat increases as well. We note that the specific heat intensity seems to be stronger for x somewhere between x^* and 1, though the reason for this is unclear.

The data we have seen so far seems to suggest that the phase transition to partial AIAO order in the vorticity model is *not* adiabatically connected to the 3D Ising AIAO phase transition of the Ising ferromagnet at $x = 1$. Something non-trivial appears to happen in the vicinity of x^*. This is even more clear if we look at the variance of the AIAO order parameter, $\langle \phi^2 \rangle - \langle \phi \rangle^2$, which is shown in Fig. 5.14. In (a) we show the variance as a function of T/J and x, and in (b) we show cuts for a few values of x above and below x^*. Here we see an extreme distinction between the behavior above and below x^*. For $x \gg x^*$, near the ferromagnetic model, there is a very strong peak in the variance, indicating a divergence in the thermodynamic limit as is expected of the susceptibility of an order parameter at a second order phase transition. For $x < x^*$ however we see no such divergence, in fact we see only a step at T_c with no peak whatsoever, followed by a gradual but noisy rise and what appears to be an asymmetric hump or possibly a cusp at the crossover to AIAO order. Thus we can definitively say that the AIAO phase transition is qualitatively distinct from that of the ferromagnetic model. Furthermore, the crossover to the fully ordered phase may also be a phase transition if the AIAO variance indeed has a cusp.

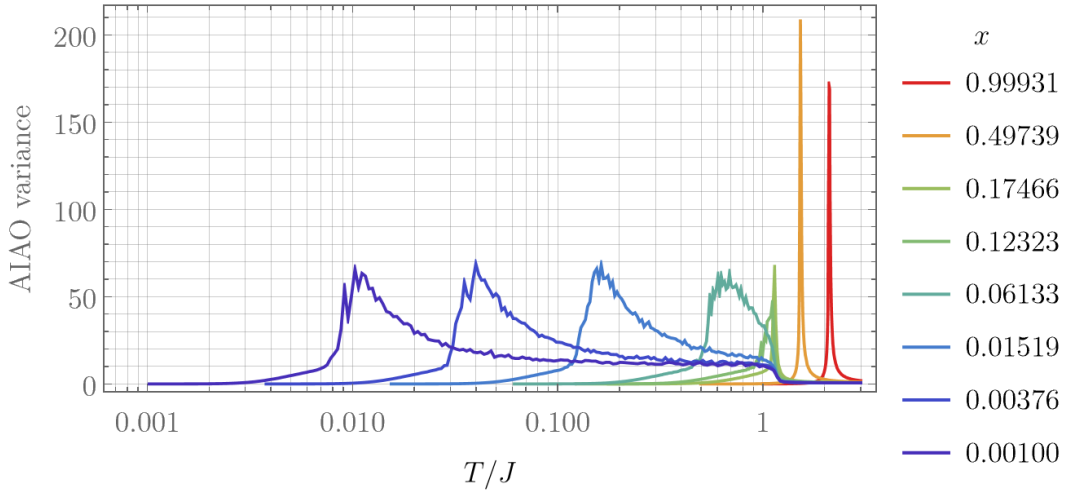
5.5.4 Finite Size Scaling

Unfortunately we do not have a full characterization of the phase transition in the vorticity model. We conclude with a collection of results looking at the finite size scaling of various quantities at the phase transition which give further evidence for the distinction between the $x = 0$ and $x = 1$ phases transitions. The results are summarized in Fig. 5.15, where we show the temperature dependence of the three quantities discussed in the previous section in the vicinity of T_c at various values of x.

In (a) we show the AIAO order parameter, and in (b) we show the specific heat, which can be compared to the $L = 16$ results of Fig. 5.13. Note in particular the specific heat, where we can see an enhancement in the peak intensity near $x = 0.6$ that we noted previously, and this appears to



(a)



(b)

Figure 5.14: The variance of the AIAO order parameter as a function of temperature and x for $L = 16$ with 10^5 samples per data point. For $x > x^*$ this quantity exhibits a strong peak, indicating diverging fluctuations of the AIAO order parameter in the thermodynamic limit. For $x < x^*$ however, it exhibits no divergence whatsoever.

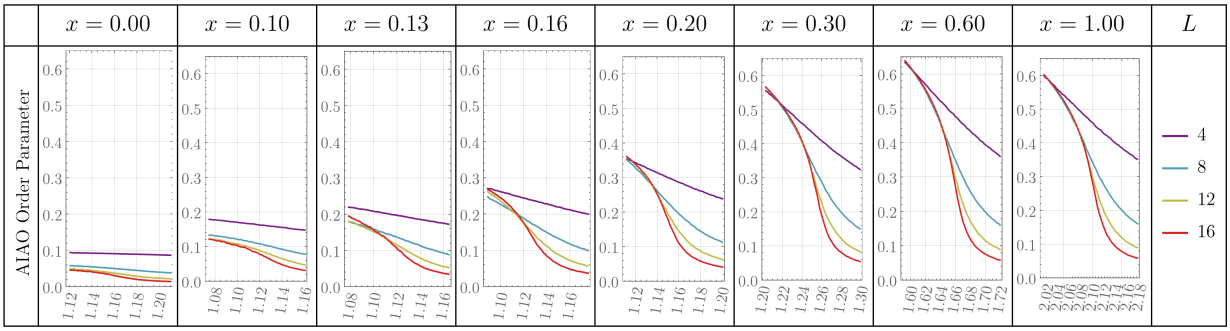
coincide with an enhancement of the AIAO order parameter as well. Note the difference in scaling behavior between large-x and small x of the specific heat. At large x the peak is very prominent and grows rapidly with system size, while for small x the peak is very much diminished. It is not clear if the specific heat actually exhibits a true divergence for $x < x^*$ in the thermodynamic limit. Next, in (c) we show the AIAO variance, which can be compared to the $L = 16$ results of Fig. 5.14. It shows a very strong peak at large x, and again we see the peak intensity is enhanced around $x = 0.6$. For $x < x^*$ the peak quickly disappears and gives way to a more gradual rise in intensity, which in Fig. 5.14 appears to become a step.

Lastly, in Fig. 5.16 we plot the fourth order Binder cumulant [176], defined as

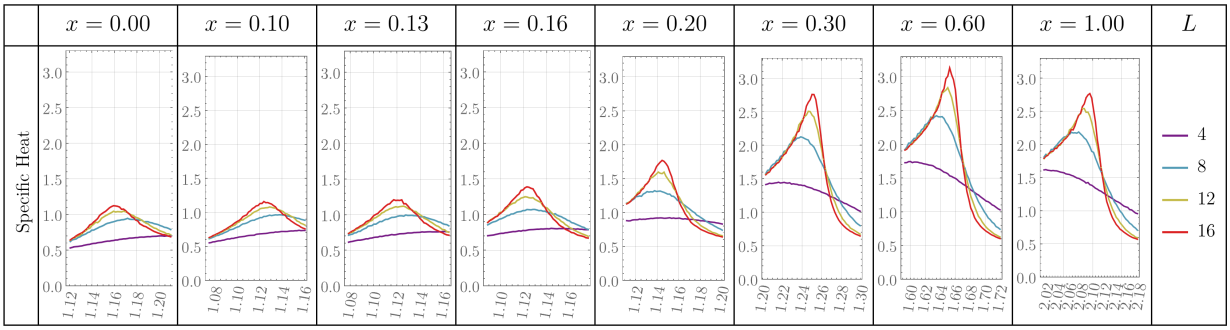
$$\kappa = 1 - \frac{\langle \phi^4 \rangle}{3\langle \phi^2 \rangle^2}, \tag{5.25}$$

which is a scaled version of the kurtosis of ϕ. The Binder cumulants for different system sizes should cross at the critical temperature in the thermodynamic limit. In (a) we plot the Binder cumulant versus temperature, and for large x we can see that they cross near a single temperature, just above $T/J = 2.10$ for $x = 1$, even for the smallest system size of $L = 4$. This makes sense, since this is a $\boldsymbol{q} = \boldsymbol{0}$ order where every unit cell has the same configuration driven by only nearest-neighbor interactions even a small system with just a few unit cells exhibits signatures of the transition. But this becomes obviously not the case when x gets smaller, where the $L = 4$ system rapidly falls away from the other three curves. Below x^* the crossings of the curves for the system sizes shown becomes worse, until at $x = 0$ the upper three curves do not cross near the same temperature. This is indicative that there is a large lengthscale emerging for $x < x^*$ which is larger than the system size, so that for smaller x and fixed L the system becomes further and further from the thermodynamic limit.

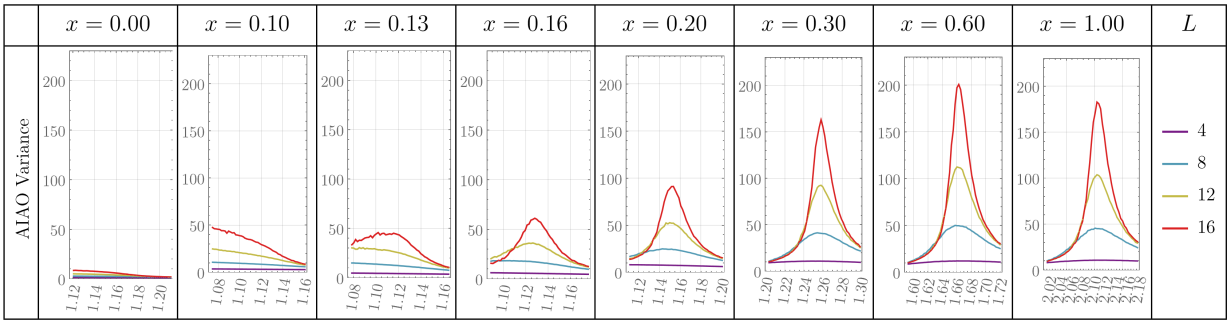
In Fig. 5.16(b) we plot the Binder cumulant vs $L^{1/\nu}(T - T_c)/T_c$ using the 3D Ising critical exponent $\nu = 0.63$, where we make a rough estimate of T_c from the location of the specific heat peak. At $x = 1$ we see a very nice collapse of the curves onto their universal scaling function [177], and we get good collapse for smaller values of x as well. For intermediate x such as $x = 0.2$-0.26 the collapse is not as good, likely due to a slight error in the critical temperature, but the curves appear clearly parallel. Below $x^* \sim 0.16$ we see the $L = 4$ curve rapidly fall away from the others, and we also note that the remaining curves are no longer parallel over a significant range and fail to collapse. This is further evidence that this phase transition is not the 3D Ising transition.



(a)

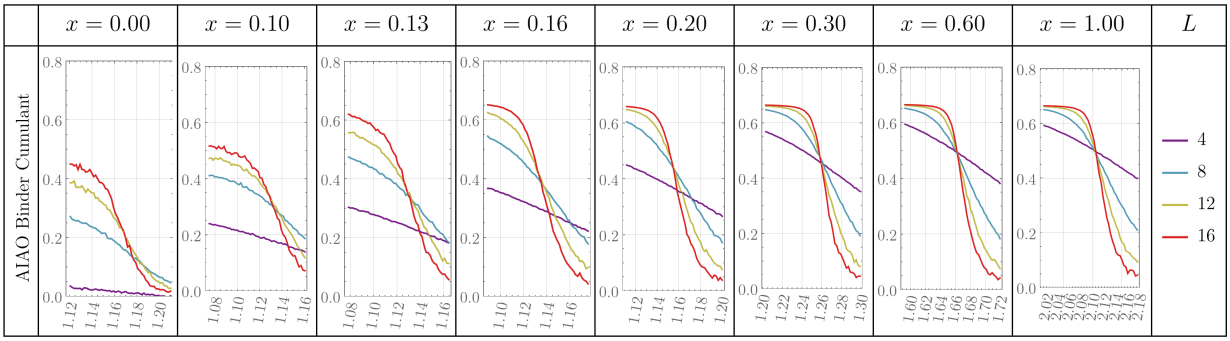


(b)

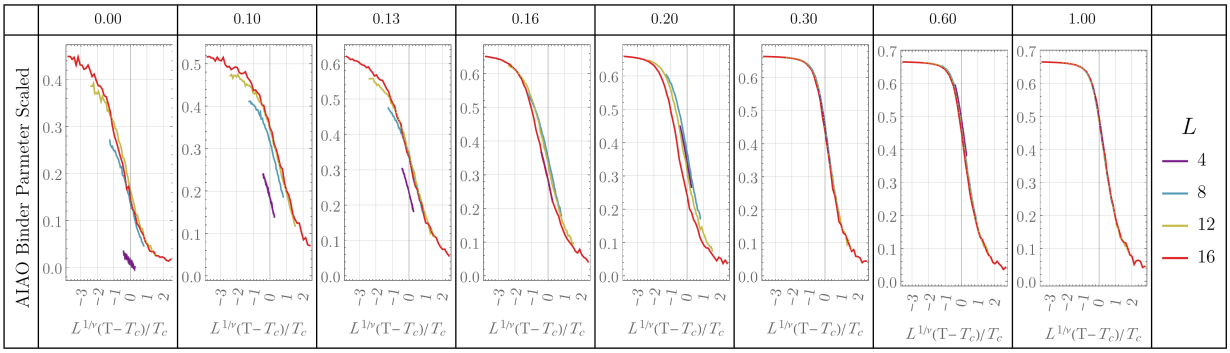


(c)

Figure 5.15: System size dependence of (a) the AIAO order parameter, (b) the specific heat, (c) the AIAO variance.



(a)



(b)

Figure 5.16: (a) The AIAO Binder cumulant defined in Eq. (5.25) at various system sizes and for various values of x. At large x the Binder cumulants for different system sizes cross near a single temperature, even when $L = 4$. (b) The Binder cumulant plotted vs $L^{1/\nu}(T - T_c)/T_c$, where T_c/J is estimated from the specific heat peak and $\nu = 0.63$ corresponding to the critical exponents of the 3D Ising universality class.

5.6 Conclusion and Outlook

Much remains to be done in the study of the vorticity model and the 2-Coulomb phase and the strange phase transition we have investigated here. The best possibility for studying this model in great numerical detail would be to find a way to construct the analog of the cluster algorithm for nearest-neighbor spin ice discussed earlier. That would allow not only a vast speedup in computing time, but would allow for a detailed study of the closed and open membranes at all temperatures and of all sizes, allowing for an parallel analysis to the loop analysis for spin ice done by [143]. In particular, it would allow harmonic sector updates. We cannot rule out the possibility that the weak AIAO symmetry breaking at $x = 0$ could be driven by some loss of ergodicity due to lacking harmonic updates, though, on the basis of the evidence currently at hand, this seems somewhat unlikely.

Also very interesting would be to construct a quantum vorticity model. Taking the ring exchange model of quantum spin ice [25] as a guide, we could immediately write down an analogous model for the vorticity model, which would be *membrane exchange*. In quantum spin ice the S^z operators are treated as the electric field 1-form, and the ring exchange term flips the six spins around a loop, giving rise to a 1-form gauge field and emergent QED. In the quantum vorticity model, the S^z operators should be treated as a 2-form field, and a membrane exchange term would look like $\prod_{e\in\partial^\dagger v} S^+(e)$ + h.c., which flips four spins on an all-in or all-out membrane of Fig. 5.3(b), i.e. performs a star move! As such we would obtain a U(1) *2-form* gauge theory, rather than the familiar U(1) 1-form gauge theory of QED. Interestingly, such higher-form gauge theories have recently generated significant interest in condensed matter [178–184].

It would also be worthwhile to construct the analog of the gauge mean field theory (GMFT) construction [130] for the vorticity model. In the usual GMFT for quantum spin ice, one introduces spinons through an auxiliary Hilbert space representing bosons located on the vertices and couples it to the 1-form S^z operators through a Gauss' law constraint $Q = \delta_1 S^z$, where now Q is treated as a physically distinct matter field. In a sense, the GMFT construction builds in the dumbbell-type fractionalization of spins into a pair of spinons by hand, representing the irrotational portion of the spin configuration through its coupling to the spinons and representing the rotational part of the field through a gauge potential. The analogous construction for the vorticity model would be to treat the S^z as 2-forms on the dual complex and add the vorticity 1-form $\tilde{\omega}$ on the dual edges "by hand", coupled through the 2-form Gauss' law $\tilde{\delta}_2 \tilde{S}^z = \tilde{\omega}$. This can be seen as a different way of fractionalizing spins as "current loops" rather than pairs of "charges," analogous to the two corresponding representations of a dipole in electromagnetism. The excitations of this 2-form gauge theory would then by charged strings rather than charged point particles.[11]

[11] Consider the example of a complex-valued 0-form matter field which has whose action is invariant under global

$U(1)$ phase rotation. A local excitation (particle) of the 0-form field is a small region where the field takes a coherent non-zero value, and thus carries a phase. The phase is called the charge of the particle. A 1-form gauge field is described by a 1-form vector potential A which encodes infinitesimal phase rotations along 1-dimensional displacements. As the particle (a 0-dimensional object) moves along a 1-dimensional path, its phase is rotated by its interaction with the gauge field. The physical degrees of freedom of the gauge field are encoded in the field strength 2-form $F = \mathrm{d}A$, which is a generalized version of geometric curvature. The phase rotation of the particle then reflects the curvature of the field. A 2-form gauge field then is described by a 2-form vector potential, a 3-form curvature, and couples to 1-dimensional charged strings.

Chapter 6

Coulombic Physics In Symmetry-Broken Long Range Ordered Phases

Throughout the earlier chapters of this thesis we focused on Ising models, which have a single spin component and can be modeled by a simple real or complex-valued k-form. In this chapter we will take up the case of $N = 3$ component Heisenberg spins on the pyrochlore lattice. In the case of perfectly isotropic interactions of Heisenberg spins, in the SCGA the spin components decouple and we simply get three copies of the $N = 1$ description in terms of 1-forms, which Ref. [76] refer to as "colors", in the sense of color charge in non-Abelian gauge theories. In the presence of anisotropy, however, these three copies become mixed in non-trivial ways, and it is not *a priori* obvious how to generalize the Hodge formalism developed and utilized thus far to multi-component spins. We will not present a full account of how to do so, but we will use much of the intuition gained thus far to understand the physics of Heisenberg spin systems. It seems clear that such a microscopic description in terms of discrete forms should exist, because as we will see the physics of these systems is closely related, and in particular they exhibit generalized Coulomb phases characterized by multi-fold pinch points and described in the long wavelength approximation by higher-rank gauge theories with matrix-valued fields [120, 185].[1]

Our primary interest in this chapter is not these generalized or higher-rank Coulomb phases, however. Our interest will be in what happens to the spin liquid physics when the system orders into a symmetry broken phase. This problem is important to understand because spin liquids often require some level of fine tuning to maintain stability. In the case of spin ice, for example, we saw that dipole interactions weakly lift the degeneracy of the flat bands and should eventually cause

[1]Note: higher rank gauge theories are usually described by symmetric matrix valued 1-form fields, they are distinct from the higher-form gauge theories discussed in the previous chapter.

the system to order into a symmetry broken phase where the entropy of the Coulomb phase is entirely removed. In general we might expect that if the lifting of the degeneracy is very weak, then quantum fluctuations could prevent the system from ordering and the spin liquid could survive to zero temperature, but this still leaves a relatively narrow window of parameter space where spin liquids can exist to low enough temperatures to be observable experimentally. In most materials near a spin liquid phase we can expect that the spin liquid likely will not survive all the way to zero temperature, and at some temperature there will be a phase transition to long range order. A significant amount of work has gone into characterizing the *transitions* between a spin liquid phase and a conventionally ordered phase, because some of these transitions can be understood as being driven by an emergent Higgs mechanism [168, 186, 187]. Some works have also studied the behavior of such systems on the paramagnetic side of the transition before long range order has set in [187–189]. Our question is not about the transition itself, however, but whether any of the strongly correlated Coulombic physics can survive in the ordered phase that co-opts a spin liquid ground state at low temperature, and in particular in the fluctuations around the ground state.

There is some precedent for this idea in the pyrochlore literature already in Nd compounds [133, 190, 191]. In particular the dipolar-octupolar compound $Nd_2Zr_2O_7$ has a finite temperature phase transition, below which it exhibits a coexistence of Bragg peaks and pinch points in the energy-integrated neutron scattering intensity. The explanation for this phenomenon came from Benton [191], who showed via a spin wave expansion about the ordered ground state that the equations of motion for spin fluctuations can be decoupled via a Helmholtz-Hodge decomposition, and the rotational and irrotational components do not mix at the non-interacting magnon level, which has been dubbed "dynamical fragmentation" (as opposed to the related fragmentation scenario in spin ice iridates where there is a partial condensation of the irrotational modes while the rotational modes remain free to fluctuate [133, 134]). This can be seen as a material three-dimensional realization of the decoupling demonstrated for a two-dimensional toy model in Ref. [192]. $Nd_2Zr_2O_7$ has a dipolar-octupolar doublet [158], and its pseudo-spin Hamiltonian therefore can be decoupled into an XYZ form. As a consequence, the spin wave bands decouple in a fairly trivial way into rotational and irrotational bands, following directly from the corresponding adjacency matrix bands. While this clearly deserves a detailed analysis using the Hodge theory we have developed thus far, we will not provide it here. In this chapter we will instead consider the richer physics arising in fully anisotropic Hamiltonians, which do not so trivially decouple.

6.1 Pinch Point Features in an Ordered Phase

The most general nearest-neighbor Hamiltonian allowed by symmetry on the pyrochlore lattice for Heisenberg spins can be written in the following form [36, 193]

$$H = \frac{1}{2}\sum_{\langle ij\rangle}\Big[\underbrace{J\boldsymbol{S}_i\cdot\boldsymbol{S}_j}_{\text{Heisenberg}} + \underbrace{KS_i^{(ij)}S_j^{(ij)}}_{\substack{\text{Bond-dependent}\\ \text{Kitaev-like}}} + \underbrace{J_{\text{PD}}\boldsymbol{S}_i\cdot P_{ij}\cdot\boldsymbol{S}_j}_{\text{Pseudo-dipolar}} + \underbrace{J_{\text{DM}}\boldsymbol{D}_{ij}\cdot(\boldsymbol{S}_i\times\boldsymbol{S}_j)}_{\text{Dzyaloshinkii-Moriya}}\Big] + \frac{D}{2}\sum_i \underbrace{(\boldsymbol{S}_i\cdot\hat{\boldsymbol{z}}_i)^2}_{\substack{\text{Single ion}\\ \text{anisotropy}}}, \tag{6.1}$$

where $S_i^{(ij)}$ is the Cartesian component of the spin orthogonal to the cubic plane in which pyrochlore bond $\langle ij\rangle$ lies,[2] and $P_{ij} = \hat{\boldsymbol{r}}_{ij}\hat{\boldsymbol{r}}_{ij}^T$ projects out the spin components orthogonal to $\boldsymbol{r}_j - \boldsymbol{r}_i$, which arises when truncating dipole interaction as nearest-neighbor. The Dzyaloshinskii-Moriya (DM) vectors form a spiral when viewed along the [111] axis [194]. This Hamiltonian can be expressed in the compact form

$$H = \frac{1}{2}\sum_{ij}\sum_{\alpha\beta} S_i^\alpha \mathcal{J}_{ij}^{\alpha\beta} S_j^\beta, \tag{6.2}$$

and the interaction matrix $\mathcal{J}$ has *twelve* bands, instead of the four bands we have seen for the Ising models so far.

Spin liquids on the pyrochlore lattice with Heisenberg spins can generally be understood as descending from the fully isotropic Heisenberg antiferromagnet. When all of the anisotropy terms are absent this Hamiltonian decouples into three copies of the nearest-neighbor $N = 1$ model and has a massive ground state degeneracy when $J > 0$, with $\mathcal{J}$ having *six* flat bands and six dispersive bands. Adding anisotropy terms will in general lift this degeneracy, but it is possible that some combinations of anisotropy only lift some of the flat (rotational) bands and leave others behind, resulting in an anisotropic classical spin liquid phase described by a higher-rank field theory [185], which we will delve into more later in this chapter. Such models require fine tuning in general, however, and in general the bands would have a subextensive minimum and the system will have a phase transition at a finite temperature.

The model that we will use to investigate such a co-opted spin liquid is to start from the isotropic Heisenberg anti-ferromagnet and add a small single ion anisotropy $D < 0$, which will force the system into AIAO long range order. In Fig. 6.1 we show the unpolarized neutron scattering structure factor computed in Monte Carlo at three temperatures. At high temperature the system is paramagnetic with weak antiferromagnetic correlations. At a lower temperature

[2] See Fig. 3.13 or Fig. 4.2(a) where a tetrahedron of the pyrochlore lattice is shown inscribed in a cube, with each edge of the tetrahedron (bond of the pyrochlore) lying in one of the cubic faces, with opposite bonds in parallel faces.

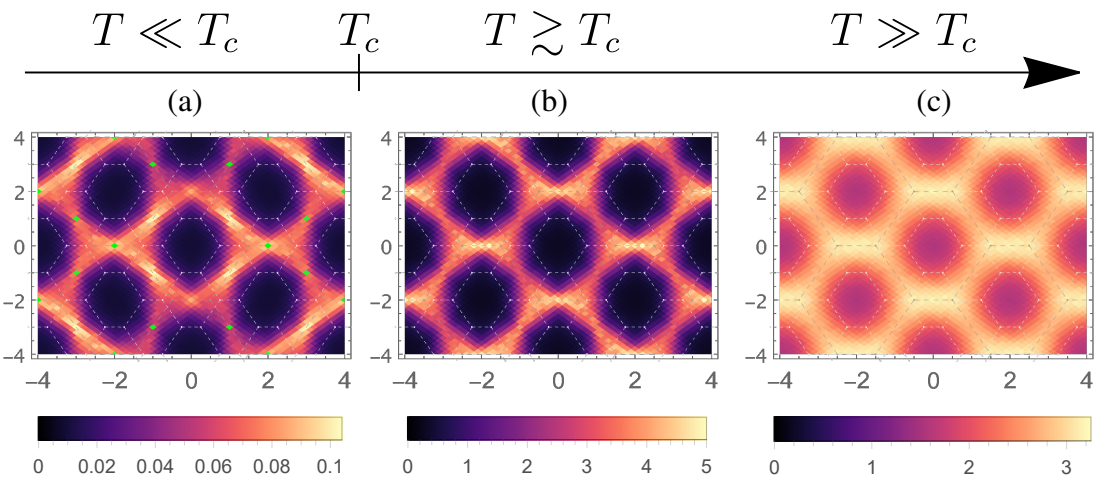


Figure 6.1: Unpolarized neutron scattering cross section in the (hhl) plane measured in Monte Carlo for the easy axis model. At high temperatures $T \gg T_c$ the system is paramagnetic with weak antiferromagnetic correlations. Just above the critical temperature $T \gtrsim T_c$ the system has begun forming pinch points indicative of the correlations of a Coulomb phase, but this spin liquid is co-opted by a phase transition to long range order. Deep in the ordered phase $T \ll T_c$ sharp Bragg peaks have formed (green points), but remnant pinch-point-like features remain in the correlations of fluctuations about the ordered configuration (note the intensity scale).

just above the transition to long range order, the structure factor begins to develop the pinch points associated to the zero-divergence constraint of the Coulomb phase. This Coulomb phase is co-opted by a phase transition to long range AIAO order, with the extensive entropy of the spin liquid removed and the paramagnetic correlations replaced by sharp Bragg peaks (green dots). But we see in the *background*, beneath the Bragg peaks, broad correlations that look similar to those of the Coulomb phase above T_c (note the intensity scale). Our primary goal will be to understand what these correlations are and whether these pinch-point like features are truly indicative of strong Coulombic correlations.

6.2 3-Component Spins on the Pyrochlore Lattice

Before we do, however, we will take the time to discuss in more detail the Heisenberg antiferromagnet, its Coulomb phase, and the rank-2 spin liquids that derive from it and how they can be understood by the long-wavelength expansion prescription laid out in Chapter 2 utilizing the irreps

of the point group T_d. We will use these tools to understand the physics in the ordered phase.

6.2.1 The Heisenberg Antiferromagnet Coulomb Phase

The pyrochlore nearest-neighbor Heisenberg AFM Hamiltonian is

$$H_{\text{Heis}} = \frac{J}{2} \sum_{ij} \mathcal{A}_{ij} \boldsymbol{S}_i \cdot \boldsymbol{S}_j, \tag{6.3}$$

where $\boldsymbol{S}_i$ is a classical three-component vectors of fixed-length S at the pyrochlore lattice site i, $\langle ij \rangle$ denotes nearest-neighbor pairs of the pyrochlore lattice and $J > 0$ is the exchange interaction. Note that this Hamiltonian decouples into three copies of the $N = 1$ Hamiltonian,

$$H_{\text{Heis}} = \sum_{\alpha} \left(\frac{J}{2} \sum_{ij} \mathcal{A}_{ij} S_i^{\alpha} S_j^{\alpha} \right). \tag{6.4}$$

We will not be treating H_{Heis} in this decoupled form, however, and instead keep the spins as 3-component vectors.[3] The Hamiltonian can be written as a sum over tetrahedra,

$$H_{\text{Heis}} = E_0 + 2J \sum_t |\boldsymbol{m}_t|^2, \tag{6.5}$$

where t labels tetrahedra, $E_0 = -2JN_tS^2$ is the classical ground state energy and N_t is the number of tetrahedra, and

$$\boldsymbol{m}_t \equiv \frac{1}{2} \sum_{i \in t} \boldsymbol{S}_i, \tag{6.6}$$

which we will refer to as the net moment of tetrahedron t. Expressed in the form of Eq. (6.5), it is clear that ground states must satisfy the zero-moment constraint $\boldsymbol{m}_t = 0$ on every tetrahedron. This yields six constraints per tetrahedron on the twelve spin components, leaving the system with an extensive ground state degeneracy. The zero temperature phase is a Coulomb phase like that of spin ice, but now the zero-divergence constraint reads $\boldsymbol{m}_t = 0$ instead of $\sum_i S_i = 0$. The distinction is small but important: in spin ice the ground states are *ferromagnetic*, i.e. the 2-in/2-out ground states are the maximally polarized states and have non-zero net moment on each tetrahedron,

[3] In the language of differential forms, we are treating our spins as vector-valued 1-forms. We will not be using the language of forms much in this chapter, however, because we do not have a fully developed theory to treat the anisotropy terms in the Hamiltonian, which likely will require the introduction of a generalized "covariant" exterior derivative/differential d.

but in terms of the Ising variables it is an Ising *anti*ferromagnet. In contrast, the Heisenberg antiferromagnet Coulomb phase has *zero moment* on each tetrahedron.

The long-wavelength expansion discussed for spin ice in Section 2.5.2 was originally developed by Conlon and Chalker for the Heisenberg model [76, 81]. They identified the net moment vector $\boldsymbol{m}_t$ with a 3-color charge Q_c on each tetrahedron, where c indexes colors, and combined the remaining nine degrees of freedom into three 3-component vector fluxes $\boldsymbol{B}_c$. These are just three copies of the corresponding long-wavelength modes for spin ice. The long-wavelength expansion of the Hamiltonian is then just three copies of the long-wavelength expansion for spin ice,

$$\beta H_{\text{Heis, eff}} = \sum_c \frac{1}{2} \int \mathrm{d}^3\boldsymbol{r}\, (\kappa |\boldsymbol{B}_c|^2 + \xi^2 |\nabla \cdot \boldsymbol{B}_c|^2). \tag{6.7}$$

In the limit of zero temperature, the lengthscale ξ diverges and each flux obeys the constraint $\nabla \cdot \boldsymbol{B}_c = 0$. This constraint results in dipolar correlations of the $\boldsymbol{B}_c$ fields, producing the sharp pinch point pattern in the spin-spin correlation functions [55]. At non-zero temperature, ξ is finite and $\nabla \cdot \boldsymbol{B}_c \neq 0$, i.e. a finite density of divergence-full field appears, whose source is the local violation of the constraint $Q_c = 0$, which results in a broadening of the pinch points.

6.2.2 Irreducible Representations Analysis

In this section, we review how deviations from the fully isotropic model can be classified by the irreducible representations of the tetrahedral point group T_d. The most general 3-component spin model on the pyrochlore lattice with nearest-neighbor bilinear interactions can be written as

$$H = \frac{1}{2} \sum_t \sum_{ij \in t} \sum_{\alpha,\beta} S_i^\alpha \mathcal{J}_{ij}^{\alpha\beta} S_j^\beta, \tag{6.8}$$

where t labels tetrahedra (vertices of the diamond cell complex), α, β label spin components, and the 12×12 interactions matrix $\mathcal{J}$ is identical on each tetrahedron. The twelve eigenvectors of $\mathcal{J}$ are indexed by irreducible representations (irreps) of T_d, and the corresponding eigenvalues are degenerate within each irrep. Consequently, in the eigenbasis of $\mathcal{J}$, the Hamiltonian Eq. (6.8) decomposes into a sum over irreps,

$$H = \frac{1}{2} \sum_t \sum_I J_I \left| \boldsymbol{m}_t^{(I)} \right|^2 \tag{6.9}$$

where I labels the irreps, the $\boldsymbol{m}^{(I)}$ are linear combinations of the twelve spin components for a given tetrahedron, and J_I are the respective energy eigenvalues for each irrep. The irrep analysis

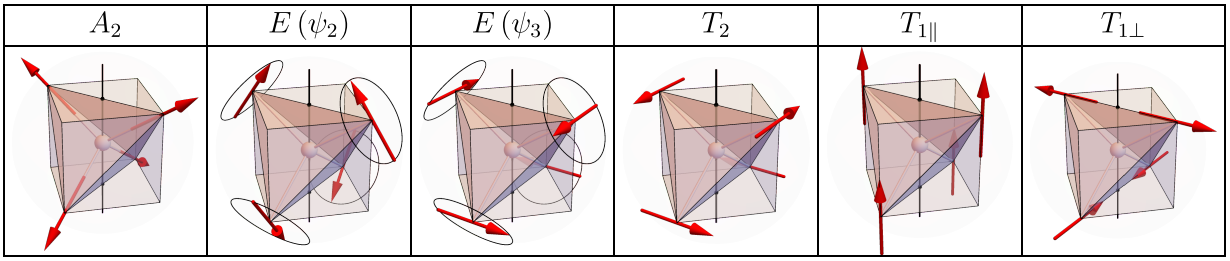


Figure 6.2: Symmetry broken ground states of the nearest-neighbor anisotropic Heisenberg model labeled by their corresponding irreps. The E irrep has a continuously degenerate ground state manifold, where all four spins can simultaneously rotate about their local [111] axes, indicated by circles, from which the high-symmetry ψ_2 and ψ_3 are selected by order by disorder. The T_2 irrep ground states (also called Palmer Chalker states) form a spiral when viewed along a Cartesian axis (black vertical line).

of the representation of T_d acting on this twelve dimensional space decomposes it into a direct sum of five irreps, $A_2 \oplus E \oplus T_1 \oplus T_1 \oplus T_2$, where A is 1-dimensional, E is 2-dimensional, and the T's are 3-dimensional (cf. Sections 2.3.3 and 2.5.2 as well as Fig. 2.13). In this decomposition, each irrep describes a different normal mode for a single tetrahedron, or in other words a different subspace of the twelve-dimensional soft spin configuration space for a single tetrahedron. With fixed-length spins, the intersection of the hard spin manifold (a Cartesian product of four spheres) with each of these subspaces produces a set of symmetry broken ground states for each irrep. Since the Hamiltonian is the same on every tetrahedron, the minimum J_I corresponds to a $\boldsymbol{q} = \boldsymbol{0}$ long range ordered ground state for the entire system in which every tetrahedron orders the same way. We will always take the convention that the ground state energy has been shifted to zero so that the minimum J_I is zero.

The hard spin ground states are as follows, illustrated in Fig. 6.2:

- A_2 corresponds to the two all-in/all-out (AIAO) ground state.
- E corresponds to the Γ_5 manifold, which is topologically a circle, out of which the highest-symmetry configurations are the six ψ_2 and six ψ_3 configurations, which are generally selected via order by disorder when J_E has the minimum energy [36, 195, 196].
- T_2 corresponds to the six chiral Palmer-Chalker (PC) ground states [197].
- The remaining six degrees of freedom split into two T_1 irreps, which we will call $T_{1\parallel}$ and $T_{1\perp}$. $T_{1\parallel}$ corresponds to the fully co-linear ferromagnetic configurations, whose manifold is topologically a sphere. $T_{1\perp}$ corresponds to a set of coplanar states in the Cartesian planes,

where when viewed along the axis orthogonal to the plane two opposite spins point towards the axis and the other two away from the axis.

In general, the $T_{1\parallel}$ and $T_{1\perp}$ modes can mix as they are both T_1 irreps. We will discuss this further in Section 6.3.

The Heisenberg antiferromagnet is characterized by the hierarchy of energy scales

$$J_{A_2} = J_E = J_{T_2} = J_{T_{1\perp}} = 0 < J_{T_{1\parallel}}. \tag{6.10}$$

The bands of this model are just three degenerate copies of the bands of NNSI, and the gapped $J_{T_{1\parallel}}$ corresponds to the threefold degenerate band maximum, while the remaining irreps correspond to the ninefold band touching at $\boldsymbol{q} = \mathbf{0}$, which correspond to the harmonic modes in NNSI.[4] It is clear that for temperatures $T \ll J_{T_{1\parallel}}$, the gapped $T_{1\parallel}$ irrep (the net moment) is suppressed on each tetrahedron, while the remaining zero-moment irreps are free to fluctuate at zero energy cost, which describes the Coulomb spin liquid phase.

Using the maximally degenerate Heisenberg AFM as a starting point, we can construct a variety of frustrated spin models which may realize various spin liquid phases by rearranging the hierarchy of irrep energies in Eq. (6.10). Gapping additional irreps introduces additional constraints to the original fully degenerate Heisenberg Coulomb spin liquid. So long as the constraints are not sufficiently strong to remove the ground state entropy and collapse the system into a long-range ordered phase, this procedure will yield novel descendent spin liquids. This conceptual approach was first employed by Benton et al. [120], who considered gapping A_2 in addition to $T_{1\parallel}$, yielding a spin liquid with non-analytic (pinched) correlations everywhere along the high-symmetry (111) *lines* in reciprocal space. In another work, Yan et al. considered gapping T_2 in addition to A_2 and $T_{1\parallel}$ [185], leaving only $T_{1\perp}$ and E to fluctuate. The degeneracy of these two remaining irreps, plus the addition of breathing anisotropy (removing the constraints on half of the tetrahedra) yields a rank-2 $U(1)$ spin liquid with sub-dimensional fracton excitations [185, 198, 199].[5] This phase is described by a coarse-grained symmetric traceless tensor field and exhibits *fourfold* pinch point singularities in the spin structure factor, which are direct reflections of the rank-2 divergence-free constraint [200].

While the procedure in [120, 185] is useful for exploring the space of interesting spin liquid phases, degeneracy of the irreps requires fine-tuning. Small perturbations, which are always

[4] Note a key difference here, in spin ice the threefold band touching point corresponds to the three components of the net magnetization, but here the net magnetization corresponds to the $T_{1\parallel}$ irrep and the band *maximum*, no longer corresponding to the harmonic modes.

[5] Although Ref. [185]sleu do not directly address the issue, it turns out that without the breathing anisotropy the constraints are just a bit too strong and the system orders via order by disorder.

present in material realizations of these models, will generically lift the degeneracy of the lowest energy modes. One irrep will generally have the lowest energy, leading it to condense into long-range order at low temperature, co-opting a zero-temperature spin liquid phase. Before we address the question of whether and to what extent Coulomb correlations can exist in the ordered phase, we are going to demonstrate how Conlon and Chalker's method of charges and fluxes is related to the irreducible representations through the rank-2 formalism first laid out in Ref. [185], because the connection between the two has not yet been identified in the literature. Then we will define a Helmholtz-Hodge decomposition for multi-component spins which we will use to study the physics above and below T_c.

6.3 Rank-2 Formalism, Irreps, and Fields

Yan et al. [185] demonstrated that the nine degrees of freedom labeled A_2, E, T_2 and $T_{1\perp}$ can be organized into a 3×3 matrix. In their construction, these irreps are identified with the anti-symmetric (T_2), symmetric off-diagonal ($T_{1\perp}$), diagonal trace-free (E), and scalar trace (A_2) parts of the matrix. In this section we demonstrate that this tensorial description can be derived directly from the flux description of Conlon and Chalker, and that it is in fact the natural way to generalize the flux method when perturbing the isotropic, O(3)-symmetric Heisenberg model.

In the approach laid out in Section 2.3.3, the long-wavelength expansion begins by studying the $\boldsymbol{q} = \mathbf{0}$ sector, which is equivalent to studying a single unit cell containing four spins. We treat the spin components S_i^α in a single unit cell as a 12-component vector. This vector is acted on by a 12-dimensional representation of the group T_d, under which the four corner indices indexed by $\mu \in \{1, 2, 3, 4\}$ transform as a 4-dimensional permutation representation which we denoted $\mathbf{4}$, and the spin component index α transforms as a 3-dimensional rotation representation denoted $\bar{\mathbf{3}}$ (see Fig. 2.13 for the character table and notation for the representations). The bar denotes that the spins transform as angular momentum *pseudo*-vectors rather than vectors, meaning under an inversion they change sign, see the discussion in Section 2.5.2. We denote the twelve-dimensional

representation **12**, which is a tensor product representation of the permutations and rotations,[6]

$$\mathbf{12} = \mathbf{4} \otimes \bar{\mathbf{3}}. \tag{6.11}$$

From here one could apply basic tools of character theory to decompose this tensor product representation into a direct sum of irreps of T_d as we discussed in the previous section. Alternatively, however, we can note that the permutation representation **4** is already reducible, as we did in Section 2.5.2: the symmetric linear combination $\sum_\mu S^\alpha_\mu$ is invariant under permutations and thus forms a 1-dimensional irreducible representation of **4**. Therefore, we can express $\mathbf{4} = \mathbf{1} \oplus \mathbf{3}$. The resulting three-dimensional representation **3** is an irreducible rotation representation: let $\hat{z}_\mu$ be four 3-component unit vectors pointing from the center of a tetrahedron to the corners, then a permutation of the corners indexed by μ is equivalent to an O(3) rotation of the tetrahedron and thus of the vectors $\hat{z}_\mu$. Thus, the twelve-dimensional representation Eq. (6.11) decomposes into a vector and matrix representation:

$$\mathbf{12} = (\mathbf{1} \oplus \mathbf{3}) \otimes \bar{\mathbf{3}} = \bar{\mathbf{3}} \oplus (\mathbf{3} \otimes \bar{\mathbf{3}}). \tag{6.12}$$

This demonstrates that the decomposition of the single-tetrahedron degrees of freedom into a 3-component vector and rank-2 tensor is in fact a natural consequence of the symmetries.

In the decomposition Eq. (6.12), the pseudo-vector $\bar{\mathbf{3}}$ corresponds to the net-moment irrep $T_{1\parallel}$, spanned by the linear combinations

$$m_\alpha^{(T_{1\parallel})} \equiv \frac{1}{2} \sum_\mu S^\alpha_\mu, \tag{6.13}$$

which corresponds in the flux language to the color charge Q_α, with the spin component being the color index. This is a pseudovector representation since, again, the spins $\boldsymbol{S}_\mu$ transform as angular momenta. The factor of half is a normalization, chosen such that the coefficients of the S^α_μ in the sum form a unit-length vector.

[6] In what follows, it is useful to consider the analogous problem of a the vibrational modes of a tetrahedral molecule. The analysis would be exactly the same, with the only difference being that $\bar{\mathbf{3}}$ in Eq. (6.11) is replaced with **3**. Then the various irreps correspond to the different symmetry modes of vibration. In particular, A_2 corresponds to the breathing mode, T_2 corresponds to rotations, ψ_2 and ψ_3 would correspond to a "twist" and a "stretch", $T_{1\parallel}$ is just the net displacement and $T_{1\perp}$ corresponds to a mode where, when viewed along a Cartesian axis, two spins move towards the axis and two move away. The analogy between continuum rank-2 spin liquids and elasticity is described by Yan and Reuther in Ref. [201], and was also used by Yan and Nevidomskyy to propose a method of realizing a rank-2 $U(1)$ gauge theory via phonon-nematic coupling in Ref. [202]. It also underlies the so-called fracton-elastic duality descibed by Ref. [203].

The remaining $(\mathbf{3} \otimes \bar{\mathbf{3}})$ tensor product representation is spanned by the nine components of

$$\mathcal{B}_{\alpha\beta} \equiv \frac{1}{2} \sum_{\mu} z_{\mu}^{\alpha} S_{\mu}^{\beta}, \tag{6.14}$$

where z_μ transform as vectors and $\boldsymbol{S}_\mu$ transform as pseudovectors. In this expression we define $z_\mu \equiv \sqrt{3}\hat{z}_\mu$, such that $z_\mu^\alpha \in \{-1, 1\}$. The $\mathcal{B}$ rank-2 tensor in this expression is precisely equivalent to the three vector fluxes $\boldsymbol{B}_c$ defined by Conlon and Chalker [76, 81], arranged into a 3×3 matrix, where α, the spin component index, is the color. From here we can naturally identify the remaining irreps by decomposing this matrix into its trace, anti-symmetric, and symmetric parts,

$$(\mathbf{3} \otimes \bar{\mathbf{3}}) = \bar{\mathbf{1}} \oplus \mathbf{3} \oplus \mathbf{5}, \tag{6.15}$$

where $\mathbf{1}$ is the scalar trace, $\mathbf{3}$ is the anti-symmetric part, and $\mathbf{5}$ is the trace-free symmetric part. This decomposition is well-known when studying the tensor product representations of $O(3)$, in which case it is just the Clebsh-Gordon decomposition of a tensor product of two spin-1 representations (if we ignore the psuedovector part). Restricting to the cubic symmetries of $T_d \subset O(3)$, the symmetric part decomposes further into its off-diagonal and diagonal parts, $\mathbf{5} = \mathbf{2} \oplus \bar{\mathbf{3}}$.

To see this in action, let us start with the trace of $\mathcal{B}$, from which we obtain the A_2 irrep,

$$m^{(A_2)} \equiv \frac{1}{\sqrt{3}} \sum_{\alpha} \mathcal{B}_{\alpha\alpha} = \frac{1}{2\sqrt{3}} \sum_{\mu} (\boldsymbol{S}_\mu \cdot z_\mu), \tag{6.16}$$

which is the sum of the projection of the spin vectors along their local [111] easy axes and is maximized in the all-in all-out configuration. Under inversion symmetry this changes sign since the $\boldsymbol{S}_\mu$ transform as pseudovectors, and we therefore denote this pseudoscalar irrep by $\bar{\mathbf{1}}$. Similarly, the anti-symmetric part of $\mathcal{B}$ corresponds to the T_2 irrep,

$$m_{\alpha}^{(T_2)} \equiv \frac{1}{\sqrt{2}} \epsilon_{\alpha\beta\gamma} \mathcal{B}_{\beta\gamma} = \frac{1}{2\sqrt{2}} \sum_{\mu} (\boldsymbol{S}_\mu \times z_\mu)_\alpha \tag{6.17}$$

where ϵ is the fully antisymmetric Levi-Civita symbol. Under an inversion Eq. (6.17) picks up two signs, one from the pseudovector transformation of $\boldsymbol{S}_\mu$ and another from the antisymmetric Levi-Civita symbol, and thus this transforms as the vector irrep denoted $\mathbf{3}$. Additionally, it is clear in Eq. (6.17) that the components of the $\boldsymbol{S}_\mu$ along the local z_μ drop out of the cross product and thus do not contribute to this irrep, i.e. the irrep only contains local xy components of the spins.

The remaining five degrees of freedom are contained in the trace-free diagonal and the symmetric off-diagonal components of the tensor $\mathcal{B}$. Starting with the trace-free diagonal, we

define the E irrep normal mode as

$$m_\alpha^{(E)} \equiv \sqrt{\frac{3}{2}}\left(\mathcal{B}_{\alpha\alpha} - \frac{1}{3}\operatorname{Tr}(\mathcal{B})\right) \tag{6.18}$$

$$= \frac{\sqrt{3}}{2\sqrt{2}}\sum_\mu \left(z_\mu^\alpha S_\mu^\alpha - \frac{1}{3}(\boldsymbol{S}_\mu \cdot \boldsymbol{z}_\mu)\right). \tag{6.19}$$

This has three components, but they are not linearly independent, since $\sum_\alpha m_\alpha^{(E)} = 0$, leaving us with the two-dimensional E irrep. A compact expression can be obtained for this order parameter by expanding each $\boldsymbol{S}_\mu$ as

$$\boldsymbol{S}_\mu \equiv \boldsymbol{S}_\mu^\perp + (\boldsymbol{S}_\mu \cdot \hat{\boldsymbol{z}}_\mu)\hat{\boldsymbol{z}}_\mu, \tag{6.20}$$

where $\boldsymbol{S}_\mu^\perp$ is the projection of $\boldsymbol{S}_\mu$ orthogonal to $\hat{\boldsymbol{z}}_\mu$. Then, using $(z_\mu^\alpha)^2 = |\boldsymbol{z}_\mu|^2/3$, $(\boldsymbol{S}_\mu \cdot \boldsymbol{z}_\mu)$ cancels in Eq. (6.19), and we are left with

$$m_\alpha^{(E)} = \frac{\sqrt{3}}{2\sqrt{2}}\sum_\mu z_\mu^\alpha (S_\mu^\perp)^\alpha, \tag{6.21}$$

demonstrating that this two-dimensional irrep, like T_2, does not contain the spin components along $\hat{\boldsymbol{z}}_\mu$. Finally, the remaining three degrees of freedom corresponding to the $T_{1\perp}$ irrep are contained in the off-diagonal symmetric part of $\mathcal{B}$. We define

$$m_\alpha^{(T_{1\perp})} \equiv \frac{1}{\sqrt{2}}|\epsilon_{\alpha\beta\gamma}|\mathcal{B}_{\beta\gamma}. \tag{6.22}$$

Let $\boldsymbol{e}_\alpha$ be the Cartesian basis vectors, then define

$$\tilde{\boldsymbol{e}}_x \equiv (\boldsymbol{e}_y + \boldsymbol{e}_z)/\sqrt{2},$$
$$\tilde{\boldsymbol{e}}_{\bar{x}} \equiv (\boldsymbol{e}_y - \boldsymbol{e}_z)/\sqrt{2},$$

and similarly for y, z by cyclic permutation, with which we can write

$$m_\alpha^{(T_{1\perp})} \equiv \frac{1}{2\sqrt{2}}\sum_\mu [(\boldsymbol{S}_\mu \cdot \tilde{\boldsymbol{e}}_\alpha)(\boldsymbol{z}_\mu \cdot \tilde{\boldsymbol{e}}_\alpha) - (\boldsymbol{S}_\mu \cdot \tilde{\boldsymbol{e}}_{\bar{\alpha}})(\boldsymbol{z}_\mu \cdot \tilde{\boldsymbol{e}}_{\bar{\alpha}})]. \tag{6.23}$$

These three modes indexed by α correspond to coplanar configuration in the cubic plane orthogonal to α, with two spins pointing towards each other along an edge of the tetrahedron and the other two pointing away from each other. Under inversion the spins change sign, so Eq. (6.23) transforms as a pseudovector, denoted $\bar{\mathbf{3}}$.

To summarize, under the action of the symmetry group T_d the four spins on a tetrahedron transform in a twelve dimensional representation which decomposes as

$$\begin{aligned}
\mathbf{12} &= \bar{\mathbf{3}} \oplus (\mathbf{3} \otimes \bar{\mathbf{3}}) && (6.24)\\
&= \bar{\mathbf{3}} \oplus (\bar{\mathbf{1}} \oplus \mathbf{3} \oplus \mathbf{5}) && (6.25)\\
&= \bar{\mathbf{3}} \oplus (\bar{\mathbf{1}} \oplus \mathbf{3} \oplus \mathbf{2} \oplus \bar{\mathbf{3}}) && (6.26)\\
&\equiv T_{1\parallel} \oplus (A_2 \oplus T_2 \oplus E \oplus T_{1\perp}). && (6.27)
\end{aligned}$$

Thus we have demonstrated that (i) the rank-2 description of the spin degrees of freedom can be naturally derived from a basic representation analysis, and (ii) the rank-2 tensor $\mathcal{B}$ can be understood as Conlon and Chalker's three color-copies of the 3-component vector $\boldsymbol{B}$. [7] The vector part of the irrep decomposition corresponds to the net tetrahedron moment $\boldsymbol{m}$, and the matrix part $\mathcal{B}$ contains the remaining zero-moment degrees of freedom. In the language of Conlon and Chalker, the net moment $\boldsymbol{m}^{(T_{1\parallel})}$ is interpreted as the three-color charge Q_c, whereas the three corresponding fluxes $\boldsymbol{B}_c$ underlying the gauge field description should be interpreted as the 3×3 matrix $\mathcal{B}$. This decomposition therefore tells us how the different colors couple in the presence of anisotropic interactions.

The fact that the resulting decomposition has the $\bar{\mathbf{3}}$ irrep appearing twice means that these two irreps can couple quadratically in the Hamiltonian (6.9). This coupling can be removed by making an orthogonal rotation between $T_{1\parallel}$ and $T_{1\perp}$ (Eq. (6.13) and Eq. (6.23)) by an angle θ [193]

$$\begin{pmatrix} m_\alpha^{(T_{1\theta})} \\ m_\alpha^{(T_{1\bar{\theta}})} \end{pmatrix} = \begin{pmatrix} \cos\theta & \sin\theta \\ -\sin\theta & \cos\theta \end{pmatrix} \begin{pmatrix} m_\alpha^{(T_{1\parallel})} \\ m_\alpha^{(T_{1\perp})} \end{pmatrix}. \tag{6.28}$$

The angle θ is called the canting angle, because the ground states associated to the two order parameters defined in Eq. (6.28) are so-called "splayed ferromagnet" configurations, where the spins are ferromagnetically aligned, but canted towards or away from the global cubic axes by precisely θ. This is illustrated in Fig. 6.3. The maximum canting angle is $\cos^{-1}(1/\sqrt{3})$, and in this limit we denote the two irreps T_{1z} and T_{1xy}, respectively, to indicate that these irreps correspond to ferromagnetic modes with all spins lying either along their local z axis or in their local xy planes. In terms of the vector plus rank-2 tensor decomposition, however, we note that the vector mode always corresponds to $T_{1\parallel}$ while the symmetric matrix mode corresponds to $T_{1\perp}$. If the system orders into a T_1 configuration with a non-zero canting angle the resulting configuration can be

[7] We note that the organization of the various irreps into a rank-2 tensor was first done by Yan et al. in Ref. [185], building on the previous work of Ref. [120], but the connection to the fluxes of Conlon and Chalker [76] was not established, and the authors of Ref. [185] did not provide an explanation of how the rank-2 description can be derived directly from the irrep analysis as we have done in this section.

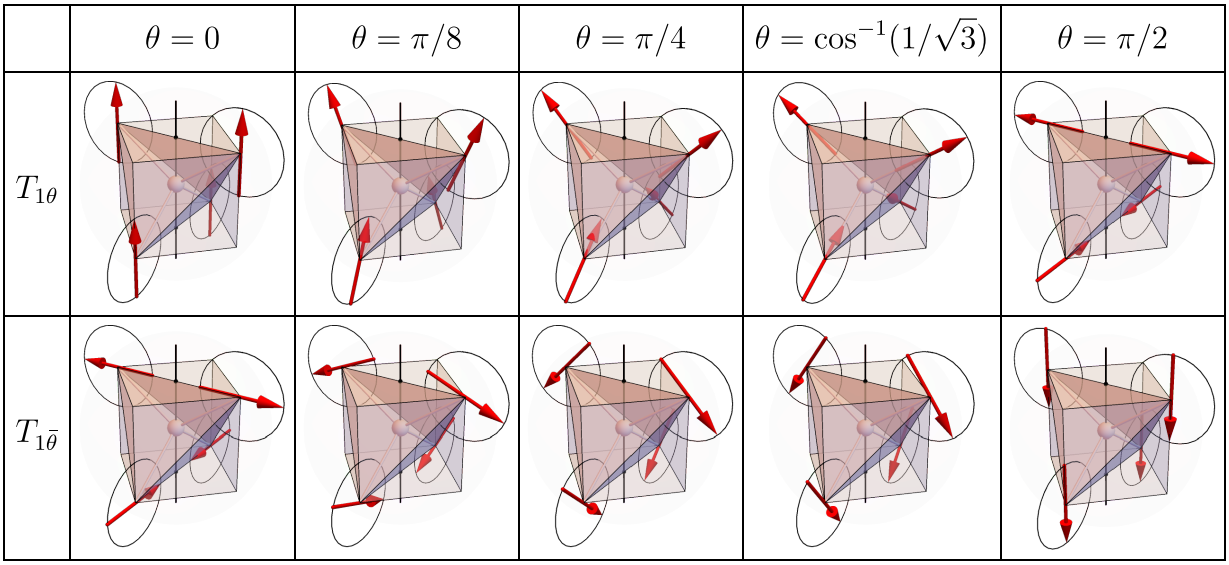


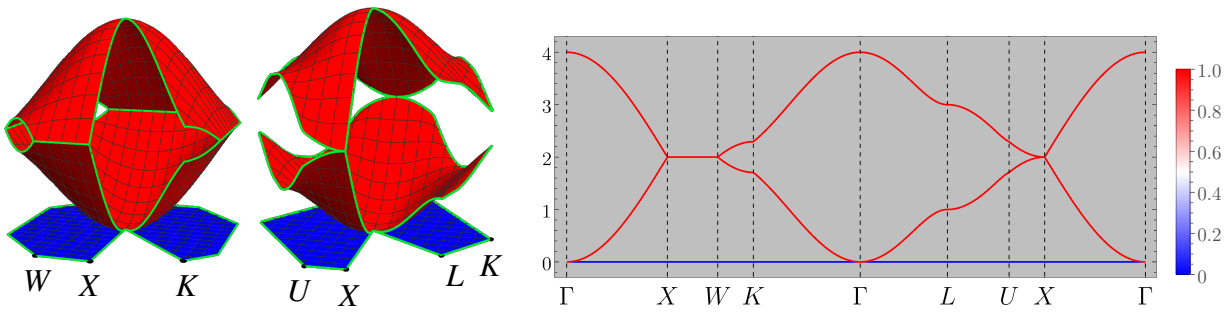
Figure 6.3: Illustration of the T_1 ground states as the canting angle is varied. At the maximum canting angle $\theta = \cos^{-1}(1/\sqrt{3})$, $T_{1\theta}$ corresponds to the ice configurations with spins along the local [111] axes, and $T_{1\bar{\theta}}$ has spins in the local xy planes. At $\theta = \pi/2$ the $T_{1\parallel}$ and $T_{1\perp}$ irrep swap.

interpret as a partial condensation of both the $T_{1\parallel}$ vector (the net moment) and the $T_{1\perp}$ symmetric matrix, which results in the spins canting away from the collinear configurations to configurations with cubic symmetry, yielding a reduced total magnetic moment proportional to the contribution of the $T_{1\parallel}$ irrep. We will discuss this further when we consider fluctuations about a symmetry broken ground state.

6.4 Helmholtz-Hodge Decomposition for 3-Component Spins

Thus far we have seen that the isotropic Heisenberg antiferromagnet decouples into three copies of the nearest-neighbor $N = 1$ antiferromagnet (Eq. (6.4)). Its Coulomb phase can therefore essentially be viewed as three decoupled copies of the single component Coulomb phase of spin ice,[8] and its long wavelength physics is described by three decoupled divergence-free vector fields $\boldsymbol{B}_c$. In the presence of anisotropy, the colors and vector components become coupled, and we

[8]The only caveat to keep in mind is that the spin ice ground states are actually ferromagnetic, and $\boldsymbol{B}$ corresponds to the net moment of a tetrahedron, while the Heisenberg antiferromagnet ground states have zero moment and the three components of Q_c correspond to the three component of the net magnetization of a tetrahedron.



(a) (b)

Figure 6.4: (a) The four bands of the pyrochlore adjacency matrix $\mathcal{A}$, or equivalently the twelve bands of $\mathcal{J}_{\text{Heis}}$, in the $(hk0)$ and $(hh\ell)$ planes, restricted for the first Brillouin zone (FBZ) of the fcc lattice. (b) Line cuts of the bands, with minimum eigenvalue shifted to zero. The two degenerate flat bands (blue) correspond to rotational modes, while the two dispersive bands (red) correspond to irrotational modes.

demonstrated how the nine components combine into a tensor which decomposes $\mathcal{B}$, which the anisotropy will decompose into its trace, anti-symmetric, trace-free diagonal, and off-diagonal symmetric components.

Unfortunately, what we do not have is a full treatment of this problem using the discrete differential forms approach utilized extensively in the previous chapters. In principle what we need is vector-valued differential forms, but it is not clear how to generalize the definition of the differential and codifferential in the anisotropic case, and therefore how to define the Helmholtz-Hodge decomposition. In lieu of such a full theoretical formalism to rigorously define the HHD for spin configurations, we will resort to defining the HHD in the most trivial way that we possibly can: by extending its definition isotropically to spin space, i.e. we define the following projectors:

$$\mathcal{P}^{\text{irr}} \equiv \mathcal{P}^{\text{im}\,\mathrm{d}_0} \otimes \mathbb{1}_{\text{spin}}, \tag{6.29}$$

$$\mathcal{P}^{\text{rot}} \equiv \mathcal{P}^{\text{im}\,\delta_2} \otimes \mathbb{1}_{\text{spin}}, \tag{6.30}$$

$$\mathcal{P}^{\text{har}} \equiv \mathcal{P}^{H^1} \otimes \mathbb{1}_{\text{spin}}. \tag{6.31}$$

We cannot say necessarily that these projectors are *per se* the "correct" HHD projectors for anisotropic spin systems, and indeed there may be a better definition once a full discrete Hodge theoretical treatment has been developed for $N = 3$ component spins, but what we can say is that they are the "correct" HHD projectors for the *isotropic* Heisenberg antiferromagnet. In Fig. 6.4 we show the bands in the isotropic limit, and have colored the flat rotational bands blue and dispersive

irrotational bands red. Now by *rotational* we mean "zero moment on every tetrahedron". Despite the lack of a formal definition for anisotropic models, we will see that this definition is still a very useful tool for studying models which are close to the isotropic Heisenberg antiferromagnet.

To provide a quantitative measure of the rotational-irrotational character of a spin configuration, we define a scalar quantity associated to a given spin configuration,

$$\mathbb{P}(S) \equiv \frac{|\mathcal{P}^{\text{irr}} S|^2}{|S|^2}, \tag{6.32}$$

which measures the violation of the "zero-divergence" or zero-moment constraint on each tetrahedron, such that $\mathbb{P}(S) = 1$ indicates a fully irrotational spin configuration and $\mathbb{P}(S) = 0$ indicates a purely rotational plus harmonic configuration.

6.5 Coulomb Physics Below T_c

In the previous section, we described in detail the AFM Heisenberg model, which is maximally frustrated and hosts a Coulomb phase characterized by non-analytic pinch points in the correlation functions. Many interesting phases can emerge from the massive degeneracy of the Heisenberg AFM, by tuning the degeneracy and hierarchy of the irrep energies J_I in Eq. (6.9). The introduction of a generic perturbation will break this degeneracy, resulting in the condensation of the lowest irrep and spontaneous symmetry breaking. Our aim now is to study what happens well below the ordering temperature in a model which is near to the spin liquid point. To that end, starting from the AFM Heisenberg model, we consider the introduction of a perturbation which results in the situation where the irrep energies satisfy

$$J_{A_2} < J_E \approx J_{T_2} \approx J_{T_{1\perp}} \ll J_{T_{1\parallel}}, \tag{6.33}$$

such that the system has a phase transition AIAO order (condensing the A_2 irrep). We will investigate the behavior of the spin wave fluctuations around the ground state via a low-temperature expansion.

6.5.1 The Easy Axis Model

The model we will study is the AFM Heisenberg model plus a weak single-ion anisotropy (SIA), described by the Hamiltonian

$$H = J \sum_t \sum_{\langle ij \rangle \in t} \boldsymbol{S}_i \cdot \boldsymbol{S}_j - \frac{D}{2} \sum_i (\boldsymbol{S}_i \cdot \hat{\boldsymbol{z}}_i)^2 , \tag{6.34}$$

with $J \gg D > 0$. This is a small perturbation to the Heisenberg AFM, with the SIA term favoring alignment of spins along their local [111] axes.[9] The irrep energies for this model are (after shifting the ground state energy to zero)

$$J_{A_2} = 0\,, \quad J_E = D\,, \quad J_{T_2} = D\,, \tag{6.35}$$

$$J_{T_{1\perp}} \approx \frac{D}{3}\,, \quad J_{T_{1\parallel}} \approx J + \frac{2D}{3}\,, \tag{6.36}$$

which satisfy the conditions laid out in Eq. (6.33), i.e. $J_E \approx J_{T_2} \approx J_{T_{1\perp}}$. The single ion anisotropy selects out of the zero-moment Coulomb phase configurations the AIAO configuration as the ground state [10]

As can be seen in the temperature evolution of the structure factor in Fig. 6.1, at high temperatures $T \gg J$, the system is paramagnetic, then as temperature is lowered it crosses over into a Coulomb regime in the range $D \ll T \ll J$. Here the net moment of every tetrahedron is strongly suppressed, and we see the formation of not-quite-sharp pinch point features. At low temperatures $T \ll T_c \sim D$, the system orders into the AIAO phase, forming sharp Bragg peaks (green points) which contain the vast majority of the spectral weight. Our interest is in trying to understand the weak background correlations which are reminiscent of the partially formed pinch points above T_c.

6.5.2 Low Temperature Expansion

To study the low-temperature physics and these pinch-point-like features, we perform a low-temperature-expansion for small Gaussian fluctuations about the AIAO ground state, the details of which are laid out in Appendix D.1. Doing so yields an effective quadratic Hamiltonian,

$$H_{\text{eff}} = \frac{J}{2} \sum_t \Big| \sum_{i \in t} \boldsymbol{S}_i^{\perp} \Big|^2 + \frac{D}{2} \sum_i |\boldsymbol{S}_i^{\perp}|^2, \tag{6.37}$$

where $\boldsymbol{S}_i^{\perp} \equiv \boldsymbol{S}_i - (\boldsymbol{S}_i \cdot \hat{\boldsymbol{z}}_i)\hat{\boldsymbol{z}}_i$ is the portion of the spin vector $\boldsymbol{S}_i$ transverse to the local ordering axes $\hat{\boldsymbol{z}}_i$. In Eq. (6.37), the first term is almost exactly equivalent to the Heisenberg AFM Hamiltonian, Eqs. (6.3) and (6.5), except that the $\boldsymbol{S}_i^{\perp}$ are vectors constrained to the local xy planes. This already implies that fluctuations about the ground state energetically prefer to remain within the

[9] This can also be used as a model for spin ice if we take J to be ferromagnetic [204].

[10] Note again that the AIAO configuration is a ground state of the Heisenberg AFM, whereas it was the maximum energy configuration of NNSI. Also note that the single ion anisotropy term induces a weak coupling between $T_{1\parallel}$ and $T_{1\perp}$ and thus a small canting. This will not be relevant to our discussion since $D/J \ll 1$ and neither of these modes is critical for the AIAO ordering discussed here.

zero-moment irrotational sector, which suggests the origin of the pinch-point-like features: the fluctuations are energetically constrained to lie in the Coulomb phase ground state manifold.

However, unlike the original spin model, the $\boldsymbol{S}_i^{\perp}$ do not have a conserved spin length. They describe continuous rotations of the spin vectors $\boldsymbol{S}_i$ away from their ordering axes, and at $T = 0$ all of the $\boldsymbol{S}_i^{\perp} = \mathbf{0}$ (classically). Instead of spin length conservation, the size of fluctuations is controlled by the second term, with *all* fluctuations now gapped by D. Note that D plays the same role as $\lambda_0 T$ does in the SCGA description of spin liquids.

It is clear that the physics of this model must be analogous to that of the Heisenberg AFM. Specifically, a subset of excitations can be described as irrotational, producing a non-zero net moment on tetrahedra with energy cost proportional to J. The remaining excitations are rotational, maintaining zero moment everywhere. These must be described by a closed loop containing two spins per tetrahedron, such that the $\boldsymbol{S}_i^{\perp}$ in the loop cancel pairwise on each tetrahedron. These could be interpreted as "remnant" rotational modes inherited from the parent Coulomb phase above T_c. However, these closed-string modes cost an energy proportion to their length, since any deviation of the $\boldsymbol{S}_i^{\perp}$ from zero costs energy proportional to D. Thus the symmetry broken phase is similar to a Coulomb phase, in that the excitations are string-like, but with the difference that (i) excitations are now confined, and (ii) the lack of conserved spin length means that the the total spectral weight of excitations is not conserved, meaning that in the ground state the system has no excitations and is not liquid like. This is analogous to a Higgsed phase [205].

6.5.3 HHD Analysis of the Hessian Band Structure

At this point it might be tempting to conclude that the low temperature phase, while retaining some properties of a Coulomb phase (or perhaps a Higgsed Coulomb phase) is really rather trivial. The lack of conserved spin length means that there will always be a constant relative density of irrotational defects in the system of the order D/J, whereas in the Coulomb phase the number density of irrotational defects goes as $\lambda_0 T/J$, which goes to zero at low temperature. Therefore naively, the fluctuations around the ground state can never achieve the singular correlations associated to having a strict zero-divergence constraint. Upon closer analysis, however, we can observe that this system does have some rather interesting properties.

As we have done throughout this thesis, let us look at the spectrum of the effective interaction matrix in the effective Hamiltonian (6.37), and utilize them using the HHD projectors we previously introduced. To do so, we start by writing the quadratic effective Hamiltonian as as a bilinear form (see Appendix D.1.1 for details),

$$H_{\text{eff}} = \frac{1}{2} \sum_{ij} \sum_{ab} (S_i^{\perp})^a \tilde{\mathcal{H}}_{ij}^{ab} (S_j^{\perp})^b, \tag{6.38}$$

where $a, b \in \{x, y\}$ index the two components of $\boldsymbol{S}_i^\perp$ in the local xy planes. The matrix $\tilde{\mathcal{H}}$ is the Hessian, which describes the quadratic term in a Taylor series expansion of Eq. (6.34) about the fully ordered ground state, and plays the analogous role of the interaction matrix $\mathcal{J}$ in Eq. (6.8).

We denote the eigenvalues of $\tilde{\mathcal{H}}$ by $\epsilon_n(\boldsymbol{q})$, with corresponding eigenvectors $\phi_{\boldsymbol{q},n}$, indexed by crystal momenta $\boldsymbol{q}$ and a band index $n \in \{1, \ldots, 8\}$. The eigenvectors $\phi_{\boldsymbol{q},n}$ define the low-energy normal modes for fluctuations of the $\boldsymbol{S}_i^\perp$, via the linear combinations

$$S^\perp_{\boldsymbol{q},n} \equiv \sum_{i,a} [\hat{\phi}_{\boldsymbol{q},n}]_i^a (S_i^\perp)^a, \tag{6.39}$$

where the hat on $\hat{\phi}_{\boldsymbol{q},n}$ denotes normalization. H_{eff} is then diagonalized into a sum of decoupled quadratic modes,

$$H_{\text{eff}} = \frac{1}{2} \sum_{\boldsymbol{q}} \sum_{n=1}^{8} \epsilon_n(\boldsymbol{q}) |S^\perp_{\boldsymbol{q},n}|^2. \tag{6.40}$$

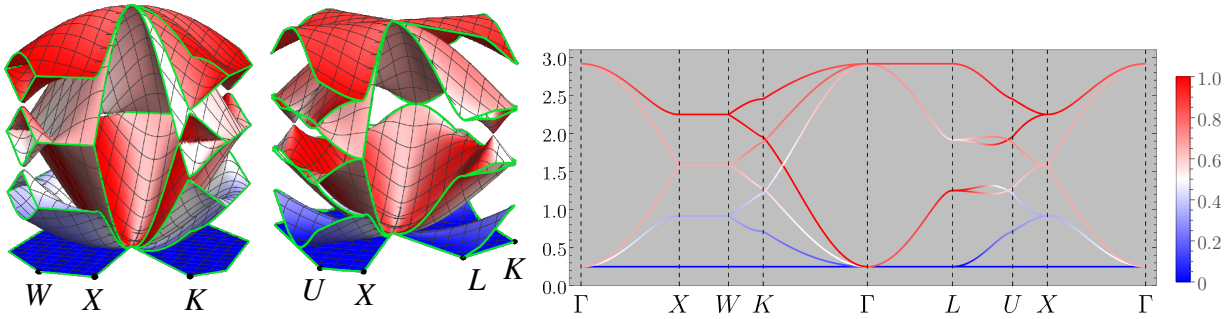
Since we are working within a Gaussian approximation, the thermal populations of the normal modes can be directly read off,

$$\langle |S^\perp_{\boldsymbol{q},n}|^2 \rangle = \frac{T}{\epsilon_n(\boldsymbol{q})}. \tag{6.41}$$

The eigenvalues are bounded below by the diagonal "mass" term D in Eq. (6.37), $\epsilon_n(\boldsymbol{q}) > D$, and as such the population of the lowest energy modes tends to zero as the temperature is decreased. This contrasts to the Coulomb spin liquid, where the lowest energy modes have energy $\lambda_0 T$ (in the SCGA) and the conserved spin length along with the flat band degeneracy allows these modes to continue fluctuating down to zero temperature.

To expose Coulombic physics in the low temperature fluctuations, we look at the bands of the Hessian and utilize the HHD projections previously defined. In Fig. 6.5, we plot $\epsilon_n(\boldsymbol{q})$, and we color the bands according to the irrotational projection of the eigenvectors Eq. (6.32), $\mathbb{P}(\phi_{\boldsymbol{q},n})$, where blue indicates rotational modes and red indicates irrotational modes.[11] We observe two completely flat bands with eigenvalue D, which are fully rotational (blue). These flat rotational bands correspond to the closed-loop fluctuations, and the confinement is reflected in the bands sitting at finite energy D. The six dispersive bands have *mixed* rotational and irrotational character, with the higher bands being strongly irrotational (red). One of the most interesting observations we can make from this is that the HHD reveals singularities in the *eigenvectors* of the dispersive bands near the five-fold band touching point at zone center, indicating that the long-wavelength physics is non-trivial in a way distinct from the usual Coulomb phase.

[11] Note that to do so we first rotate the vectors $\phi_{\boldsymbol{q},n}$, which have $2 \times 4L^3$ components, back to the global frame where the projections are defined, with $3 \times 4L^3$ components. See Appendix D.1 for more details.



(a) (b)

Figure 6.5: (a) The eight bands of the Hessian for the easy axis model with $D = 0.25J$ in the FBZ in the $(hk0)$ and $(hh\ell)$ planes, colored by the irrotational projection of the corresponding eigenvectors, $\mathbb{P}(\phi_{\boldsymbol{q},n})$. (b) Corresponding line cuts of the Hessian bands along a high-symmetry path, the same as Fig. 6.4, again colored by HHD projection. There are two degenerate fully rotational (blue) flat bands at the bottom of the spectrum with $\epsilon_{\text{flat}}(\boldsymbol{q}) = D$. The upper bands, with energy of order J, are dominant irrotational (red).

6.5.4 Correlations and Pinch Points

The correlation functions for the $\boldsymbol{S}_i^{\perp}$ can be obtained directly by inverting the Hessian,

$$\langle (S^{\perp}_{-\boldsymbol{q}})^a_{\mu} (S^{\perp}_{\boldsymbol{q}})^b_{\nu} \rangle = T \sum_{n=1}^{8} \frac{[\hat{\phi}^*_{\boldsymbol{q},n}]^a_{\mu} [\hat{\phi}_{\boldsymbol{q},n}]^b_{\nu}}{\epsilon_n(\boldsymbol{q})}. \tag{6.42}$$

Mapping the local transverse spin components, indexed by a, b, back to global 3-component vector indices α, β, we define the transverse spin-spin susceptibility tensor[12] as

$$\chi^{\perp}_{\alpha\beta}(\boldsymbol{q}) \equiv \frac{1}{T} \sum_{\mu,\nu} \langle (S^{\perp}_{-\boldsymbol{q}})^{\alpha}_{\mu} (S^{\perp}_{\boldsymbol{q}})^{\beta}_{\nu} \rangle = \frac{1}{T} \frac{1}{L^3} \sum_{ij} \langle (S^{\perp}_i)^{\alpha} (S^{\perp}_j)^{\beta} \rangle e^{-i\boldsymbol{q}\cdot\boldsymbol{r}_i}. \tag{6.43}$$

We then define the transverse spin structure factor as the trace of this tensor,

$$S^{\perp}_{\text{Spin}}(\boldsymbol{q}) \equiv \sum_{\alpha} \chi^{\perp}_{\alpha\alpha}(\boldsymbol{q}) = \frac{1}{T} \sum_{\mu,\nu} \langle (\boldsymbol{S}^{\perp}_{-\boldsymbol{q}})_{\mu} \cdot (\boldsymbol{S}^{\perp}_{\boldsymbol{q}})_{\nu} \rangle. \tag{6.44}$$

[12] Transverse in the sense of only containing contributions from the spin components transverse to the local ordering axes.

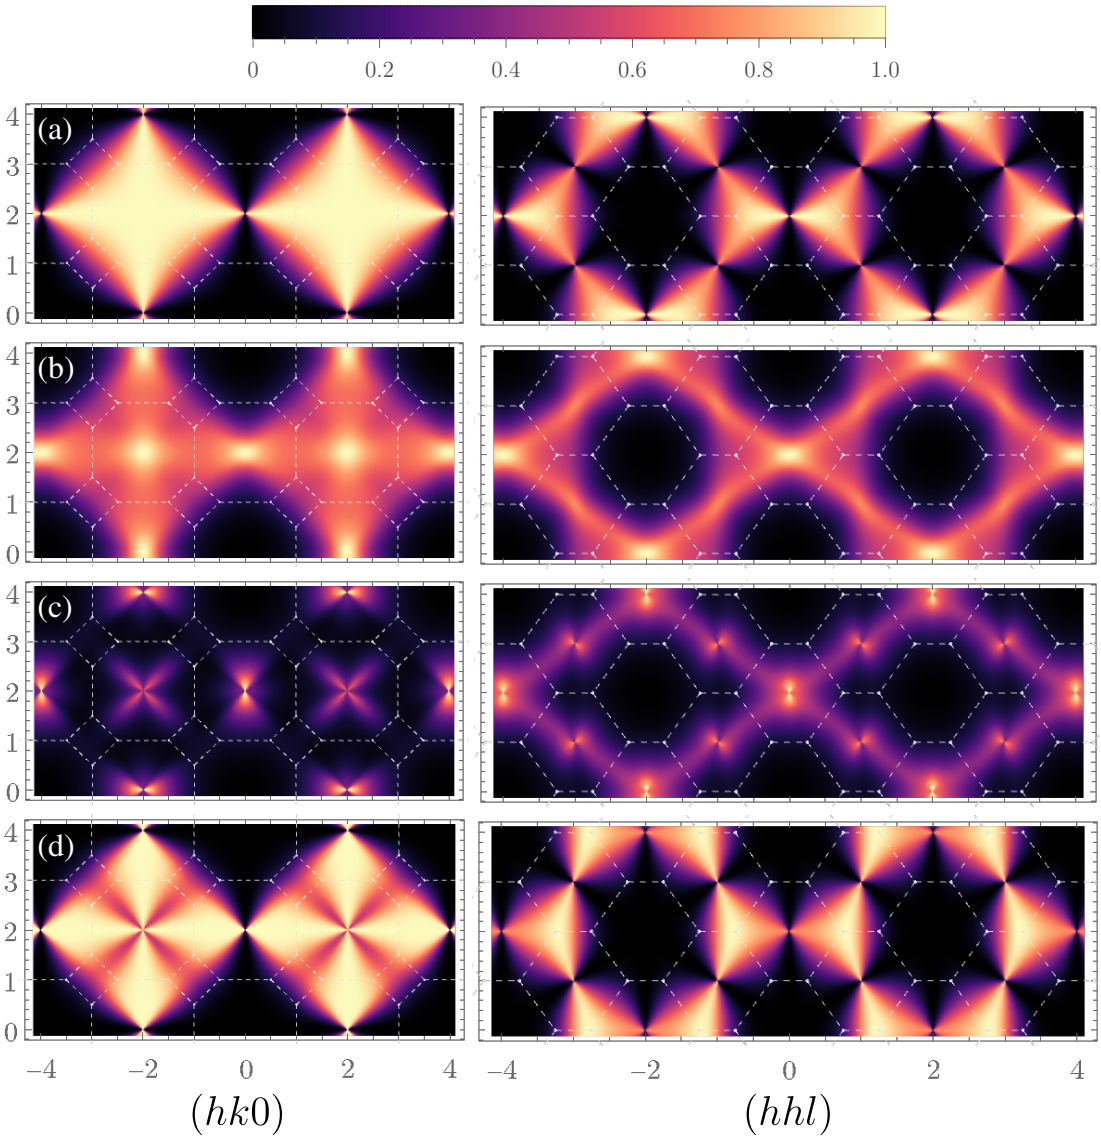


Figure 6.6: Spin structure factors in the $(hk0)$ (left) and (hhl) (right) planes for (a) the Heisenberg AFM Coulomb phase; (b) the easy axis model below the critical temperature, computed from the Hessian using Eq. (6.44); (c) the easy axis model restricting the sum in Eq. (6.42) to only the upper six dispersive bands; (d) the same but for only the bottom two flat bands. Adding the contributions of (c) and (d) yields the total structure factor in (b), with the complementary fourfold pinch points combining to create non-singular, broadened pinch point features.

Fig. 6.6(b) depicts the spin structure factor Eq. (6.44) in the $(hk0)$ and (hhl) planes (left and right columns, respectively). The structure factor exhibits features similar to the Heisenberg AFM Coulomb phase, whose structure factor is shown in (a), but with quite broadened pinch features rather than sharp singularities. One might have naively wanted to stop here, assuming that the lack of singularities precludes the presence of genuine Coulomb physics. However, motivated by the HHD analysis of the bands in Fig. 6.4, we separate the contribution of the the six predominantly irrotational dispersive bands, shown in (c), from the contribution of the two flat rotational bands, shown in (d), by restricting the sum over bands in Eq. (6.42). Doing so reveals *fourfold* pinch point singularities in the respective correlations.

The singularities occur at the five-fold band touching point at zone center where bands with different rotational and irrotational character meet: the flat bands are purely rotational while the bands dispersing away from zone center have non-zero irrotational content as seen in Fig. 6.5. Thus the low energy fluctuations about the symmetry broken ground state do in fact contain non-trivial Coulombic correlations, and the presence of fourfold pinch points implies that this structure is actually *richer* than the standard Coulomb physics of the parent Heisenberg model, which only exhibits twofold pinch singularities, with four-fold pinch points associated to rank-2 $U(1)$ spin liquids [185, 200, 202]

Note however that unlike a classical spin liquid, where the upper bands are depopulated at low temperature and only the correlations of the lower bands remain, in this case all modes are populated according to Eq. (6.41). The pinch points of the upper bands in (c) and lower bands in (d) are complementary: when added together, they yield the non-singular features in (b), which are what one would observe in a static (energy integrated) neutron scattering experiment. Thus it seems that the fluctuations in this system are perpetually stuck in what amounts to a finite-temperature spin liquid phase, the pinch points never sharpening to true singularities. We will demonstrate, however, that singular correlations can in fact be exposed in *inelastic* energy-resolved neutron scattering experiments. Before we get to dynamics, however, we will first show that the fourfold pinch points are indeed associated to a long-wavelength tensor field theory using the rank-2 prescription we developed in Section 6.3.

6.5.5 Rank-2 Tensor Field Theory

The observation of fourfold pinch points singularities in the correlation functions would seem to imply that the Hamiltonian can be expressed in terms of a rank-2 field theory in the long-wavelength limit similar to a rank-2 spin liquid. Note that under T_d symmetry the $\boldsymbol{S}_i^{\perp}$ transform as pseudovectors just like the $\boldsymbol{S}_i$, meaning that we can use the machinery laid out in Section 6.3 to perform the same analysis for the Hessian below T_c as we would do for the interaction matrix

above T_c. Thus, we define a vector and rank-2 tensor to describe the fluctuating degrees of freedom for each tetrahedron

$$Q_\alpha \equiv \frac{1}{2}\sum_{i\in t}(S_i^\perp)^\alpha, \tag{6.45}$$

$$\mathcal{M}_{\alpha\beta} \equiv \frac{1}{2}\sum_{i\in t} z_i^\alpha (S_i^\perp)^\beta. \tag{6.46}$$

Note that this construction naturally organizes the degrees of freedom into a three-component vector and a nine-component tensor, even though there are only eight degrees of freedom total corresponding to the components of the vector orthogonal to the local easy axis. We know from the earlier discussion that the trace-free diagonal (E irrep) and the anti-symmetric part (T_2 irreps) faithfully reflect five of these eight degrees of freedom, because they only involve contributions from the spin components orthogonal to $\hat{\boldsymbol{z}}_i$ (cf. Eq. (6.17) and Eq. (6.21)). We also know that $\mathcal{M}$ has zero trace since $\boldsymbol{S}_i^\perp \cdot \boldsymbol{z}_i = 0$, which means the remaining three degrees of freedom are split between the two T_1 modes, the vector Q and the off-diagonal symmetric part of $\mathcal{M}$. The only T_1 irrep which can be constructed with only the local xy components of the spins is T_{1xy}, given by the linear combination defined in Eq. (6.28) with $\theta = \cos^{-1}(1/\sqrt{3})$,

$$m_\alpha^{(T_{1xy})} \equiv \sqrt{\frac{1}{3}} m_\alpha^{(T_{1\perp})} - \sqrt{\frac{2}{3}} m_\alpha^{(T_{1\parallel})}. \tag{6.47}$$

The insight of the rank-2 irrep construction is that a portion of the three T_{1xy} degrees of freedom must be represented in the vector Q while the rest go into the symmetric part of $\mathcal{M}$. Fluctuations of T_{1xy} have a net-moment $T_{1\parallel}$ component which is suppressed by the large energy scale J and a zero-moment $T_{1\perp}$ component which is suppressed by the much smaller energy scale D.

To see this in action, we now rewrite the effective Hamiltonian Eq. (6.37) in terms of Q and $\mathcal{M}$, in terms of which it takes the remarkably simple form,

$$H_{\text{eff}} = \sum_t \left[2J\,|Q_t|^2 + D\,\text{Tr}(\mathcal{M}_t^{\text{T}}\mathcal{M}_t)\right]. \tag{6.48}$$

This cleanly demonstrates that the net moment Q_t is gapped by J, while the remaining degrees of freedom in $\mathcal{M}$ are gapped by D. The T_{1xy} modes obtain a contribution from *both* terms as a consequence of the decomposition in Eq. (6.47), while E and T_2 only appear in the second term. Correspondingly, the eigenvalues of the $\boldsymbol{q} = \boldsymbol{0}$ Hessian are

$$J_E = J_{T_2} = D,\; J_{T_{1xy}} = D + \frac{8J}{3}, \tag{6.49}$$

demonstrating that the T_{1xy} irrep modes are split between the vector net-moment and the rank-2 antisymmetric tensor. The three T_{1xy} modes correspond to the band maximum in Fig. 6.5 which is gapped away from the lower fivefold $E \oplus T_2$ band touching point by $8J/3 \gg D$. Thus the thermal population of these non-zero moment, irrotational, modes will be diminished by a factor $D/J \ll 1$ relative to the zero-moment, rotational, modes.

Our aim now is to perform a long-wavelength expansion of the Hessian to construct an effective field theory paralleling Eq. (6.7). Since the T_{1xy} modes have a much lower thermal population than the E and T_2 modes, we focus only on the lower fivefold $E \oplus T_2$ band touching point. By expanding the Fourier transformed Hessian $\tilde{\mathcal{H}}$ to quadratic order in $\boldsymbol{q}$ and restricting to the five lower bands, we obtain the effective long-wavelength action

$$\beta H_{\text{eff}}^{E\oplus T_2} = \frac{1}{2}\int \mathrm{d}^3\boldsymbol{q}\left[\beta D\,\mathrm{Tr}\left[\mathcal{M}^T\mathcal{M}\right] + \xi^2|\boldsymbol{q}\cdot\mathcal{M}|^2\right] + O(\boldsymbol{q}^4), \tag{6.50}$$

where $\xi \propto \sqrt{\beta J}a$ has dimensions of length, where a is the pyrochlore nearest-neighbor separation distance. We note one rather interesting fact, which is that the tensor $\mathcal{M}$ describing the low energy long-wavelength modes of this system has both a trace-free diagonal and an *anti*-symmetric part, meaning that it does not describe the usual $U(1)$ rank-2 spin liquid described by a symmetric trace free tensor [199, 200]

Note that if this were a true spin liquid the first term would independent of T at low temperature, but the ratio between the two terms is constant, with the second term larger than the first by a factor J/D. In the limit $D \to 0$ we obtain the rank-2 zero-divergence constraint $\sum_\alpha q^\alpha \mathcal{M}^{\alpha\beta} = 0$ for each β, but for finite D/J this constraint is never satisfied exactly, so the pinch points in Fig. 6.6 never become sharp.

6.5.6 Dynamics, HHD, and Inelastic Neutron Scattering

So far we have demonstrated that the band structure of the Hessian for small fluctuations about the symmetry-broken AIAO ground state of the easy axis model can be analyzed using the same methods employed to study Coulombic spin liquids on the pyrochlore lattice. In particular, the Hessian describing the low-temperature fluctuations of this model exhibits purely rotational flat bands inherited from the parent Heisenberg AFM, with a five-fold band touching point with three dispersive bands, corresponding to the $E \oplus T_2$ zero-moment fluctuating modes, and this band touching point can be described in the long wavelength limit by utilizing the same formalism used to describe rank-2 Coulombic spin liquids. Correspondingly, there are non-analyticities in the eigenvector structure of the Hessian around the band touching point, which we exposed using the HHD in Fig. 6.5.

The singular structure can be exposed by separating the contributions of the lower and upper Hessian bands, revealing fourfold pinch point correlations in Fig. 6.6(c) and (d), but these singularities are not experimentally accessible with static (energy integrated) neutron scattering, which receives contribution from all bands and would reveal only the broad non-singular integrated intensity reminiscent of broadened pinch points, seen in Fig. 6.6(b). In this section we consider how these singularities can be exposed using *inelastic* energy resolved neutron scattering. To do so, we consider the spin wave dynamics governed by the Hessian, and characterize the dynamical spin wave modes by the application of the HHD projection.

We consider classical spin precession dynamics for small fluctuations around the ground state, governed by the equation of motion

$$\frac{\mathrm{d}}{\mathrm{d}t}\boldsymbol{S}_i = -\boldsymbol{S}_i \times \boldsymbol{H}_i, \tag{6.51}$$

where $\boldsymbol{H}_i$ is the local molecular field, given by $\boldsymbol{H}_i \equiv \partial H/\partial \boldsymbol{S}_i$. Linearizing the equations of motion about the ground state and seeking periodic solutions, the equations of motion are reduced to a generalized eigenvalue problem,

$$\tilde{\mathcal{H}}\psi_{\boldsymbol{q},m} = -i\omega_m(\boldsymbol{q})\Gamma\psi_{\boldsymbol{q},m} \tag{6.52}$$

where the generalized eigenvalues $\omega_m(\boldsymbol{q})$ are spin wave precession frequencies, indexed by crystal momentum $\boldsymbol{q}$ and band index $m \in \{1, \ldots, 8\}$, $\tilde{\mathcal{H}}$ is the Hessian matrix appearing in Eq. (6.38), and Γ is a symplectic matrix.[13] Note that for each solution $\psi_{\boldsymbol{q},m}$ with frequency $\omega_m(\boldsymbol{q})$, its complex conjugate $\psi^*_{\boldsymbol{q},m}$ is another solution with frequency $-\omega_m(\boldsymbol{q})$. We therefore restrict our attention to the four positive-frequency bands.

For each eigenvector the corresponding normal mode is given by the linear combination

$$[S^{\perp}_{\boldsymbol{q},m}(t)]^a_i = [\psi_{\boldsymbol{q},m}]^a_i e^{i\omega_m(\boldsymbol{q})t}. \tag{6.53}$$

We wish to analyze these modes utilizing the HHD, since at any time t the spin configuration can be decomposed into rotational and irrotational components according to the HHD, and in Appendix D.3 we show that $\mathbb{P}(\psi_{\boldsymbol{q},m})$ is equivalent to the average irrotational projection over a

[13] See Appendix D.2 for the derivation and details of the approximations required to arrive at this equation and see Ref. [206] for a Lagrangian perspective on this problem and a comparison with the equivalent quantum spin wave results. We note that the quantum spin wave calculation using Holstein-Primakoff bosons gives the same band structure but slightly different intensities for the neutron scattering due to Bose statistics for magnons rather than Maxwell-Boltzmann statistics for spin waves. We use the semi-classical spin wave approximation so that we can analyze the dynamical eigenvectors, which is complicated by a Bugoliubov transformation in the quantum spin wave calculation.

full period. The result is shown in Fig. 6.7(a,b), where we show the magnon bands $\omega_m(\boldsymbol{q})$ for the easy axis model colored by the average irrotational projection. This reveals that lower two bands are almost completely rotational, meaning that the precessional dynamics do not violate the zero-moment condition, with every tetrahedron maintaining almost zero moment throughout each period. The upper bands on the other hand are strongly mixed (white), except near the $\{hhh\}$ lines (from Γ to L), where the top band is strongly rotational and the lower band is strongly irrotational. Indeed in the second plot in (a) we can see that the HHD structure of these two upper bands forms a fourfold pinch. The upper two bands touch one of the lower bands at the zone center, and their HHD content does not smoothly match at the band touching point, indicating a singularity in the eigenvector structure near this point and the corresponding modes. We conclude that the singularities at the five-fold band touching point in the Hessian spectrum have translated to a singularity at this threefold band touching point in the dynamical spin wave spectrum.

In Fig. 6.7(c) we show the inelastic neutron scattering intensity on the bands, in which we can see complementary pinch points on the lower and upper bands dispersing away from the band touching point. This is a clear signature that the dynamical spin wave modes of this system exhibit highly non-trivial Coulombic correlations. Integrating up the scattering intensity for the top two mixed bands and the bottom two rotational bands separately in (d) (top and bottom, respectively), we observe complementary pinch point patterns. Thus we have established that, not only is there Coulombic physics of the excitations below the critical temperature, but in fact the dynamics of the fluctuations obey strong Coulombic constraints and are thus highly correlated in ways that are not obvious from the static, energy-integrated structure factors.

6.6 Conclusion

To summarize, in this chapter we considered the question of whether Coulombic physics can survive below a phase transition to long range order out of a Coulomb phase. To study this, we considered a minimal model which can be described as the Heisenberg antiferromagnet plus a small perturbation selecting one of the ground states and forcing the system into long range at low temperatures. Performing an HHD analysis on the band structure of Hessian reveals singularities in the decomposition of the bands at the band touching points. By performing a long-wavelength expansion we showed that the low temperature fluctuations deep in the ordered phase can be described using the same mathematical tools used to describe rank-2 $U(1)$ Coulombic spin liquids, but that the emergent Gauss' law (zero divergence constraint) will never be fully realized. We then solved the precessional equations of motion (numerically) and applied the HHD to measure the average irrotational content of these modes over a full period, and found that of the four spin wave bands the two lower energy bands are almost fully rotational, meaning that these modes fluctuate

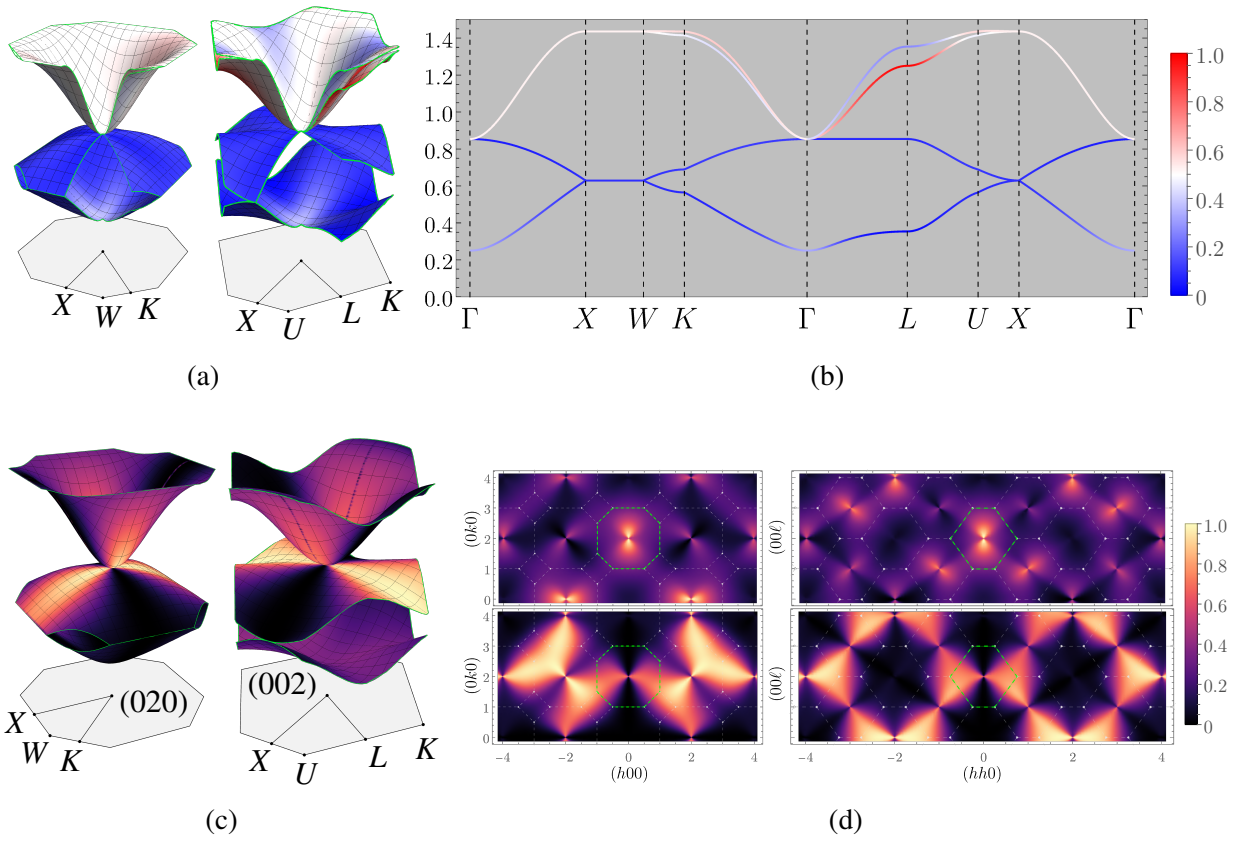


Figure 6.7: (a,b) Magnon frequency bands colored by irrotational projection of their eigenvectors in (a) the $(hk0)$ plane and the $(hh\ell)$ plane, (b) along with line cuts along a high symmetry path in reciprocal space (cf. Fig. 2.12). The two lower bands are strongly rotational (blue), while the upper two bands have mixed character, with rotational and irrotational modes meeting at a threefold band touching point. (c) Inelastic neutron scattering intensity from the bottom three bands (upper band removed for clarity) in the $(hk0)$ plane centered at (020) and the $(hh\ell)$ plane centered at (002), showing singularities at the threefold band touching point. (d) integrated scattering intensity of the upper two (top) and lower two (bottom) bands in the $(hk0)$ (left) and (hhl) (right) planes, demonstrating that singular pinch points are observable via energy integration of inelastic neutron scattering data. Intensity units are arbitrary.

within the zero-moment Coulomb manifold. The upper bands were found to be strongly mixed, with a fourfold pinch point in their HHD structure at the band touching point with one of the two lower bands, indicating that the dynamical modes exhibit singular correlations. We find that the inelastic neutron scattering intensity from these bands exhibits complementary pinch points in the lower and upper bands indicating non-trivial singular correlations in the dynamical spin-spin correlation functions. These singularities are not visible in the energy-integrated structure factors.

The obvious next step is to find a way to express the equations of motion in terms of the long-wavelength tensor and vector degrees of freedom and find a way to expose the rank-2 Coulombic correlations directly in the long wavelength limit at the threefold band touching point in the spin wave spectrum. We note that the results here are similar to the dynamical fragmentation scenario [191] in $Nd_2Zr_2O_7$ and related compounds, which has strong experimental support [187, 188, 190, 207, 208]. However our results are much more subtle than those in the simple dipolar-octupolar Hamiltonian, whose anisotropy can be removed by a coordinate transformation [158]. Our result is the first known example of dynamical *rank-2* fragmentation. Its analysis is not so straightforward as the complete decoupling of rotational and irrotational modes, with the spinwave bands showing strong mixing of rotational and irrotational content.

One additional aspect which sets our results apart from the dipolar-octupolar case is that the two rotational bands in the spectrum are not degenerate. In fact the two rotational bands touch at the X point and there are additional singularities associated with this band touching point which apparently have no precedent in our understanding of the physics of Coulomb phases, because in practically all examples of Coulombic phases studied until now (e.g. nearest-neighbor spin ice and the vorticity model) the rotational bands are degenerate. We saw that in dipolar spin ice there is a very weak lifting of this degeneracy, and it would be good to develop an understanding of what perturbations are capable of lifting this degeneracy and why they are degenerate to begin with. Furthermore, of course it would be worthwhile to develop a full Hodge theoretic understanding of anisotropic Heisenberg spin models on the pyrochlore lattice, in order to have a full microscopic theory of the rank-2 spin liquid phases analogous to the microscopic description of spin ice that we developed in Chapter 3.

Chapter 7

Conclusion and Synthesis

The primary thrust of this thesis has been the introduction and utilization of a theoretical methodology for the study geometrically frustrated magnets on line graph lattices built on a discrete version of Hodge theory, the details of which were laid out in Chapter 3. They key fact that is exploited is that the geometry of these lattices allows us to treat the spin degrees of freedom as a 1-form field, or as a 2-form field on the dual lattice, by identifying the parent lattice of the line graph as a cellular complex. This perspective is known to some extent in the literature through the mapping of quantum spin ice to a lattice gauge theory [25], and what we have done is to carefully formulate the corresponding classical models as a lattice field theory, which is an exact description within the classical soft spin approximations described in Chapter 2. Mathematically, the theory is constructed by defining the appropriate notion of exterior derivative on a cellular complex, which we did by utilizing Stoke's theorem as its definition, thus obtaining the lattice analogues of divergence, gradient, and curl. The most important insight of the Hodge-theoretic approach is that it emphasizes the importance of the Helmholtz-Hodge decomposition (HHD) in understanding the *microscopic* physics of these models, which until now has been appreciated primarily in the restricted context of spin ice fragmentation [133, 134]. In particular, it makes clear the way that the spin degrees of freedom become fractionalized into rotational, irrotational, and harmonic components. Furthermore, we established that the flat bands appearing in the adjacency matrix of line graph lattices correspond to the rotational degrees of freedom, the dispersive bands to irrotational degrees of freedom, and the band touching point to the harmonic degrees of freedom.

In Chapter 3, the primary application of the Hodge theory was in the analysis of the physics of dipolar spin ice. The dominant approach to describing the physics of dipolar spin ice is through the dumbbell model [119], by treating the magnetic moments of the ions not as ideal dipoles but as a pair of magnetic monopoles. We demonstrated that the dumbbell model is equivalent to the analytical approach dubbed "projective equivalence" [142], which observes that the dipole-dipole

interaction matrix can be approximated by the projector to the irrotational subspace. In particular, it is now clear that the dumbbell approximation only captures the effects of dipole-dipole interaction in the irrotational sector, but leaves the perfectly degenerate rotational flat bands intact. Utilizing the connection, we were able to study the (small) corrections to the dumbbell approximation, via the projective equivalence model, and show that the leading qualitative correction is the lifting of the flat band degeneracy, which is physically driven by the tendency of the dipole interactions to polarize the magnetic moments into small loops. We dubbed this correction the vorticity interaction, as it prefers to maximize the vorticity (circulation) around every minimal plaquette of the lattice.

In Chapter 4, we utilized the insights gained from Chapter 3 to interpret the results of polarized neutron scattering experiment in the context of the HHD, re-contextualizing the results of Ref. [149] in the language of the now fully developed Hodge theory. In particular, we demonstrated that scattering at certain wavevectors reflects either the irrotational or rotational fluctuations, while at other wavevectors there is a combination of both contributions. In particular, we demonstrated that the (hhl) non-spin-flip scattering channel contains contributions from *only* the rotational modes. We further demonstrated that the origin of this fact is that the (hhl) NSF probes a type of loop-loop correlation function which is insensitive to the irrotational sector. In particular, this means that for the nearest-neighbor spin ice, for which the rotational modes are perfectly degenerate, the (hhl) NSF is completely flat. We concluded then that the non-flat (hhl) NSF of dipolar spin ice measured experimentally by Fennel et al. [148] reflects the vorticity correction to the dumbbell model and its tendency to weakly anti-ferromagnetically polarize the alpha-chains of the pyrochlore lattice, yielding increased scattering intensity near the zone boundary and decreased intensity near the zone center. This demonstrates that the Hodge theory approach provides new insights into the interpretation of experimental data.

In Chapter 5 we introduced a new classical Ising model on the pyrochlore model, based on the vorticity interaction term, dubbed the vorticity model. This model is a natural complement to nearest-neighbor spin ice, and hosts a new exotic classical spin liquid phase which we dubbed a 2-Coulomb phase. The ground states of this phase can be described as a "membrane soup", and its excitations are one-dimensional closed strings in the dual lattice, or equivalently, vortex tubes in the primary lattice. We confirmed the existence of this phase through classical Monte Carlo simulations, but also found a surprising and still enigmatic all-in-all-out symmetry breaking phase transition which coexists with the spin liquid phase. We hypothesized that this transition is driven by an order-by-disorder mechanism and confirmed that it is not adiabatically connected to the 3D Ising universality class. In the end, the vorticity model and its novel spin liquid ground state are a testament to the power of the Hodge theory approach in guiding our intuition for the design of new models.

Finally, in Chapter 6 we changed gears somewhat to study anisotropic Heisenberg models on

the pyrochlore lattice, in order to study Coulombic physics in ordered, rather than spin liquid, phases. We laid out an analytical approach to apply an HHD analysis to the spin wave fluctuations about an ordered ground state at low temperature, and demonstrated singularities in the dynamical spin correlations which are in principle observable in inelastic neutron scattering experiments. In particular, our results can be seen as the first observation of dynamical *rank-2* fragmentation, described by the fluctuations of a rank-2 tensor. This demonstrates that the Hodge approach can be used to study the physics of anisotropic spin models in addition to Ising models.

Many paths for new research opportunities lie open, building off of the work presented in this thesis. Going in reverse order, based on the results of Chapter 6 it is clear that a full extension of the Hodge theory to describe anisotropic systems should be possible, which should likely culminate in a definition of a *rank-2 HHD*. From Chapter 5, the most obvious next step is to write down the equivalent of the ring-exchange model for quantum spin ice and construct the corresponding 2-form lattice gauge theory. The results from Chapter 4 would be most readily applied to the interpretation of recent polarized neutron scattering studies of the dipolar-octupolar quantum spin ice candidate $Ce_2Zr_2O_7$ [162], and a detailed analysis of the effects of quantum fluctuations on the rotational and irrotational sector correlations is needed. Finally, the Hodge theory presented in Chapter 3 opens many new possibilities for building interesting models, and may have applications beyond the localized magnetic ion systems that we have studied here. Before we finish, we leave the reader with one intriguing thought. Perhaps the most exciting possibility, in the author's opinion, is the possibility of realizing so-called Kähler-Dirac fermions [209–213] in a spin model, which could in principle be built with two interpenetrating copies of the centered pyrochlore lattice [147], whose degrees of freedom would be described as 0-, 1-, 2-, and 3-form fields. The core idea being that the operator $(\mathrm{d} + \delta)$ serves as a square root of the Laplacian and therefore as a geometric analogue of the Dirac operator. We will end here without further explication, simply to say that it seems clear that the Hodge approach, with its emphasis on *geometry*, serves as a powerful guide for our intuition, and no doubt holds many possibilities for future exploration.

References

[1] J. Baez and J. P. Muniain, *Gauge Fields, Knots and Gravity*, Series on Knots and Everything, Vol. 4 (Worlds Scientific, 1994).

[2] M. Nakahara, *Geometry, Topology and Physics*, 2nd ed. (Taylor & Francis Group, Boca Raton, 2003).

[3] R. Bott and L. W. Tu, *Differential Forms in Algebraic Topology* (Springer New Yok, NY, 1982).

[4] M. F. Atiyah, N. J. Hitchin, and I. M. Singer, "Deformations of instantons," Proceedings of the National Academy of Sciences **74**, 2662–2663 (1977).

[5] E. Witten, "Quantum field theory and the Jones polynomial," Communications in Mathematical Physics **121**, 351–399 (1989).

[6] M. Atiyah, "Topological quantum field theories," Publications Mathématiques de l'Institut des Hautes Études Scientifiques **68**, 175–186 (1988).

[7] M. Atiyah and G. W. Moore, "A shifted view of fundamental physics," Surveys in Differential Geometry **15**, 1–16 (2010).

[8] K. G. Wilson, "Confinement of quarks," Phys. Rev. D **10**, 2445–2459 (1974).

[9] K. G. Wilson, "Quarks and Strings on a Lattice," in *New Phenomena in Subnuclear Physics: Part A*, The Subnuclear Series, edited by A. Zichichi (Springer US, Boston, MA, 1977) pp. 69–142.

[10] M. Creutz, "Monte Carlo study of quantized SU(2) gauge theory," Physical Review D **21**, 2308–2315 (1980).

[11] M. Creutz, *Quarks, Gluons and Lattices* (Cambridge University Press, Cambridge, 1984).

[12] X.-G. Wen, "Topological orders and edge excitations in fractional quantum Hall states," Advances in Physics **44**, 405–473 (1995).

[13] M. A. Levin and X.-G. Wen, "String-net condensation: A physical mechanism for topological phases," Phys. Rev. B **71**, 045110 (2005).

[14] H. C. Po, H. Watanabe, C.-M. Jian, and M. P. Zaletel, "Lattice Homotopy Constraints on Phases of Quantum Magnets," Phys. Rev. Lett. **119**, 127202 (2017).

[15] S. D. Huber, "Topological mechanics," Nature Phys **12**, 621–623 (2016).

[16] P. A. McClarty, "Topological Magnons: A Review," Annual Review of Condensed Matter Physics **13**, 171–190 (2022).

[17] B. Bradlyn, L. Elcoro, J. Cano, M. G. Vergniory, Z. Wang, C. Felser, M. I. Aroyo, and B. A. Bernevig, "Topological quantum chemistry," Nature **547**, 298–305 (2017).

[18] M. Sato and Y. Ando, "Topological superconductors: A review," Rep. Prog. Phys. **80**, 076501 (2017).

[19] D. V. Else and R. Thorngren, "Topological theory of Lieb-Schultz-Mattis theorems in quantum spin systems," Phys. Rev. B **101**, 224437 (2020).

[20] É. Lantagne-Hurtubise and M. Franz, "Topology in abundance," Nature Reviews Physics **1**, 183–184 (2019).

[21] L. Fu, C. L. Kane, and E. J. Mele, "Topological Insulators in Three Dimensions," Phys. Rev. Lett. **98**, 106803 (2007).

[22] X.-G. Wen, "Systematic construction of gapped nonliquid states," Phys. Rev. Research **2**, 033300 (2020).

[23] J. Wang, "Nonliquid cellular states: Gluing gauge-higher-symmetry-breaking versus gauge-higher-symmetry-extension interfacial defects," Phys. Rev. Research **4**, 023258 (2022).

[24] Y. Fuji and A. Furusaki, "Bridging three-dimensional coupled-wire models and cellular topological states: Solvable models for topological and fracton orders," (2021), arXiv:2112.07926 .

[25] M. Hermele, M. P. A. F. Fisher, and L. Balents, "Pyrochlore photons: The U(1) spin liquid in a S=1/2 three-dimensional frustrated magnet," Phys. Rev. B **69**, 064404 (2004).

[26] E. Witten, "Symmetry and emergence," Nat. Phys. **14**, 116–119 (2018).

[27] J. Jensen and a. Mackintosh, *Rare Earth Magnetism: Structures and Excitations* (Oxford University Press, Oxford, 1991).

[28] P. Fazekas, *Electron Correlation and Magnetism*, Series in Condensed Matter Physics, Vol. 5 (World Scientific, Singapore, 1999).

[29] R. M. White, *Quantum Theory of Magnetism*, 3rd ed., Springer Series in Solid-State Sciences, Vol. 32 (Springer Berlin, Heidelberg, 1970).

[30] A. Bertin, Y. Chapuis, P. D. de Réotier, and A. Yaouanc, "Crystal electric field in the R2Ti22O7 pyrochlore compounds," J. Phys.: Condens. Matter **24**, 256003 (2012).

[31] H. D. Zhou, C. R. Wiebe, J. A. Janik, L. Balicas, Y. J. Yo, Y. Qiu, J. R. D. Copley, and J. S. Gardner, "Dynamic Spin Ice: Pr2Sn2O7," Phys. Rev. Lett. **101**, 227204 (2008).

[32] M. C. Hatnean, M. R. Lees, O. A. Petrenko, D. S. Keeble, G. Balakrishnan, M. J. Gutmann, V. V. Klekovkina, and B. Z. Malkin, "Structural and magnetic investigations of single-crystalline neodymium zirconate pyrochlore Nd2Zr2O7," Phys. Rev. B **91**, 174416 (2015).

[33] R. Sibille, E. Lhotel, V. Pomjakushin, C. Baines, T. Fennell, and M. Kenzelmann, "Candidate Quantum Spin Liquid in the Ce3+ Pyrochlore Stannate Ce2Sn2O7," Phys. Rev. Lett. **115**, 097202 (2015).

[34] S. Onoda and Y. Tanaka, "Quantum fluctuations in the effective pseudospin-1/2 model for magnetic pyrochlore oxides," Phys. Rev. B **83**, 094411 (2011).

[35] J. G. Rau and M. J. P. Gingras, "Magnitude of quantum effects in classical spin ices," Phys. Rev. B **92**, 144417 (2015).

[36] J. G. Rau and M. J. Gingras, "Frustrated Quantum Rare-Earth Pyrochlores," Annual Review of Condensed Matter Physics **10**, 357–386 (2019).

[37] L. Landau, "The Theory of Phase Transitions," Nature **138**, 840–841 (1936).

[38] P. M. Chaikin and T. C. Lubensky, *Principles of Condensed Matter Physics* (Cambridge University Press, Cambridge, 1995).

[39] T. W. B. Kibble, "Topology of cosmic domains and strings," J. Phys. A: Math. Gen. **9**, 1387–1398 (1976).

[40] G. Toulouse, "The frustration model," Modern Trends in the Theory of Condensed Matter , 195–203 (1980).

[41] G. Toulouse, "How 'Frustration' Set In," Physics Today **42**, 97–97 (1989).

[42] G. Toulouse, J. Vannimenus, and J. Maillard, "Spin glasses and roughening transition," J. Phyique Lett. **38**, 459–461 (1977).

[43] R. Moessner and A. P. Ramirez, "Geometrical frustration," Physics Today **59**, 24–29 (2006).

[44] J. Knolle and R. Moessner, "A field guide to spin liquids," Annual Review of Condensed Matter Physics **10**, 451–472 (2019).

[45] Y. Zhou, K. Kanoda, and T.-K. Ng, "Quantum spin liquid states," Rev. Mod. Phys. **89**, 025003 (2017).

[46] L. Savary and L. Balents, "Quantum spin liquids: A review," Reports on Progress in Physics **80** (2017), 10.1088/0034-4885/80/1/016502.

[47] M. Oshikawa and T. Senthil, "Fractionalization, topological order, and quasiparticle statistics," Physical Review Letters **96**, 060601 (2006).

[48] B. J. Powell, "Emergent particles and gauge fields in quantum matter," Contemporary Physics **61**, 96–131 (2020).

[49] T. T. Wu and C. N. Yang, "Concept of nonintegrable phase factors and global formulation of gauge fields," Phys. Rev. D **12**, 3845–3857 (1975).

[50] S. T. Bramwell and M. J. Harris, "Frustration in Ising-type spin models on the pyrochlore lattice," Journal of Physics Condensed Matter **10** (1998), 10.1088/0953-8984/10/14/002.

[51] M. J. Harris, S. T. Bramwell, D. F. McMorrow, T. Zeiske, and K. W. Godfrey, "Geometrical Frustration in the Ferromagnetic Pyrochlore Ho2Ti2O7," Phys. Rev. Lett. **79**, 2554–2557 (1997).

[52] M. J. P. Gingras and B. C. Den Hertog, "Origin of spin-ice behavior in Ising pyrochlore magnets with long-range dipole interactions: An insight from mean-field theory," Canadian Journal of Physics **79**, 1339–1351 (2001).

[53] S. T. Bramwell and M. J. Harris, "The history of spin ice," Journal of Physics: Condensed Matter **32**, 374010 (2020).

[54] M. J. P. Gingras, P. A. McClarty, and J. G. Rau, "Spin Ice: Microscopic Physics," in *Spin Ice*, Springer Series in Solid-State Sciences, edited by M. Udagawa and L. Jaubert (Springer International Publishing, Cham, 2021) pp. 1–18.

[55] C. L. Henley, "The "Coulomb Phase" in Frustrated Systems," Annual Review of Condensed Matter Physics **1**, 179–210 (2010).

[56] R. Moessner, "Spin Ice As a Coulomb Liquid: From Emergent Gauge Fields to Magnetic Monopoles," in *Spin Ice*, Springer Series in Solid-State Sciences, edited by M. Udagawa and L. Jaubert (Springer International Publishing, Cham, 2021) pp. 37–70.

[57] T. Mizoguchi, L. D. C. Jaubert, R. Moessner, and M. Udagawa, "Magnetic clustering, half-moons, and shadow pinch points as signals of a proximate Coulomb phase in frustrated Heisenberg magnets," Phys. Rev. B **98**, 144446 (2018).

[58] C. Castelnovo, R. Moessner, and S. Sondhi, "Spin Ice, Fractionalization, and Topological Order," Annual Review of Condensed Matter Physics **3**, 35–55 (2012).

[59] T. Mizoguchi, L. D. C. Jaubert, and M. Udagawa, "Clustering of Topological Charges in a Kagome Classical Spin Liquid," Phys. Rev. Lett. **119**, 077207 (2017).

[60] O. Benton, O. Sikora, and N. Shannon, "Seeing the light: Experimental signatures of emergent electromagnetism in a quantum spin ice," Phys. Rev. B **86**, 075154 (2012).

[61] S. D. Pace, S. C. Morampudi, R. Moessner, and C. R. Laumann, "Emergent Fine Structure Constant of Quantum Spin Ice is Large," Phys. Rev. Lett. **127**, 117205 (2021).

[62] L. D. C. Jaubert, "Topology of the Vacuum," in *Spin Ice*, Springer Series in Solid-State Sciences, edited by M. Udagawa and L. Jaubert (Springer International Publishing, Cham, 2021) pp. 117–141.

[63] C. Nisoli, "Topological order of the Rys F-model and its breakdown in realistic square spin ice: Topological sectors of Faraday loops," EPL **132**, 47005 (2020).

[64] A. Y. Kitaev, "Fault-tolerant quantum computation by anyons," Annals of Physics **303**, 2–30 (2003).

[65] K. Slagle and Y. B. Kim, "X-cube model on generic lattices: Fracton phases and geometric order," Phys. Rev. B **97**, 165106 (2018).

[66] X.-G. Wen, "Topological Order: From Long-Range Entangled Quantum Matter to a Unified Origin of Light and Electrons," ISRN Condensed Matter Physics **2013**, 1–20 (2013).

[67] R. M. Nandkishore and M. Hermele, "Fractons," Annual Review of Condensed Matter Physics **10**, 295–313 (2019).

[68] É. Lantagne-Hurtubise, J. G. Rau, and M. J. P. Gingras, "Spin-Ice Thin Films: Large-N Theory and Monte Carlo Simulations," Phys. Rev. X **8**, 021053 (2018).

[69] T. H. Berlin and M. Kac, "The Spherical Model of a Ferromagnet," Phys. Rev. **86**, 821–835 (1952).

[70] M. Kac, "The work of T.H. Berlin in statistical mechanics . . . a personal reminiscence," Physics Today **17**, 40–42 (1964).

[71] H. E. Stanley, "Spherical model as the limit of infinite spin dimensionality," Physical Review **176**, 718–722 (1968).

[72] P. A. Pearce and C. J. Thompson, "The spherical limit for n-vector correlations," J Stat Phys **17**, 189–196 (1977).

[73] O. Benton, "Analytical Approaches to Quantum Spin Ice," in *Spin Ice*, Springer Series in Solid-State Sciences, edited by M. Udagawa and L. Jaubert (Springer International Publishing, Cham, 2021) pp. 303–323.

[74] G. Alexanian, R. MacKenzie, M. B. Paranjape, and J. Ruel, "Problems with complex actions," Can. J. Phys. **85**, 699–705 (2007).

[75] G. Alexanian, R. MacKenzie, M. B. Paranjape, and J. Ruel, "Path integration and perturbation theory with complex Euclidean actions," Phys. Rev. D **77**, 105014 (2008).

[76] P. H. Conlon and J. T. Chalker, "Absent pinch points and emergent clusters: Further neighbor interactions in the pyrochlore Heisenberg antiferromagnet," Phys. Rev. B **81**, 224413 (2010).

[77] G. H. Wannier, "Antiferromagnetism. The Triangular Ising Net," Phys. Rev. **79**, 357–364 (1950).

[78] O. Götze, J. Richter, R. Zinke, and D. J. J. Farnell, "Ground-state properties of the triangular-lattice Heisenberg antiferromagnet with arbitrary spin quantum number s," Journal of Magnetism and Magnetic Materials **397**, 333–341 (2016).

[79] L. Balents, M. P. A. Fisher, and S. M. Girvin, "Fractionalization in an easy-axis Kagome antiferromagnet," Phys. Rev. B **65**, 224412 (2002).

[80] L. Beineke and J. Bagga, *Line Graphs and Line Digraphs*, 1st ed., Developments in Mathematics, Vol. 68 (Springer Cham, 2021).

[81] P. H. Conlon, *Aspects of Frustrated Magnetism*, Ph.D. thesis, University of Oxford, Oxford (2010).

[82] G.-W. Chern, "Artificial Spin Ice: Beyond Pyrochlores and Magnetism," in *Spin Ice*, Springer Series in Solid-State Sciences, edited by M. Udagawa and L. Jaubert (Springer International Publishing, Cham, 2021) pp. 419–453.

[83] C. Nisoli, R. Moessner, and P. Schiffer, "Colloquium: Artificial spin ice: Designing and imaging magnetic frustration," Rev. Mod. Phys. **85**, 1473–1490 (2013).

[84] A. Farhan, M. Saccone, C. F. Petersen, S. Dhuey, R. V. Chopdekar, Y. L. Huang, N. Kent, Z. Chen, M. J. Alava, T. Lippert, A. Scholl, and S. van Dijken, "Emergent magnetic monopole dynamics in macroscopically degenerate artificial spin ice," Science Advances **5**, 1–9 (2019).

[85] A. Ortiz-Ambriz and P. Tierno, "Engineering of frustration in colloidal artificial ices realized on microfeatured grooved lattices," Nat Commun **7**, 10575 (2016).

[86] D. Levis, *Two-Dimensional Spin Ice and the Sixteen-Vertex Model*, Ph.D. thesis, Université Pierre et Marie Curie, Paris (2012).

[87] L. Foini, D. Levis, M. Tarzia, and L. F. Cugliandolo, "Static properties of 2D spin-ice as a sixteen-vertex model," J. Stat. Mech. , P02026 (2013).

[88] R. W. Youngblood and J. D. Axe, "Polarization fluctuations in ferroelectric models," Phys. Rev. B **23**, 232–238 (1981).

[89] R. Youngblood, J. D. Axe, and B. M. McCoy, "Correlations in ice-rule ferroelectrics," Phys. Rev. B **21**, 5212–5220 (1980).

[90] R. J. Baxter, *Exactly Solved Model in Statistical Mechanics* (Academic Press, London, 1982).

[91] R. J. Baxter, "Partition function of the eight-vertex lattice model," Annals of Physics **281**, 187–222 (2000).

[92] E. Witten, "Gauge theories, vertex models, and quantum groups," Nuclear Physics B **330**, 285–346 (1990).

[93] I. K. Kostov, "Exact solution of the six-vertex model on a random lattice," Nuclear Physics B **575**, 513–534 (2000).

[94] J. Rehn, A. Sen, and R. Moessner, "Fractionalized Z2 Classical Heisenberg Spin Liquids," Phys. Rev. Lett. **118**, 047201 (2017).

[95] Y. Okamoto, M. Nohara, H. Aruga-Katori, and H. Takagi, "Spin-Liquid State in the S=1/2 Hyperkagome Antiferromagnet Na4Ir3O8," Phys. Rev. Lett. **99**, 137207 (2007).

[96] J. M. Hopkinson, S. V. Isakov, H.-Y. Kee, and Y. B. Kim, "Classical Antiferromagnet on a Hyperkagome Lattice," Phys. Rev. Lett. **99**, 037201 (2007).

[97] J. A. Quilliam, S. Meng, H. A. Craig, L. R. Corruccini, G. Balakrishnan, O. A. Petrenko, A. Gomez, S. W. Kycia, M. J. P. Gingras, and J. B. Kycia, "Juxtaposition of spin freezing and long range order in a series of geometrically frustrated antiferromagnetic gadolinium garnets," Phys. Rev. B **87**, 174421 (2013).

[98] H.-K. Jin and Y. Zhou, "Classical and quantum order in hyperkagome antiferromagnets," Phys. Rev. B **101**, 054408 (2020).

[99] H. Jacobsen, O. Florea, E. Lhotel, K. Lefmann, O. A. Petrenko, C. S. Knee, T. Seydel, P. F. Henry, R. Bewley, D. Voneshen, A. Wildes, G. Nilsen, and P. P. Deen, "Spin dynamics of the director state in frustrated hyperkagome systems," Phys. Rev. B **104**, 054440 (2021).

[100] S. Chillal, Y. Iqbal, H. O. Jeschke, J. A. Rodriguez-Rivera, R. Bewley, P. Manuel, D. Khalyavin, P. Steffens, R. Thomale, A. T. M. N. Islam, J. Reuther, and B. Lake, "Evidence for a three-dimensional quantum spin liquid in PbCuTe2O6," Nat Commun **11**, 2348 (2020).

[101] L. E. Chern and Y. B. Kim, "Theoretical study of quantum spin liquids in S-1/2 hyper-hyperkagome magnets: Classification, heat capacity, and dynamical spin structure factor," Phys. Rev. B **104**, 094413 (2021).

[102] J. M. Hopkinson and H.-Y. Kee, "Geometric frustration inherent to the trillium lattice, a sublattice of the B20 structure," Phys. Rev. B **74**, 224441 (2006).

[103] D. L. Bergman, C. Wu, and L. Balents, "Band touching from real-space topology in frustrated hopping models," Phys. Rev. B **78**, 125104 (2008).

[104] T. Mizoguchi and M. Udagawa, "Flat-band engineering in tight-binding models: Beyond the nearest-neighbor hopping," Phys. Rev. B **99**, 235118 (2019).

[105] A. Mielke, "Ferromagnetism in the Hubbard model on line graphs and further considerations," J. Phys. A: Math. Gen. **24**, 3311–3321 (1991).

[106] A. Mielke, "Exact ground states for the Hubbard model on the Kagome lattice," J. Phys. A: Math. Gen. **25**, 4335–4345 (1992).

[107] C. S. Chiu, D.-S. Ma, Z.-D. Song, B. A. Bernevig, and A. A. Houck, "Fragile topology in line-graph lattices with two, three, or four gapped flat bands," Phys. Rev. Research **2**, 043414 (2020).

[108] D.-S. Ma, Y. Xu, C. S. Chiu, N. Regnault, A. A. Houck, Z. Song, and B. A. Bernevig, "Spin-Orbit-Induced Topological Flat Bands in Line and Split Graphs of Bipartite Lattices," Phys. Rev. Lett. **125**, 266403 (2020).

[109] S. Lisi, X. Lu, T. Benschop, T. A. de Jong, P. Stepanov, J. R. Duran, F. Margot, I. Cucchi, E. Cappelli, A. Hunter, A. Tamai, V. Kandyba, A. Giampietri, A. Barinov, J. Jobst, V. Stalman, M. Leeuwenhoek, K. Watanabe, T. Taniguchi, L. Rademaker, S. J. van der Molen, M. P. Allan, D. K. Efetov, and F. Baumberger, "Observation of flat bands in twisted bilayer graphene," Nat. Phys. **17**, 189–193 (2021).

[110] Z. Liu, F. Liu, and Y.-S. Wu, "Exotic electronic states in the world of flat bands: From theory to material," Chinese Phys. B **23**, 077308 (2014).

[111] C. Danieli, A. Andreanov, and S. Flach, "Many-body localization transition from flat-band fine tuning," Phys. Rev. B **105**, L041113 (2022).

[112] J.-W. Rhim and B.-J. Yang, "Classification of flat bands according to the band-crossing singularity of Bloch wave functions," Phys. Rev. B **99**, 045107 (2019).

[113] H. Yang, H. Nakano, and H. Katsura, "Symmetry-protected topological phases in spinful bosons with a flat band," Phys. Rev. Research **3**, 023210 (2021).

[114] F. C. de Lima, G. J. Ferreira, and R. H. Miwa, "Topological flat band, Dirac fermions and quantum spin Hall phase in 2D Archimedean lattices," Phys. Chem. Chem. Phys. **21**, 22344–22350 (2019).

[115] J. Cano and B. Bradlyn, "Band Representations and Topological Quantum Chemistry," Annu. Rev. Condens. Matter Phys. **12**, 225–246 (2021).

[116] N. Regnault, Y. Xu, M.-R. Li, D.-S. Ma, M. Jovanovic, A. Yazdani, S. S. P. Parkin, C. Felser, L. M. Schoop, N. P. Ong, R. J. Cava, L. Elcoro, Z.-D. Song, and B. A. Bernevig, "Catalogue of flat-band stoichiometric materials," Nature **603**, 824–828 (2022).

[117] H. Wondratschek and U. Müller, *International Tables for Crystallography - A*, edited by T. Hahn, Vol. A (International Union of Crystallography, Chester, England, 2006).

[118] A. Zee, *Group Theory in a Nutshell for Physicists* (Princeton University Press, Princeton, 2016).

[119] C. Castelnovo, R. Moessner, and S. L. Sondhi, "Magnetic monopoles in spin ice," Nature **451**, 42–45 (2008).

[120] O. Benton, L. D. C. Jaubert, H. Yan, and N. Shannon, "A spin-liquid with pinch-line singularities on the pyrochlore lattice," Nature Communications **7**, 1–7 (2016).

[121] S. W. Lovesey, "Theory of neutron scattering by electrons in magnetic materials," Phys. Scr. **90**, 108011 (2015).

[122] B. Roessli and P. Böni, "Polarized Neutron Scattering," (2000), arXiv:cond-mat/0012180 .

[123] B. Dalla Piazza, M. Mourigal, N. B. Christensen, G. J. Nilsen, P. Tregenna-Piggott, T. G. Perring, M. Enderle, D. F. McMorrow, D. A. Ivanov, and H. M. Rønnow, "Fractional excitations in the square-lattice quantum antiferromagnet," Nature Phys **11**, 62–68 (2015).

[124] G. T. Barkema and M. E. J. Newman, "Monte Carlo simulation of ice models," Phys. Rev. E **57**, 1155–1166 (1998).

[125] R. G. Melko and M. J. P. Gingras, "Monte Carlo studies of the dipolar spin ice model," J. Phys.: Condens. Matter **16**, R1277–R1319 (2004).

[126] R. G. Melko, *Low Temperature Physics of Dipolar Spin Ice : A Monte Carlo Study*, Ph.D. thesis, University of Waterloo (2001).

[127] H. G. Evertz, "The loop algorithm," Advances in Physics **52**, 1–66 (2003).

[128] H. Otsuka, "Cluster algorithm for Monte Carlo simulations of spin ice," Phys. Rev. B **90**, 220406 (2014).

[129] J. D. Alzate-Cardona, D. Sabogal-Suárez, R. F. L. Evans, and E. Restrepo-Parra, "Optimal phase space sampling for Monte Carlo simulations of Heisenberg spin systems," J. Phys.: Condens. Matter **31**, 095802 (2019).

[130] L. Savary and L. Balents, "Coulombic Quantum Liquids in Spin-1/2 Pyrochlores," Phys. Rev. Lett. **108**, 037202 (2012).

[131] M. Kardar, *Statistical Physics of Fields* (Cambridge University Press, Cambridge, 2007).

[132] I. Rousochatzakis, A. M. Läuchli, and F. Mila, "Highly frustrated magnetic clusters: The kagom\'e on a sphere," Phys. Rev. B **77**, 094420 (2008).

[133] M. E. Brooks-Bartlett, S. T. Banks, L. D. C. Jaubert, A. Harman-Clarke, and P. C. W. Holdsworth, "Magnetic-Moment Fragmentation and Monopole Crystallization," Phys. Rev. X **4**, 011007 (2014).

[134] E. Lhotel, L. D. C. Jaubert, and P. C. Holdsworth, "Fragmentation in Frustrated Magnets: A Review," Journal of Low Temperature Physics **201**, 710–737 (2020).

[135] D. Slobinsky, L. Pili, and R. A. Borzi, "Polarized monopole liquid: A Coulomb phase in a fluid of magnetic charges," Phys. Rev. B **100**, 020405 (2019).

[136] C. L. Henley, "Power-law spin correlations in pyrochlore antiferromagnets," Phys. Rev. B **71**, 014424 (2005).

[137] O. Benton and R. Moessner, "Topological Route to New and Unusual Coulomb Spin Liquids," Phys. Rev. Lett. **127**, 107202 (2021).

[138] G. Moore, "A Minicourse on Generalized Abelian Gauge Theory, Self-Dual Theories, and Differential Cohomology," (2011).

[139] J. G. Rau and M. J. P. Gingras, "Spin slush in an extended spin ice model," Nat Commun **7**, 12234 (2016).

[140] D. Kiese, F. Ferrari, N. Astrakhantsev, N. Niggemann, P. Ghosh, T. Müller, R. Thomale, T. Neupert, J. Reuther, M. J. P. Gingras, S. Trebst, and Y. Iqbal, "Pinch-points to half-moons and up in the stars: The kagome skymap," (2022), arXiv:2206.00264 .

[141] B. C. den Hertog and M. J. P. Gingras, "Dipolar Interactions and Origin of Spin Ice in Ising Pyrochlore Magnets," Phys. Rev. Lett. **84**, 3430–3433 (2000).

[142] S. V. Isakov, R. Moessner, and S. L. Sondhi, "Why Spin Ice Obeys the Ice Rules," Phys. Rev. Lett. **95**, 217201 (2005).

[143] L. D. C. Jaubert, M. Haque, and R. Moessner, "Analysis of a Fully Packed Loop Model Arising in a Magnetic Coulomb Phase," Phys. Rev. Lett. **107**, 177202 (2011).

[144] R. G. Melko, B. C. den Hertog, and M. J. P. Gingras, "Long-Range Order at Low Temperatures in Dipolar Spin Ice," Phys. Rev. Lett. **87**, 067203 (2001).

[145] L. D. C. Jaubert and P. C. W. Holdsworth, "Magnetic monopole dynamics in spin ice," J. Phys.: Condens. Matter **23**, 164222 (2011).

[146] A. M. Samarakoon, S. A. Grigera, D. A. Tennant, A. Kirste, B. Klemke, P. Strehlow, M. Meissner, J. N. Hallén, L. Jaubert, C. Castelnovo, and R. Moessner, "Anomalous magnetic noise in an imperfectly flat landscape in the topological magnet Dy2Ti2O7," PNAS **119**, e2117453119 (2022).

[147] R. Nutakki, R. Röß-Ohlenroth, D. Volkmer, A. Jesche, H.-A. K. von Nidda, A. A. Tsirlin, P. Gegenwart, L. Pollet, and L. D. C. Jaubert, "Frustration on a centred pyrochlore lattice in metal-organic frameworks," (2022), arXiv:2203.08780 .

[148] T. Fennell, P. P. Deen, A. R. Wildes, K. Schmalzl, D. Prabhakaran, A. T. Boothroyd, R. J. Aldus, D. F. McMorrow, and S. T. Bramwell, "Magnetic Coulomb Phase in the Spin Ice Ho2Ti2O7," Science **326**, 415–417 (2009).

[149] K. T. K. Chung, J. S. K. Goh, A. Mukherjee, W. Jin, D. Lozano-Gómez, and M. J. P. Gingras, "Probing Flat Band Physics in Spin Ice Systems via Polarized Neutron Scattering," Phys. Rev. Lett. **128**, 107201 (2022).

[150] Y. Kato and S. Onoda, "Numerical Evidence of Quantum Melting of Spin Ice: Quantum-to-Classical Crossover," Phys. Rev. Lett. **115**, 077202 (2015).

[151] M. Taillefumier, O. Benton, H. Yan, L. D. C. Jaubert, and N. Shannon, "Competing Spin Liquids and Hidden Spin-Nematic Order in Spin Ice with Frustrated Transverse Exchange," Phys. Rev. X **7**, 041057 (2017).

[152] C. Castelnovo and R. Moessner, "Rod motifs in neutron scattering in spin ice," Phys. Rev. B **99**, 121102 (2019).

[153] J. P. Clancy, J. P. C. Ruff, S. R. Dunsiger, Y. Zhao, H. A. Dabkowska, J. S. Gardner, Y. Qiu, J. R. D. Copley, T. Jenkins, and B. D. Gaulin, "Revisiting static and dynamic spin-ice correlations in Ho2Ti2O7 with neutron scattering," Phys. Rev. B **79**, 014408 (2009).

[154] J. P. C. Ruff, R. G. Melko, and M. J. P. Gingras, "Finite-Temperature Transitions in Dipolar Spin Ice in a Large Magnetic Field," Phys. Rev. Lett. **95**, 097202 (2005).

[155] Z. Hiroi, K. Matsuhira, and M. Ogata, "Ferromagnetic Ising Spin Chains Emerging from the Spin Ice under Magnetic Field," J. Phys. Soc. Jpn. **72**, 3045–3048 (2003).

[156] P. A. McClarty, O. Sikora, R. Moessner, K. Penc, F. Pollmann, and N. Shannon, "Chain-based order and quantum spin liquids in dipolar spin ice," Phys. Rev. B **92**, 094418 (2015).

[157] D. A. Huse, W. Krauth, R. Moessner, and S. L. Sondhi, "Coulomb and Liquid Dimer Models in Three Dimensions," Phys. Rev. Lett. **91**, 167004 (2003).

[158] Y.-P. Huang, G. Chen, and M. Hermele, "Quantum Spin Ices and Topological Phases from Dipolar-Octupolar Doublets on the Pyrochlore Lattice," Phys. Rev. Lett. **112**, 167203 (2014).

[159] B. Gao, T. Chen, D. W. Tam, C.-L. Huang, K. Sasmal, D. T. Adroja, F. Ye, H. Cao, G. Sala, M. B. Stone, C. Baines, J. A. T. Verezhak, H. Hu, J.-H. Chung, X. Xu, S.-W. Cheong, M. Nallaiyan, S. Spagna, M. B. Maple, A. H. Nevidomskyy, E. Morosan, G. Chen, and P. Dai, "Experimental signatures of a three-dimensional quantum spin liquid in effective spin-1/2 Ce2Zr2O7 pyrochlore," Nat. Phys. **15**, 1052–1057 (2019).

[160] J. Gaudet, E. M. Smith, J. Dudemaine, J. Beare, C. R. C. Buhariwalla, N. P. Butch, M. B. Stone, A. I. Kolesnikov, G. Xu, D. R. Yahne, K. A. Ross, C. A. Marjerrison, J. D. Garrett, G. M. Luke, A. D. Bianchi, and B. D. Gaulin, "Quantum Spin Ice Dynamics in the Dipole-Octupole Pyrochlore Magnet Ce2Zr2O7," Phys. Rev. Lett. **122**, 187201 (2019).

[161] A. Bhardwaj, S. Zhang, H. Yan, R. Moessner, A. H. Nevidomskyy, and H. J. Changlani, "Sleuthing out exotic quantum spin liquidity in the pyrochlore magnet Ce2Zr2O7," npj Quantum Mater. **7**, 1–8 (2022).

[162] E. M. Smith, O. Benton, D. R. Yahne, B. Placke, R. Schäfer, J. Gaudet, J. Dudemaine, A. Fitterman, J. Beare, A. R. Wildes, S. Bhattacharya, T. DeLazzer, C. R. C. Buhariwalla, N. P. Butch, R. Movshovich, J. D. Garrett, C. A. Marjerrison, J. P. Clancy, E. Kermarrec, G. M. Luke, A. D. Bianchi, K. A. Ross, and B. D. Gaulin, "Case for a U(1)-pi Quantum Spin Liquid Ground State in the Dipole-Octupole Pyrochlore Ce2Zr2O7," Phys. Rev. X **12**, 021015 (2022).

[163] M. Hosoi, E. Z. Zhang, A. S. Patri, and Y. B. Kim, "Uncovering Footprints of Dipolar-Octupolar Quantum Spin Ice from Neutron Scattering Signatures," Phys. Rev. Lett. **129**, 097202 (2022).

[164] R. Aldus, *Neutron Scattering Studies of Frustrated Magnets*, Ph.D. thesis, University College London, London (2010).

[165] R. J. Aldus, T. Fennell, P. P. Deen, E. Ressouche, G. C. Lau, R. J. Cava, and S. T. Bramwell, "Ice rule correlations in stuffed spin ice," New J. Phys. **15**, 013022 (2013).

[166] L. J. Chang, Y. Su, H. Schneider, R. Mittal, M. R. Lees, G. Balakrishnan, and T. Brückel, "Magnetic correlations in pyrochlore spin ice as probed by polarized neutron scattering," in *Journal of Physics: Conference Series*, Vol. 211 (2010).

[167] J. S. Gardner, G. Ehlers, A. Faraone, and V. García Sakai, "High-resolution neutron spectroscopy using backscattering and neutron spin-echo spectrometers in soft and hard condensed matter," Nature Reviews Physics **2**, 103–116 (2020).

[168] S. Powell, "Higgs transitions of spin ice," Phys. Rev. B **84**, 094437 (2011).

[169] L. Pauling, "The Structure and Entropy of Ice and of Other Crystals with Some Randomness of Atomic Arrangement," J. Am. Chem. Soc. **57**, 2680–2684 (1935).

[170] R. R. P. Singh and J. Oitmaa, "Corrections to Pauling residual entropy and single tetrahedron based approximations for the pyrochlore lattice Ising antiferromagnet," Phys. Rev. B **85**, 144414 (2012).

[171] S. Coleman and E. Weinberg, "Radiative Corrections as the Origin of Spontaneous Symmetry Breaking," Phys. Rev. D **7**, 1888–1910 (1973).

[172] K. Gregor, D. A. Huse, R. Moessner, and S. L. Sondhi, "Diagnosing deconfinement and topological order," New J. Phys. **13**, 025009 (2011).

[173] C. Lacroix, P. Mendels, and F. Mila, eds., *Introduction to Frustrated Magnetism*, 1st ed., Springer Series in Solid-State Sciences, Vol. 164 (Springer-Verlag, Heidelberg, 2011).

[174] C. L. Henley, "Ordering due to disorder in a frustrated vector antiferromagnet," Phys. Rev. Lett. **62**, 2056–2059 (1989).

[175] J. Villain, R. Bidaux, J.-P. Carton, and R. Conte, "Order as an effect of disorder," J. Phys. France **41**, 1263–1272 (1980).

[176] K. Binder, "Finite size scaling analysis of ising model block distribution functions," Z. Physik B - Condensed Matter **43**, 119–140 (1981).

[177] A. W. Sandvik, "Computational studies of quantum spin systems," in *AIP Conference Proceedings*, Vol. 1297 (2010) pp. 135–338.

[178] N. Iqbal and J. McGreevy, "Mean string field theory: Landau-Ginzburg theory for 1-form symmetries," (2021), arXiv:2106.12610 .

[179] D. Hofman and N. Iqbal, "Goldstone modes and photonization for higher form symmetries," SciPost Physics **6**, 006 (2019).

[180] E. Lake, "Higher-form symmetries and spontaneous symmetry breaking," (2018), arXiv:1802.07747 .

[181] J. A. Damia, R. Argurio, and L. Tizzano, "Continuous Generalized Symmetries in Three Dimensions," (2022), arXiv:2206.14093 .

[182] N. Iqbal and K. Macfarlane, "Higher-form symmetry breaking and holographic flavour," (2021), arXiv:2107.00373 .

[183] A. Das, R. Gregory, and N. Iqbal, "Higher-form symmetries, anomalous magnetohydrodynamics, and holography," (2022), arXiv:2205.03619 .

[184] D. Delmastro, J. Gomis, P.-S. Hsin, and Z. Komargodski, "Anomalies and Symmetry Fractionalization," (2022), arXiv:2206.15118 .

[185] H. Yan, O. Benton, L. D. C. Jaubert, and N. Shannon, "Rank-2 U(1) Spin Liquid on the Breathing Pyrochlore Lattice," Phys. Rev. Lett. **124**, 127203 (2020).

[186] L.-J. Chang, S. Onoda, Y. Su, Y.-J. Kao, K.-D. Tsuei, Y. Yasui, K. Kakurai, and M. R. Lees, "Higgs transition from a magnetic Coulomb liquid to a ferromagnet in Yb2Ti2O7," Nat Commun **3**, 992 (2012).

[187] J. Xu, O. Benton, A. T. M. N. Islam, T. Guidi, G. Ehlers, and B. Lake, "Order out of a Coulomb Phase and Higgs Transition: Frustrated Transverse Interactions of Nd2Zr2O7," Phys. Rev. Lett. **124**, 097203 (2020).

[188] M. Léger, E. Lhotel, M. Ciomaga Hatnean, J. Ollivier, A. R. Wildes, S. Raymond, E. Ressouche, G. Balakrishnan, and S. Petit, "Spin Dynamics and Unconventional Coulomb Phase in Nd 2 Zr 2 O 7," Phys. Rev. Lett. **126**, 247201 (2021).

[189] A. Sadeghi, M. Alaei, F. Shahbazi, and M. J. P. Gingras, "Spin Hamiltonian, order out of a Coulomb phase, and pseudocriticality in the frustrated pyrochlore Heisenberg antiferromagnet FeF3," Phys. Rev. B **91**, 140407 (2015).

[190] S. Petit, E. Lhotel, B. Canals, M. Ciomaga Hatnean, J. Ollivier, H. Mutka, E. Ressouche, A. R. Wildes, M. R. Lees, and G. Balakrishnan, "Observation of magnetic fragmentation in spin ice," Nature Physics **12**, 746–750 (2016).

[191] O. Benton, "Quantum origins of moment fragmentation in Nd2Zr2O7," Phys. Rev. B **94**, 104430 (2016).

[192] H. Yan, R. Pohle, and N. Shannon, "Half moons are pinch points with dispersion," Phys. Rev. B **98**, 140402 (2018).

[193] H. Yan, O. Benton, L. Jaubert, and N. Shannon, "Theory of multiple-phase competition in pyrochlore magnets with anisotropic exchange with application to Yb2Ti2O7, Er2Ti2O7, Er2Sn2O7," Phys. Rev. B **95**, 094422 (2017).

[194] M. Elhajal, B. Canals, R. Sunyer, and C. Lacroix, "Ordering in the pyrochlore antiferromagnet due to Dzyaloshinsky-Moriya interactions," Phys. Rev. B **71**, 094420 (2005).

[195] A. M. Hallas, J. Gaudet, and B. D. Gaulin, "Experimental Insights into Ground-State Selection of Quantum XY Pyrochlores," Annual Review of Condensed Matter Physics **9**, 105–124 (2018).

[196] B. Javanparast, A. G. R. Day, Z. Hao, and M. J. P. Gingras, "Order-by-disorder near criticality in XY pyrochlore magnets," Phys. Rev. B **91**, 174424 (2015).

[197] S. E. Palmer and J. T. Chalker, "Order induced by dipolar interactions in a geometrically frustrated antiferromagnet," Phys. Rev. B **62**, 488–492 (2000).

[198] M. Pretko, "Subdimensional particle structure of higher rank U(1) spin liquids," Phys. Rev. B **95**, 115139 (2017).

[199] M. Pretko, "The fracton gauge principle," Phys. Rev. B **98**, 115134 (2018).

[200] A. Prem, S. Vijay, Y.-Z. Chou, M. Pretko, and R. M. Nandkishore, "Pinch point singularities of tensor spin liquids," Phys. Rev. B **98**, 165140 (2018).

[201] H. Yan and J. Reuther, "Low energy structure of spiral spin liquids," (2021), arXiv:2112.10676 .

[202] H. Yan and A. Nevidomskyy, "Phonon induced rank-2 U(1) nematic liquid states," (2021), arXiv:2108.11484 .

[203] M. Pretko and L. Radzihovsky, "Fracton-Elasticity Duality," Physical Review Letters **120**, 195301 (2018).

[204] J. D. M. Champion, S. T. Bramwell, P. C. W. Holdsworth, and M. J. Harris, "Competition between exchange and anisotropy in a pyrochlore ferromagnet," EPL **57**, 93 (2002).

[205] E. Fradkin and S. H. Shenker, "Phase diagrams of lattice gauge theories with Higgs fields," Phys. Rev. D **19**, 3682–3697 (1979).

[206] S. Zhang, H. J. Changlani, K. W. Plumb, O. Tchernyshyov, and R. Moessner, "Dynamical Structure Factor of the Three-Dimensional Quantum Spin Liquid Candidate NaCaNi2F7," Phys. Rev. Lett. **122**, 167203 (2019).

[207] C. Mauws, A. M. Hallas, G. Sala, A. A. Aczel, P. M. Sarte, J. Gaudet, D. Ziat, J. A. Quilliam, J. A. Lussier, M. Bieringer, H. D. Zhou, A. R. Wildes, M. B. Stone, D. Abernathy, G. M. Luke, B. D. Gaulin, and C. R. Wiebe, "Dipolar-octupolar Ising antiferromagnetism in Sm2Ti2O7 : A moment fragmentation candidate," Phys. Rev. B **98**, 100401 (2018).

[208] S. J. Gomez, P. M. Sarte, M. Zelensky, A. M. Hallas, B. A. Gonzalez, K. H. Hong, E. J. Pace, S. Calder, M. B. Stone, Y. Su, E. Feng, M. D. Le, C. Stock, J. P. Attfield, S. D. Wilson, C. R. Wiebe, and A. A. Aczel, "Absence of moment fragmentation in the mixed B-site pyrochlore Nd2GaSbO7," Phys. Rev. B **103**, 214419 (2021).

[209] T. Banks, Y. Dothan, and D. Horn, "Geometric fermions," Physics Letters B **117**, 413–417 (1982).

[210] P. Becher and H. Joos, "The Dirac-Kähler equation and fermions on the lattice," Z. Phys. C - Particles and Fields **15**, 343–365 (1982).

[211] J. M. Rabin, "Homology theory of lattice fermion doubling," Nuclear Physics B **201**, 315–332 (1982).

[212] H. Joos, "On Geometry and Physics of Staggered Lattice Fermions," in *Clifford Algebras and Their Applications in Mathematical Physics*, NATO ASI Series, edited by J. S. R. Chisholm and A. K. Common (Springer Netherlands, Dordrecht, 1986) pp. 399–423.

[213] G. Farmelo, "Dirac's hidden geometry," Nature **437**, 323–323 (2005).

[214] M. A. Subramanian, G. Aravamudan, and G. V. Subba Rao, "Oxide pyrochlores — A review," Progress in Solid State Chemistry **15**, 55–143 (1983).

APPENDICES

Appendix A

The Pyrochlore Lattice and Coordinate Conventions

The pyrochlore lattice can be described by a face-centered cubic (FCC) lattice with a four sublattice basis. With respect to a Cartesian basis $\{\hat{\boldsymbol{x}}, \hat{\boldsymbol{y}}, \hat{\boldsymbol{z}}\}$, we take our FCC lattice vectors to be

$$\boldsymbol{a}_1 = a_0(\hat{\boldsymbol{x}} + \hat{\boldsymbol{y}})/2, \qquad \boldsymbol{a}_2 = a_0(\hat{\boldsymbol{x}} + \hat{\boldsymbol{z}})/2, \qquad \boldsymbol{a}_3 = a_0(\hat{\boldsymbol{y}} + \hat{\boldsymbol{z}})/2. \tag{A.1}$$

where a_0 is the length of the corresponding conventional cubic cell. The sublattice positions for the four magnetic ions in a unit cell are given by

$$\boldsymbol{c}_1 = \boldsymbol{0}, \qquad \boldsymbol{c}_2 = \boldsymbol{a}_1/2, \qquad \boldsymbol{c}_3 = \boldsymbol{a}_2/2, \qquad \boldsymbol{c}_4 = \boldsymbol{a}_3/2. \tag{A.2}$$

The corresponding local easy axes are then along the directions

$$\begin{aligned} \hat{\boldsymbol{z}}_1 &= (+\hat{\boldsymbol{x}} + \hat{\boldsymbol{y}} + \hat{\boldsymbol{z}})/\sqrt{3}\,, & \hat{\boldsymbol{z}}_2 &= (-\hat{\boldsymbol{x}} - \hat{\boldsymbol{y}} + \hat{\boldsymbol{z}})/\sqrt{3}\,, \\ \hat{\boldsymbol{z}}_3 &= (-\hat{\boldsymbol{x}} + \hat{\boldsymbol{y}} - \hat{\boldsymbol{z}})/\sqrt{3}\,, & \hat{\boldsymbol{z}}_4 &= (+\hat{\boldsymbol{x}} - \hat{\boldsymbol{y}} - \hat{\boldsymbol{z}})/\sqrt{3}\,. \end{aligned} \tag{A.3}$$

The above conventions are illustrated in Fig. A.1, which shows the conventional cubic cell from multiple angles. The cubic Bravais lattice vectors $a_0\hat{\boldsymbol{x}}$, etc. are shown in red, and the FCC Bravais lattice vectors are shown in blue. The tetrahedra of the pyrochlore lattice are shown, with the four sublattices (tetrahedron corners) colored (1) white, (2) red, (3) yellow, and (4) green with the same numbering as Eq. (A.2). Note that the sublattice positions are exactly at the midpoints of the FCC vectors. There are sixteen pyrochlore sites in the cubic cell, and correspondingly there are eight tetrahedra—four "up" and four "down", with the "down" tetrahedra not shown in (a) for

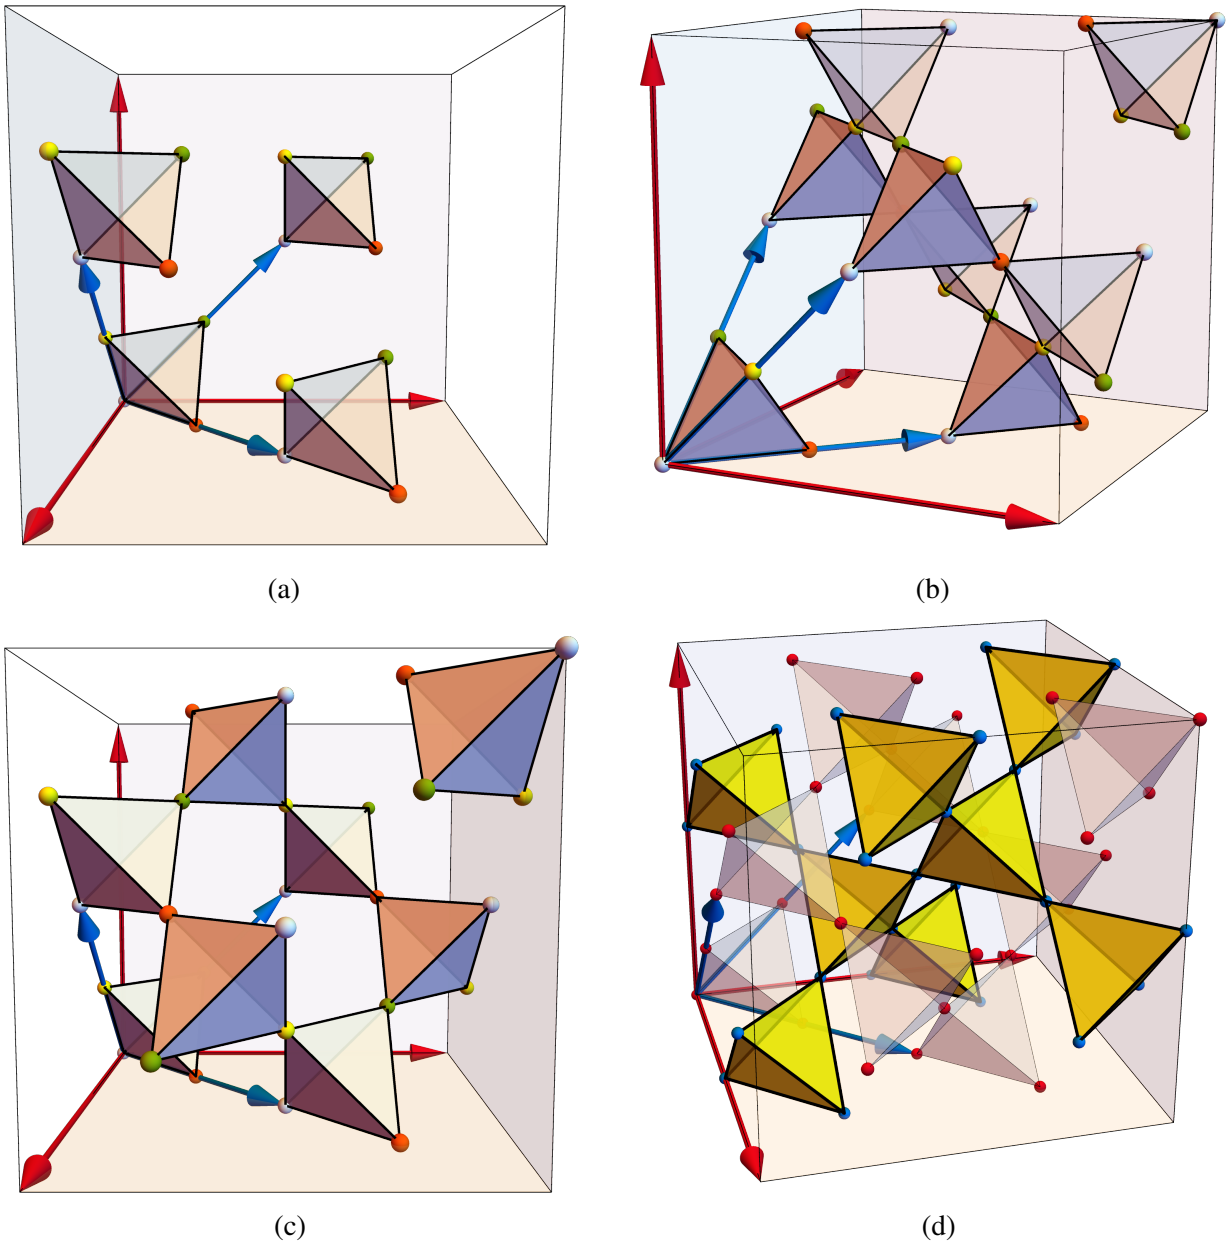
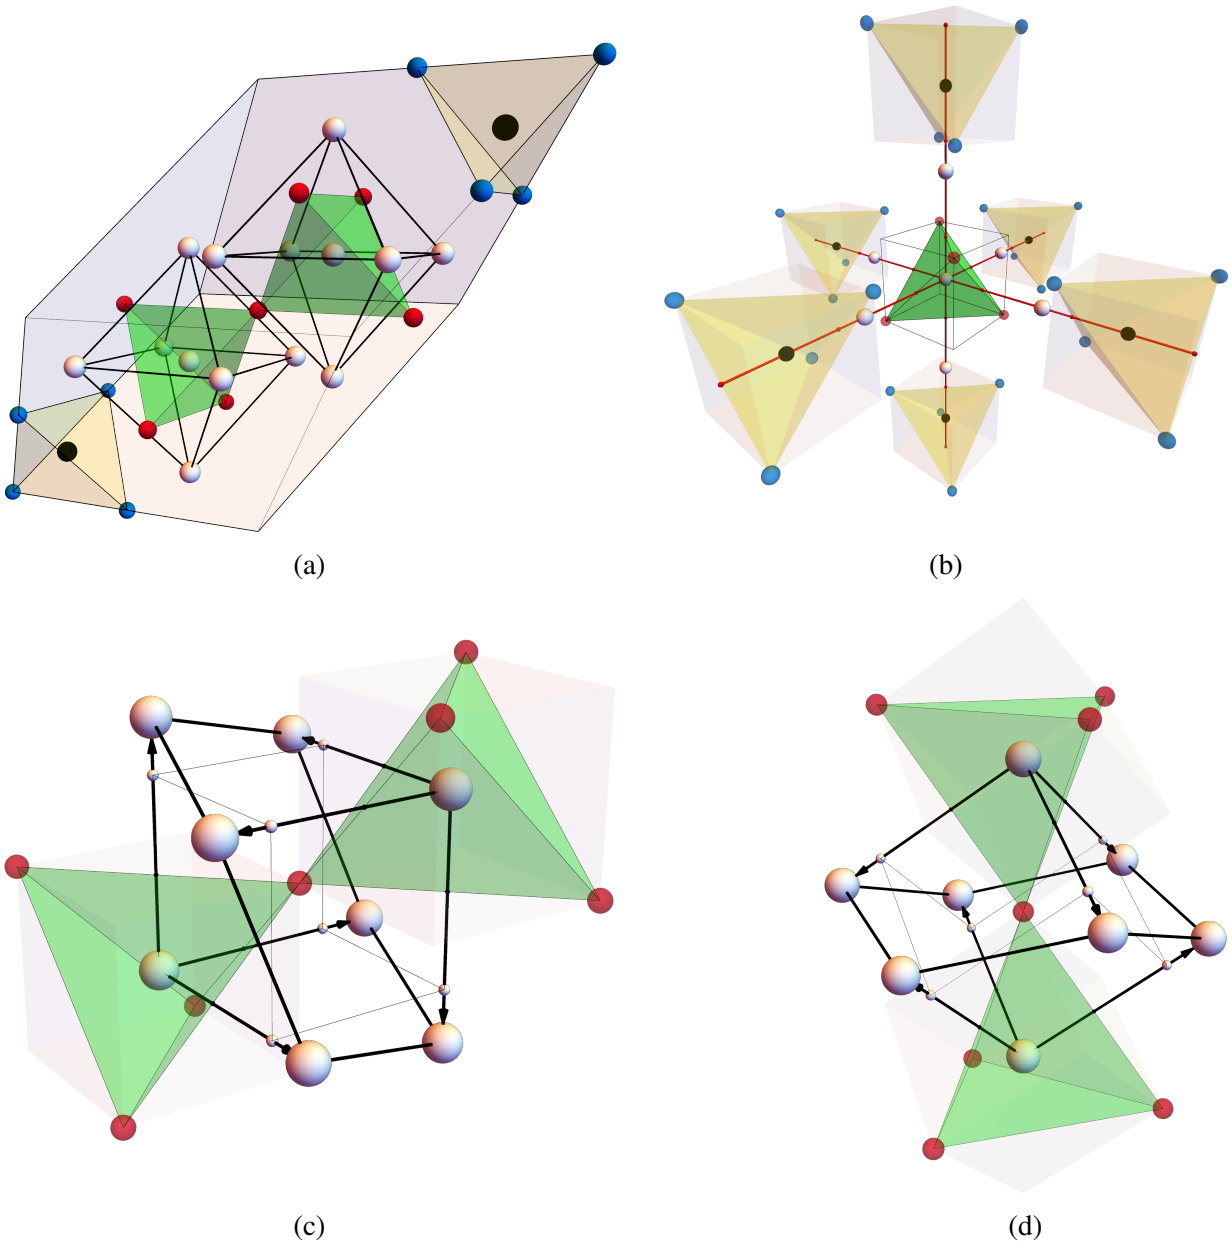


Figure A.1: Pyrochlore lattice cubic conventional cell. Cubic lattice vectors are shown in red, FCC lattice vectors are shown in blue. Sites of the pyrochlore lattice are shown as colored spheres, with four sublattices per FCC unit cell colored accordingly.

clarity. The hexagonal motif is visible in (c), and can be compared with Fig. 3.19. In (d) we show the dual pyrochlore lattice (yellow tetrahedra) which is interwoven with the primary pyrochlore lattice. Together they form an FCC lattice.

In the standard pyrochlore compounds with chemical formula $A_2B_2O_7$, the red sites correspond to magnetic *A* ions, the blue sites to non-magnetic *B* ions. There is an oxygen atom at the center of each *A*-tetrahedron but not in each *B*-tetrahedron, with additional oxygen atoms sitting midway between the *A* and *B* tetrahedron centers in each cubic direction [214]. This is shown in Fig. A.2, with *A* sites shown in red, *B* sites in blue, oxygens in white, and the vacancy in black. In (a) we show the FCC primitive cell, for which there are four *A* sites, four *B* sites, and fourteen oxygens per unit cell. The *A* tetrahedron contain a single oxygen, while the *B* tetrahedron center is vacant (black sphere). The other oxygen atoms between the *A* and *B* tetrahedron centers form octahedra surrounding the *A* tetrahedron. (b) The interstitial oxygens are slightly displaced away from the *A* tetrahedron center towards the vacancy site along the cubic axes. (c,d) The eight oxygen atoms surrounding a single *A* site form a distorted cube or hexahedron. The displacements towards the vacancies are shown by short black arrows. In the limit of zero displacement they would form a cube (small white spheres). Fig. A.2(a) and (b) should be compared with Ref. [214] Fig. 4, and Fig. A.2(c) and (d) should be compared with Ref. [214] Fig. 2(a).



(a) (b) (c) (d)

Figure A.2: Demonstration of the oxygen positions (white spheres) in the $A_2B_2O_7$ pyrochlore compounds. Red spheres are *A* sites, blue spheres are *B* sites, black spheres are oxygen vacancies.

Appendix B

Adjacency Matrices for the Pyrochlore Lattice and Projective Equivalence

Consider an isotropic model on the pyrochlore lattice whose Hamiltonian includes further than nearest-neighbor interactions. We index neighbors by n, where $n = 1$ corresponds to nearest-neighbor interactions, $n = 2$ to second-neighbor interactions, etc, as measured by their physical separation distances r_n. For a given n, there may be multiple symmetry-distinct subsets of neighbors at distance r_n, which we further index by a symmetry index $s_n = 1, \ldots, N_n$ where N_n is the number of symmetry-distinct subsets. To be precise, given two separation vectors $\boldsymbol{r}_1$ and $\boldsymbol{r}_2$ with length r_n, if these two vectors can be mapped to each other by lattice symmetries, they are in the same symmetry class s_n. The Hamiltonian can then be written as

$$H = \sum_{n=1}^{\infty} \sum_{s_n=1}^{N_n} \frac{1}{2} \sum_{ij} J_{n,s_n} \mathcal{A}_{ij}^{(n,s_n)} \boldsymbol{S}_i \cdot \boldsymbol{S}_j, \tag{B.1}$$

where J_{n,s_n} are energies and $\mathcal{A}^{(n,s_n)}$ is the adjacency matrix for n'th neighbors with symmetry type s_n on the lattice,

$$\mathcal{A}_{ij}^{(n,s_n)} = \begin{cases} 1 & i, j \text{ are neighbors of type } (n, s_n), \\ 0 & \text{otherwise.} \end{cases} \tag{B.2}$$

For n's with a single symmetry class we omit the index s_n, and it is customary to write the index $s_n = 1, 2, 3, \cdots$ with letters $a, b, c, \ldots$. In this section we give the explicit expressions for the first few adjacency matrices in the Fourier basis, in particular $\mathcal{A}^{(1)}$, $\mathcal{A}^{(2)}$, $\mathcal{A}^{(3a)}$ and $\mathcal{A}^{(3b)}$, which are discussed in the main text. Fig. 3.19 shows the second and third-type-b neighbors, while third-type-a neighbors are separated by an FCC lattice vector, see the blue vectors in Fig. A.1. We

also tabulate information for all neighbors up to $n = 100$, which are required for the purposes of including the infinite-range dipolar interactions.

B.1 Adjacency Matrices up to Third Neighbor

The adjacency matrices are translationally invariant and are thus 4×4 block diagonal in the Fourier basis, with blocks given by

$$\mathcal{A}_{\mu\nu}^{(n,s_n)}(\boldsymbol{q}) = \sum_{j\in\nu} \mathcal{A}_{ij}^{(n,s_n)} e^{i\boldsymbol{q}\cdot(\boldsymbol{r}_j - \boldsymbol{r}_i)} \quad i \in \mu, \tag{B.3}$$

For first and second neighbors, they are

$$\mathcal{A}_{\mu\nu}^{(1)}(\boldsymbol{q}) = 2\begin{pmatrix} 0 & c_{xy} & c_{xz} & c_{yz} \\ c_{xy} & 0 & c_{y\bar{z}} & c_{x\bar{z}} \\ c_{xz} & c_{y\bar{z}} & 0 & c_{x\bar{y}} \\ c_{yz} & c_{x\bar{z}} & c_{x\bar{y}} & 0 \end{pmatrix}, \tag{B.4}$$

$$\mathcal{A}_{\mu\nu}^{(2)}(\boldsymbol{q}) = 4\begin{pmatrix} 0 & c_{zz}\,c_{x\bar{y}} & c_{yy}\,c_{x\bar{z}} & c_{xx}\,c_{y\bar{z}} \\ c_{zz}\,c_{x\bar{y}} & 0 & c_{xx}\,c_{yz} & c_{yy}\,c_{xz} \\ c_{yy}\,c_{x\bar{z}} & c_{xx}\,c_{yz} & 0 & c_{zz}\,c_{xy} \\ c_{xx}\,c_{y\bar{z}} & c_{yy}\,c_{xz} & c_{zz}\,c_{xy} & 0 \end{pmatrix}, \tag{B.5}$$

where $c_{ab} \equiv \cos\left(\frac{q_a+q_b}{4}\right)$, $c_{a\bar{b}} \equiv \cos\left(\frac{q_a-q_b}{4}\right)$, with $a, b = x, y, z$. For third neighbors of symmetry type a,

$$\mathcal{A}_{\mu\nu}^{(3a)}(\boldsymbol{q}) = 2\begin{pmatrix} C_{+++} & 0 & 0 & 0 \\ 0 & C_{+--} & 0 & 0 \\ 0 & 0 & C_{--+} & 0 \\ 0 & 0 & 0 & C_{-+-} \end{pmatrix}, \tag{B.6}$$

with $C_{+--} = \cos\left(\frac{q_x+q_y}{2}\right) + \cos\left(\frac{q_y-q_z}{2}\right) + \cos\left(\frac{q_z-q_x}{2}\right)$, and so on. For third neighbors of symmetry type b,

$$\mathcal{A}_{\mu\nu}^{(3b)}(\boldsymbol{q}) = 2\begin{pmatrix} C_{---} & 0 & 0 & 0 \\ 0 & C_{-++} & 0 & 0 \\ 0 & 0 & C_{++-} & 0 \\ 0 & 0 & 0 & C_{+-+} \end{pmatrix}. \tag{B.7}$$

B.2 Adjacency Matrices Beyond Third Neighbor

Generally, exchange interactions beyond third neighbor are expected to be negligible for realistic models. However, they are important when including long-range dipole-dipole interactions, which we discuss in the next section. Therefore, we tabulate all further neighbors up to $n = 100$ (about twelve times the nearest-neighbor separation) in Tables B.1 to B.5. In these tables we provide the distance and symmetry index (n, s_n), the separation distance r_n in multiples of the nearest-neighbor separation distance r_1, and one example separation vector for each symmetry class $\boldsymbol{r}_{ij} = \boldsymbol{r}_j - \boldsymbol{r}_i$, where site i is taken to be on sublattice 1 in the convention laid out in Appendix A. All other separation vectors can be reconstructed by applying lattice point group symmetries (both D_{3d}, which fixes the sublattice, and T_d, which exchanges sublattices) to the given one (with care given to avoid double counting). The distance and symmetry index is primed for neighbors which lie in the same FCC sublattice.

B.3 Corrections to the Projective Equivalence Model of the Dipole-Dipole Interaction Matrix

As discussed in Sections 3.4.8 and 3.4.9, the dipole-dipole interaction matrix

$$\mathcal{D}_{ij} \equiv (1 - \delta_{ij}) \frac{\mu_0 \mu^2}{4\pi} \frac{\hat{\boldsymbol{\delta}}_i \cdot \hat{\boldsymbol{\delta}}_j - 3(\hat{\boldsymbol{\delta}}_i \cdot \hat{\boldsymbol{r}}_{ij})(\hat{\boldsymbol{\delta}}_j \cdot \hat{\boldsymbol{r}}_{ij})}{|\boldsymbol{r}_{ij}|^3} \tag{B.8}$$

can be approximated by the projector to the dispersive bands of the pyrochlore adjacency matrix according to Eq. (3.218). The matrix Δ_{ij} containing the corrections to the projective equivalence approximation can be expanded in terms of further-neighbor adjacency matrices according to Eq. (3.220). In Tables B.6 to B.8 we tabulate the values of $\Delta_{(n,s_n)}$ up to 100'th neighbor, which are used to approximate the dipole-dipole interaction matrix in SCGA calculations and to compute the bands of $\mathcal{D}_{ij}$.

(n, s_n)	r_n	$\boldsymbol{r}_{ij}$	(n, s_n)	r_n	$\boldsymbol{r}_{ij}$	(n, s_n)	r_n	$\boldsymbol{r}_{ij}$
1	1	$\left(\frac{1}{4}, \frac{1}{4}, 0\right)$	$14b$	$\sqrt{19}$	$\left(\frac{1}{4}, -\frac{1}{4}, \frac{3}{2}\right)$	$25a$	$\sqrt{33}$	$\left(\frac{5}{4}, \frac{5}{4}, 1\right)$
2	$\sqrt{3}$	$\left(\frac{1}{4}, -\frac{1}{4}, \frac{1}{2}\right)$	$15a'$	$\sqrt{20}$	$\left(\frac{3}{2}, \frac{1}{2}, 0\right)$	$25b$	$\sqrt{33}$	$\left(\frac{1}{4}, \frac{1}{4}, -2\right)$
$3a'$	$\sqrt{4}$	$\left(\frac{1}{2}, \frac{1}{2}, 0\right)$	$15b'$	$\sqrt{20}$	$\left(\frac{3}{2}, 0, -\frac{1}{2}\right)$	$25c$	$\sqrt{33}$	$\left(\frac{7}{4}, -\frac{1}{4}, -1\right)$
$3b'$	$\sqrt{4}$	$\left(\frac{1}{2}, 0, -\frac{1}{2}\right)$	16	$\sqrt{21}$	$\left(\frac{5}{4}, \frac{1}{4}, 1\right)$	26	$\sqrt{35}$	$\left(\frac{5}{4}, \frac{3}{4}, \frac{3}{2}\right)$
4	$\sqrt{5}$	$\left(\frac{3}{4}, -\frac{1}{4}, 0\right)$	17	$\sqrt{23}$	$\left(\frac{3}{4}, \frac{1}{4}, \frac{3}{2}\right)$	$27a'$	$\sqrt{36}$	$\left(2, \frac{1}{2}, -\frac{1}{2}\right)$
5	$\sqrt{7}$	$\left(\frac{3}{4}, \frac{1}{4}, \frac{1}{2}\right)$	$18a'$	$\sqrt{24}$	$(1, 1, -1)$	$27b'$	$\sqrt{36}$	$\left(2, \frac{1}{2}, \frac{1}{2}\right)$
$6'$	$\sqrt{8}$	$(1, 0, 0)$	$18b'$	$\sqrt{24}$	$(1, 1, 1)$	$27c'$	$\sqrt{36}$	$\left(2, -\frac{1}{2}, -\frac{1}{2}\right)$
$7a$	$\sqrt{9}$	$\left(\frac{3}{4}, \frac{3}{4}, 0\right)$	$19a$	$\sqrt{25}$	$\left(\frac{5}{4}, \frac{5}{4}, 0\right)$	$27d'$	$\sqrt{36}$	$\left(\frac{3}{2}, \frac{3}{2}, 0\right)$
$7b$	$\sqrt{9}$	$\left(\frac{1}{4}, \frac{1}{4}, -1\right)$	$19b$	$\sqrt{25}$	$\left(\frac{7}{4}, -\frac{1}{4}, 0\right)$	$27e'$	$\sqrt{36}$	$\left(\frac{3}{2}, 0, -\frac{3}{2}\right)$
8	$\sqrt{11}$	$\left(\frac{3}{4}, -\frac{3}{4}, \frac{1}{2}\right)$	$19c$	$\sqrt{25}$	$\left(\frac{5}{4}, -\frac{3}{4}, -1\right)$	$28a$	$\sqrt{37}$	$\left(\frac{7}{4}, \frac{3}{4}, 1\right)$
$9a'$	$\sqrt{12}$	$\left(1, \frac{1}{2}, -\frac{1}{2}\right)$	$20a$	$\sqrt{27}$	$\left(\frac{7}{4}, \frac{1}{4}, \frac{1}{2}\right)$	$28b$	$\sqrt{37}$	$\left(\frac{3}{4}, -\frac{1}{4}, -2\right)$
$9b'$	$\sqrt{12}$	$\left(1, \frac{1}{2}, \frac{1}{2}\right)$	$20b$	$\sqrt{27}$	$\left(\frac{3}{4}, -\frac{3}{4}, \frac{3}{2}\right)$	$28c$	$\sqrt{37}$	$\left(\frac{7}{4}, -\frac{5}{4}, 0\right)$
$9c'$	$\sqrt{12}$	$\left(1, -\frac{1}{2}, -\frac{1}{2}\right)$	$20c$	$\sqrt{27}$	$\left(\frac{5}{4}, -\frac{5}{4}, \frac{1}{2}\right)$	29	$\sqrt{39}$	$\left(\frac{7}{4}, \frac{5}{4}, \frac{1}{2}\right)$
$10a$	$\sqrt{13}$	$\left(\frac{5}{4}, \frac{1}{4}, 0\right)$	$21a'$	$\sqrt{28}$	$\left(\frac{3}{2}, 1, -\frac{1}{2}\right)$	$30a'$	$\sqrt{40}$	$(2, 1, 0)$
$10b$	$\sqrt{13}$	$\left(\frac{3}{4}, -\frac{1}{4}, -1\right)$	$21b'$	$\sqrt{28}$	$\left(\frac{3}{2}, 1, \frac{1}{2}\right)$	$30b'$	$\sqrt{40}$	$(2, 0, -1)$
11	$\sqrt{15}$	$\left(\frac{5}{4}, -\frac{1}{4}, -\frac{1}{2}\right)$	$21c'$	$\sqrt{28}$	$\left(\frac{3}{2}, \frac{1}{2}, -1\right)$	$31a$	$\sqrt{41}$	$\left(\frac{3}{4}, \frac{3}{4}, 2\right)$
$12a'$	$\sqrt{16}$	$(1, 1, 0)$	$21d'$	$\sqrt{28}$	$\left(\frac{3}{2}, -\frac{1}{2}, -1\right)$	$31b$	$\sqrt{41}$	$\left(\frac{9}{4}, \frac{1}{4}, 0\right)$
$12b'$	$\sqrt{16}$	$(1, 0, -1)$	22	$\sqrt{29}$	$\left(\frac{7}{4}, \frac{3}{4}, 0\right)$	$32a$	$\sqrt{43}$	$\left(\frac{7}{4}, \frac{1}{4}, \frac{3}{2}\right)$
$13a$	$\sqrt{17}$	$\left(\frac{3}{4}, \frac{3}{4}, 1\right)$	$23a$	$\sqrt{31}$	$\left(\frac{7}{4}, -\frac{3}{4}, -\frac{1}{2}\right)$	$32b$	$\sqrt{43}$	$\left(\frac{9}{4}, -\frac{1}{4}, -\frac{1}{2}\right)$
$13b$	$\sqrt{17}$	$\left(\frac{5}{4}, -\frac{3}{4}, 0\right)$	$23b$	$\sqrt{31}$	$\left(\frac{5}{4}, -\frac{1}{4}, -\frac{3}{2}\right)$	$32c$	$\sqrt{43}$	$\left(\frac{5}{4}, -\frac{5}{4}, \frac{3}{2}\right)$
$14a$	$\sqrt{19}$	$\left(\frac{5}{4}, \frac{3}{4}, \frac{1}{2}\right)$	$24'$	$\sqrt{32}$	$(2, 0, 0)$			

Table B.1: Further neighbor symmetry classes up to 100'th nearest neighbor.

(n, s_n)	r_n	$\boldsymbol{r}_{ij}$	(n, s_n)	r_n	$\boldsymbol{r}_{ij}$	(n, s_n)	r_n	$\boldsymbol{r}_{ij}$
$33a'$	$\sqrt{44}$	$\left(\frac{3}{2}, \frac{3}{2}, -1\right)$	$39f'$	$\sqrt{52}$	$\left(2, -\frac{1}{2}, -\frac{3}{2}\right)$	$47a'$	$\sqrt{64}$	$(2, 2, 0)$
$33b'$	$\sqrt{44}$	$\left(\frac{3}{2}, \frac{3}{2}, 1\right)$	$40a$	$\sqrt{53}$	$\left(\frac{9}{4}, \frac{5}{4}, 0\right)$	$47b'$	$\sqrt{64}$	$(2, 0, -2)$
$33c'$	$\sqrt{44}$	$\left(\frac{3}{2}, 1, -\frac{3}{2}\right)$	$40b$	$\sqrt{53}$	$\left(\frac{9}{4}, -\frac{3}{4}, -1\right)$	$48a$	$\sqrt{65}$	$\left(\frac{11}{4}, \frac{3}{4}, 0\right)$
$34a$	$\sqrt{45}$	$\left(\frac{7}{4}, -\frac{5}{4}, -1\right)$	$41a$	$\sqrt{55}$	$\left(\frac{7}{4}, \frac{5}{4}, \frac{3}{2}\right)$	$48b$	$\sqrt{65}$	$\left(\frac{9}{4}, -\frac{7}{4}, 0\right)$
$34b$	$\sqrt{45}$	$\left(\frac{9}{4}, -\frac{3}{4}, 0\right)$	$41b$	$\sqrt{55}$	$\left(\frac{3}{4}, \frac{1}{4}, -\frac{5}{2}\right)$	$49a$	$\sqrt{67}$	$\left(\frac{9}{4}, \frac{7}{4}, \frac{1}{2}\right)$
$34c$	$\sqrt{45}$	$\left(\frac{5}{4}, \frac{1}{4}, -2\right)$	$41c$	$\sqrt{55}$	$\left(\frac{9}{4}, -\frac{5}{4}, -\frac{1}{2}\right)$	$49b$	$\sqrt{67}$	$\left(\frac{11}{4}, -\frac{3}{4}, -\frac{1}{2}\right)$
$35a$	$\sqrt{47}$	$\left(\frac{9}{4}, \frac{3}{4}, \frac{1}{2}\right)$	$42a$	$\sqrt{57}$	$\left(\frac{7}{4}, \frac{7}{4}, 1\right)$	$49c$	$\sqrt{67}$	$\left(\frac{5}{4}, \frac{3}{4}, -\frac{5}{2}\right)$
$35b$	$\sqrt{47}$	$\left(\frac{7}{4}, -\frac{3}{4}, -\frac{3}{2}\right)$	$42b$	$\sqrt{57}$	$\left(\frac{7}{4}, -\frac{1}{4}, -2\right)$	$49d$	$\sqrt{67}$	$\left(\frac{7}{4}, -\frac{7}{4}, \frac{3}{2}\right)$
$36a'$	$\sqrt{48}$	$(2, 1, -1)$	$42c$	$\sqrt{57}$	$\left(\frac{5}{4}, \frac{5}{4}, -2\right)$	$50a'$	$\sqrt{68}$	$\left(2, \frac{3}{2}, -\frac{3}{2}\right)$
$36b'$	$\sqrt{48}$	$(2, 1, 1)$	$43a$	$\sqrt{59}$	$\left(\frac{9}{4}, -\frac{1}{4}, -\frac{3}{2}\right)$	$50b'$	$\sqrt{68}$	$\left(\frac{5}{2}, \frac{3}{2}, 0\right)$
$36c'$	$\sqrt{48}$	$(2, -1, -1)$	$43b$	$\sqrt{59}$	$\left(\frac{3}{4}, -\frac{3}{4}, \frac{5}{2}\right)$	$50c'$	$\sqrt{68}$	$\left(\frac{5}{2}, 0, -\frac{3}{2}\right)$
$37a$	$\sqrt{49}$	$\left(\frac{9}{4}, \frac{1}{4}, 1\right)$	$44a'$	$\sqrt{60}$	$\left(\frac{5}{2}, 1, -\frac{1}{2}\right)$	$50d'$	$\sqrt{68}$	$\left(2, \frac{3}{2}, \frac{3}{2}\right)$
$37b$	$\sqrt{49}$	$\left(\frac{5}{4}, -\frac{3}{4}, -2\right)$	$44b'$	$\sqrt{60}$	$\left(\frac{5}{2}, 1, \frac{1}{2}\right)$	$50e'$	$\sqrt{68}$	$\left(2, -\frac{3}{2}, -\frac{3}{2}\right)$
$37c$	$\sqrt{49}$	$\left(\frac{7}{4}, \frac{7}{4}, 0\right)$	$44c'$	$\sqrt{60}$	$\left(\frac{5}{2}, \frac{1}{2}, -1\right)$	$51a$	$\sqrt{69}$	$\left(\frac{7}{4}, -\frac{5}{4}, -2\right)$
$38a$	$\sqrt{51}$	$\left(\frac{7}{4}, -\frac{7}{4}, \frac{1}{2}\right)$	$44d'$	$\sqrt{60}$	$\left(\frac{5}{2}, -\frac{1}{2}, -1\right)$	$51b$	$\sqrt{69}$	$\left(\frac{11}{4}, -\frac{1}{4}, -1\right)$
$38b$	$\sqrt{51}$	$\left(\frac{1}{4}, -\frac{1}{4}, \frac{5}{2}\right)$	$45a$	$\sqrt{61}$	$\left(\frac{9}{4}, \frac{5}{4}, 1\right)$	52	$\sqrt{71}$	$\left(\frac{9}{4}, -\frac{5}{4}, -\frac{3}{2}\right)$
$39a'$	$\sqrt{52}$	$\left(2, \frac{3}{2}, -\frac{1}{2}\right)$	$45b$	$\sqrt{61}$	$\left(\frac{11}{4}, -\frac{1}{4}, 0\right)$	$53a'$	$\sqrt{72}$	$(2, 2, -1)$
$39b'$	$\sqrt{52}$	$\left(\frac{5}{2}, \frac{1}{2}, 0\right)$	$45c$	$\sqrt{61}$	$\left(\frac{7}{4}, \frac{3}{4}, -2\right)$	$53b'$	$\sqrt{72}$	$(3, 0, 0)$
$39c'$	$\sqrt{52}$	$\left(\frac{5}{2}, 0, -\frac{1}{2}\right)$	$46a$	$\sqrt{63}$	$\left(\frac{9}{4}, \frac{3}{4}, \frac{3}{2}\right)$	$53c'$	$\sqrt{72}$	$(2, 2, 1)$
$39d'$	$\sqrt{52}$	$\left(2, \frac{3}{2}, \frac{1}{2}\right)$	$46b$	$\sqrt{63}$	$\left(\frac{11}{4}, \frac{1}{4}, -\frac{1}{2}\right)$	$53d'$	$\sqrt{72}$	$(2, 1, -2)$
$39e'$	$\sqrt{52}$	$\left(2, \frac{1}{2}, -\frac{3}{2}\right)$	$46c$	$\sqrt{63}$	$\left(\frac{5}{4}, -\frac{1}{4}, -\frac{5}{2}\right)$			

Table B.2: Further neighbor symmetry classes up to 100'th nearest neighbor.

(n, s_n)	r_n	$\boldsymbol{r}_{ij}$	(n, s_n)	r_n	$\boldsymbol{r}_{ij}$	(n, s_n)	r_n	$\boldsymbol{r}_{ij}$
$54a$	$\sqrt{73}$	$\left(\frac{11}{4}, \frac{3}{4}, 1\right)$	$60a$	$\sqrt{81}$	$\left(\frac{7}{4}, \frac{7}{4}, 2\right)$	$66a$	$\sqrt{89}$	$\left(\frac{9}{4}, \frac{9}{4}, 1\right)$
$54b$	$\sqrt{73}$	$\left(\frac{1}{4}, \frac{1}{4}, -3\right)$	$60b$	$\sqrt{81}$	$\left(\frac{3}{4}, \frac{3}{4}, -3\right)$	$66b$	$\sqrt{89}$	$\left(\frac{13}{4}, -\frac{3}{4}, 0\right)$
$54c$	$\sqrt{73}$	$\left(\frac{11}{4}, -\frac{5}{4}, 0\right)$	$60c$	$\sqrt{81}$	$\left(\frac{11}{4}, -\frac{5}{4}, -1\right)$	$66c$	$\sqrt{89}$	$\left(\frac{5}{4}, -\frac{3}{4}, -3\right)$
$54d$	$\sqrt{73}$	$\left(\frac{9}{4}, \frac{1}{4}, -2\right)$	$60d$	$\sqrt{81}$	$\left(\frac{9}{4}, \frac{9}{4}, 0\right)$	$67a$	$\sqrt{91}$	$\left(\frac{11}{4}, \frac{5}{4}, \frac{3}{2}\right)$
$54e$	$\sqrt{73}$	$\left(\frac{9}{4}, -\frac{7}{4}, -1\right)$	$61a$	$\sqrt{83}$	$\left(\frac{9}{4}, \frac{7}{4}, \frac{3}{2}\right)$	$67b$	$\sqrt{91}$	$\left(\frac{13}{4}, \frac{3}{4}, -\frac{1}{2}\right)$
$55a$	$\sqrt{75}$	$\left(\frac{11}{4}, \frac{5}{4}, \frac{1}{2}\right)$	$61b$	$\sqrt{83}$	$\left(\frac{11}{4}, -\frac{3}{4}, -\frac{3}{2}\right)$	$67c$	$\sqrt{91}$	$\left(\frac{9}{4}, -\frac{1}{4}, -\frac{5}{2}\right)$
$55b$	$\sqrt{75}$	$\left(\frac{7}{4}, \frac{1}{4}, -\frac{5}{2}\right)$	$61c$	$\sqrt{83}$	$\left(\frac{9}{4}, -\frac{9}{4}, \frac{1}{2}\right)$	$68a'$	$\sqrt{92}$	$\left(3, \frac{3}{2}, -\frac{1}{2}\right)$
$55c$	$\sqrt{75}$	$\left(\frac{5}{4}, -\frac{5}{4}, \frac{5}{2}\right)$	$62a'$	$\sqrt{84}$	$\left(\frac{5}{2}, 2, -\frac{1}{2}\right)$	$68b'$	$\sqrt{92}$	$\left(3, \frac{3}{2}, \frac{1}{2}\right)$
$56a'$	$\sqrt{76}$	$\left(\frac{5}{2}, \frac{3}{2}, -1\right)$	$62b'$	$\sqrt{84}$	$\left(\frac{5}{2}, 2, \frac{1}{2}\right)$	$68c'$	$\sqrt{92}$	$\left(3, \frac{1}{2}, -\frac{3}{2}\right)$
$56b'$	$\sqrt{76}$	$\left(3, \frac{1}{2}, \frac{1}{2}\right)$	$62c'$	$\sqrt{84}$	$\left(\frac{5}{2}, \frac{1}{2}, -2\right)$	$68d'$	$\sqrt{92}$	$\left(3, -\frac{1}{2}, -\frac{3}{2}\right)$
$56c'$	$\sqrt{76}$	$\left(3, \frac{1}{2}, -\frac{1}{2}\right)$	$62d'$	$\sqrt{84}$	$\left(\frac{5}{2}, -\frac{1}{2}, -2\right)$	$69a$	$\sqrt{93}$	$\left(\frac{11}{4}, \frac{7}{4}, 1\right)$
$56d'$	$\sqrt{76}$	$\left(3, -\frac{1}{2}, -\frac{1}{2}\right)$	$63a$	$\sqrt{85}$	$\left(\frac{9}{4}, \frac{5}{4}, 2\right)$	$69b$	$\sqrt{93}$	$\left(\frac{13}{4}, \frac{1}{4}, -1\right)$
$56e'$	$\sqrt{76}$	$\left(\frac{5}{2}, \frac{3}{2}, 1\right)$	$63b$	$\sqrt{85}$	$\left(\frac{13}{4}, \frac{1}{4}, 0\right)$	$69c$	$\sqrt{93}$	$\left(\frac{11}{4}, -\frac{1}{4}, -2\right)$
$56f'$	$\sqrt{76}$	$\left(\frac{5}{2}, 1, -\frac{3}{2}\right)$	$63c$	$\sqrt{85}$	$\left(\frac{5}{4}, \frac{1}{4}, -3\right)$	70	$\sqrt{95}$	$\left(\frac{9}{4}, \frac{3}{4}, \frac{5}{2}\right)$
$56g'$	$\sqrt{76}$	$\left(\frac{5}{2}, -1, -\frac{3}{2}\right)$	$63d$	$\sqrt{85}$	$\left(\frac{11}{4}, \frac{7}{4}, 0\right)$	$71a'$	$\sqrt{96}$	$(2, 2, -2)$
$57a$	$\sqrt{77}$	$\left(\frac{9}{4}, -\frac{3}{4}, -2\right)$	$64a$	$\sqrt{87}$	$\left(\frac{7}{4}, \frac{5}{4}, \frac{5}{2}\right)$	$71b'$	$\sqrt{96}$	$(2, 2, 2)$
$57b$	$\sqrt{77}$	$\left(\frac{3}{4}, -\frac{1}{4}, -3\right)$	$64b$	$\sqrt{87}$	$\left(\frac{13}{4}, -\frac{1}{4}, -\frac{1}{2}\right)$			
$58a$	$\sqrt{79}$	$\left(\frac{11}{4}, \frac{1}{4}, \frac{3}{2}\right)$	$64c$	$\sqrt{87}$	$\left(\frac{11}{4}, -\frac{7}{4}, -\frac{1}{2}\right)$			
$58b$	$\sqrt{79}$	$\left(\frac{7}{4}, -\frac{3}{4}, -\frac{5}{2}\right)$	$65a'$	$\sqrt{88}$	$(3, 1, -1)$			
$59a'$	$\sqrt{80}$	$(3, 1, 0)$	$65b'$	$\sqrt{88}$	$(3, 1, 1)$			
$59b'$	$\sqrt{80}$	$(3, 0, -1)$	$65c'$	$\sqrt{88}$	$(3, -1, -1)$			

Table B.3: Further neighbor symmetry classes up to 100'th nearest neighbor.

(n, s_n)	r_n	$\boldsymbol{r}_{ij}$	(n, s_n)	r_n	$\boldsymbol{r}_{ij}$	(n, s_n)	r_n	$\boldsymbol{r}_{ij}$
$72a$	$\sqrt{97}$	$\left(\frac{11}{4}, \frac{3}{4}, 2\right)$	$76b$	$\sqrt{103}$	$\left(\frac{3}{4}, \frac{1}{4}, -\frac{7}{2}\right)$	$81b$	$\sqrt{109}$	$\left(\frac{13}{4}, -\frac{7}{4}, 0\right)$
$72b$	$\sqrt{97}$	$\left(\frac{13}{4}, \frac{5}{4}, 0\right)$	$76c$	$\sqrt{103}$	$\left(\frac{13}{4}, -\frac{1}{4}, -\frac{3}{2}\right)$	$81c$	$\sqrt{109}$	$\left(\frac{11}{4}, -\frac{9}{4}, -1\right)$
$72c$	$\sqrt{97}$	$\left(\frac{13}{4}, -\frac{3}{4}, -1\right)$	$76d$	$\sqrt{103}$	$\left(\frac{11}{4}, -\frac{7}{4}, -\frac{3}{2}\right)$	$82a$	$\sqrt{111}$	$\left(\frac{13}{4}, \frac{7}{4}, \frac{1}{2}\right)$
$72d$	$\sqrt{97}$	$\left(\frac{7}{4}, -\frac{1}{4}, -3\right)$	$76e$	$\sqrt{103}$	$\left(\frac{9}{4}, -\frac{5}{4}, -\frac{5}{2}\right)$	$82b$	$\sqrt{111}$	$\left(\frac{5}{4}, -\frac{1}{4}, -\frac{7}{2}\right)$
$72e$	$\sqrt{97}$	$\left(\frac{5}{4}, \frac{5}{4}, -3\right)$	$77a'$	$\sqrt{104}$	$(3, 2, 0)$	$82c$	$\sqrt{111}$	$\left(\frac{11}{4}, \frac{1}{4}, -\frac{5}{2}\right)$
$72f$	$\sqrt{97}$	$\left(\frac{9}{4}, -\frac{7}{4}, -2\right)$	$77b'$	$\sqrt{104}$	$(3, 0, -2)$	$83a'$	$\sqrt{112}$	$(3, 2, -1)$
$73a$	$\sqrt{99}$	$\left(\frac{7}{4}, -\frac{7}{4}, \frac{5}{2}\right)$	$78a$	$\sqrt{105}$	$\left(\frac{13}{4}, \frac{5}{4}, 1\right)$	$83b'$	$\sqrt{112}$	$(3, 2, 1)$
$73b$	$\sqrt{99}$	$\left(\frac{1}{4}, -\frac{1}{4}, \frac{7}{2}\right)$	$78b$	$\sqrt{105}$	$\left(\frac{11}{4}, -\frac{5}{4}, -2\right)$	$83c'$	$\sqrt{112}$	$(3, 1, -2)$
$73c$	$\sqrt{99}$	$\left(\frac{13}{4}, -\frac{5}{4}, -\frac{1}{2}\right)$	$79a$	$\sqrt{107}$	$\left(\frac{13}{4}, \frac{3}{4}, \frac{3}{2}\right)$	$83d'$	$\sqrt{112}$	$(3, -1, -2)$
$73d$	$\sqrt{99}$	$\left(\frac{9}{4}, -\frac{9}{4}, \frac{3}{2}\right)$	$79b$	$\sqrt{107}$	$\left(\frac{3}{4}, -\frac{3}{4}, \frac{7}{2}\right)$	$84a$	$\sqrt{113}$	$\left(\frac{9}{4}, \frac{9}{4}, 2\right)$
$74a'$	$\sqrt{100}$	$\left(\frac{5}{2}, 2, -\frac{3}{2}\right)$	$80a'$	$\sqrt{108}$	$\left(3, \frac{3}{2}, -\frac{3}{2}\right)$	$84b$	$\sqrt{113}$	$\left(\frac{15}{4}, -\frac{1}{4}, 0\right)$
$74b'$	$\sqrt{100}$	$\left(\frac{7}{2}, \frac{1}{2}, 0\right)$	$80b'$	$\sqrt{108}$	$\left(\frac{7}{2}, 1, \frac{1}{2}\right)$	$84c$	$\sqrt{113}$	$\left(\frac{9}{4}, \frac{1}{4}, -3\right)$
$74c'$	$\sqrt{100}$	$\left(\frac{7}{2}, 0, -\frac{1}{2}\right)$	$80c'$	$\sqrt{108}$	$\left(\frac{7}{2}, 1, -\frac{1}{2}\right)$	$85a$	$\sqrt{115}$	$\left(\frac{9}{4}, \frac{7}{4}, \frac{5}{2}\right)$
$74d'$	$\sqrt{100}$	$\left(\frac{5}{2}, \frac{5}{2}, 0\right)$	$80d'$	$\sqrt{108}$	$\left(\frac{7}{2}, \frac{1}{2}, -1\right)$	$85b$	$\sqrt{115}$	$\left(\frac{15}{4}, \frac{1}{4}, -\frac{1}{2}\right)$
$74e'$	$\sqrt{100}$	$\left(\frac{5}{2}, 2, \frac{3}{2}\right)$	$80e'$	$\sqrt{108}$	$\left(\frac{7}{2}, -\frac{1}{2}, -1\right)$	$85c$	$\sqrt{115}$	$\left(\frac{5}{4}, \frac{3}{4}, -\frac{7}{2}\right)$
$74f'$	$\sqrt{100}$	$\left(\frac{5}{2}, \frac{3}{2}, -2\right)$	$80f'$	$\sqrt{108}$	$\left(3, \frac{3}{2}, \frac{3}{2}\right)$	$85d$	$\sqrt{115}$	$\left(\frac{13}{4}, -\frac{5}{4}, -\frac{3}{2}\right)$
$74g'$	$\sqrt{100}$	$\left(\frac{5}{2}, 0, -\frac{5}{2}\right)$	$80g'$	$\sqrt{108}$	$\left(3, -\frac{3}{2}, -\frac{3}{2}\right)$	$85e$	$\sqrt{115}$	$\left(\frac{11}{4}, -\frac{3}{4}, -\frac{5}{2}\right)$
$74h'$	$\sqrt{100}$	$\left(\frac{5}{2}, -\frac{3}{2}, -2\right)$	$80h'$	$\sqrt{108}$	$\left(\frac{5}{2}, \frac{5}{2}, 1\right)$	$86a'$	$\sqrt{116}$	$\left(\frac{7}{2}, \frac{3}{2}, 0\right)$
$75a$	$\sqrt{101}$	$\left(\frac{7}{4}, \frac{3}{4}, 3\right)$	$80i'$	$\sqrt{108}$	$\left(\frac{5}{2}, \frac{5}{2}, -1\right)$	$86b'$	$\sqrt{116}$	$\left(\frac{7}{2}, 0, -\frac{3}{2}\right)$
$75b$	$\sqrt{101}$	$\left(\frac{11}{4}, -\frac{9}{4}, 0\right)$	$80j'$	$\sqrt{108}$	$\left(\frac{5}{2}, 1, -\frac{5}{2}\right)$			
$76a$	$\sqrt{103}$	$\left(\frac{11}{4}, \frac{9}{4}, \frac{1}{2}\right)$	$81a$	$\sqrt{109}$	$\left(\frac{7}{4}, -\frac{5}{4}, -3\right)$			

Table B.4: Further neighbor symmetry classes up to 100'th nearest neighbor.

(n, s_n)	r_n	$\boldsymbol{r}_{ij}$	(n, s_n)	r_n	$\boldsymbol{r}_{ij}$	(n, s_n)	r_n	$\boldsymbol{r}_{ij}$
$87a$	$\sqrt{117}$	$\left(\frac{11}{4}, \frac{7}{4}, 2\right)$	$91g'$	$\sqrt{124}$	$\left(3, \frac{1}{2}, -\frac{5}{2}\right)$	$97f'$	$\sqrt{132}$	$\left(\frac{7}{2}, 2, -\frac{1}{2}\right)$
$87b$	$\sqrt{117}$	$\left(\frac{15}{4}, \frac{3}{4}, 0\right)$	$91h'$	$\sqrt{124}$	$\left(3, -\frac{1}{2}, -\frac{5}{2}\right)$	$97g'$	$\sqrt{132}$	$\left(\frac{7}{2}, \frac{1}{2}, -2\right)$
$87c$	$\sqrt{117}$	$\left(\frac{13}{4}, \frac{1}{4}, -2\right)$	$92a$	$\sqrt{125}$	$\left(\frac{9}{4}, \frac{5}{4}, 3\right)$	$97h'$	$\sqrt{132}$	$\left(\frac{7}{2}, -\frac{1}{2}, -2\right)$
$87d$	$\sqrt{117}$	$\left(\frac{13}{4}, -\frac{7}{4}, -1\right)$	$92b$	$\sqrt{125}$	$\left(\frac{15}{4}, \frac{3}{4}, -1\right)$	$97i'$	$\sqrt{132}$	$\left(\frac{5}{2}, \frac{5}{2}, 2\right)$
$87e$	$\sqrt{117}$	$\left(\frac{9}{4}, -\frac{3}{4}, -3\right)$	$92c$	$\sqrt{125}$	$\left(\frac{15}{4}, -\frac{5}{4}, 0\right)$	$97j'$	$\sqrt{132}$	$\left(\frac{5}{2}, 2, -\frac{5}{2}\right)$
$88a$	$\sqrt{119}$	$\left(\frac{11}{4}, \frac{9}{4}, \frac{3}{2}\right)$	$92d$	$\sqrt{125}$	$\left(\frac{13}{4}, \frac{9}{4}, 0\right)$	$98a$	$\sqrt{133}$	$\left(\frac{13}{4}, \frac{9}{4}, 1\right)$
$88b$	$\sqrt{119}$	$\left(\frac{15}{4}, -\frac{3}{4}, -\frac{1}{2}\right)$	$93a$	$\sqrt{127}$	$\left(\frac{13}{4}, \frac{7}{4}, \frac{3}{2}\right)$	$98b$	$\sqrt{133}$	$\left(\frac{3}{4}, -\frac{1}{4}, -4\right)$
$89a$	$\sqrt{121}$	$\left(\frac{7}{4}, \frac{7}{4}, 3\right)$	$93b$	$\sqrt{127}$	$\left(\frac{15}{4}, \frac{5}{4}, -\frac{1}{2}\right)$	$98c$	$\sqrt{133}$	$\left(\frac{15}{4}, -\frac{5}{4}, -1\right)$
$89b$	$\sqrt{121}$	$\left(\frac{15}{4}, -\frac{1}{4}, -1\right)$	$93c$	$\sqrt{127}$	$\left(\frac{7}{4}, -\frac{3}{4}, -\frac{7}{2}\right)$	$98d$	$\sqrt{133}$	$\left(\frac{11}{4}, -\frac{1}{4}, -3\right)$
$89c$	$\sqrt{121}$	$\left(\frac{13}{4}, -\frac{3}{4}, -2\right)$	$93d$	$\sqrt{127}$	$\left(\frac{13}{4}, -\frac{9}{4}, -\frac{1}{2}\right)$	$98e$	$\sqrt{133}$	$\left(\frac{11}{4}, -\frac{9}{4}, -2\right)$
$89d$	$\sqrt{121}$	$\left(\frac{11}{4}, \frac{11}{4}, 0\right)$	$94'$	$\sqrt{128}$	$(4, 0, 0)$	$99a$	$\sqrt{135}$	$\left(\frac{11}{4}, -\frac{7}{4}, -\frac{5}{2}\right)$
$90a$	$\sqrt{123}$	$\left(\frac{11}{4}, \frac{5}{4}, \frac{5}{2}\right)$	$95a$	$\sqrt{129}$	$\left(\frac{13}{4}, \frac{5}{4}, 2\right)$	$99b$	$\sqrt{135}$	$\left(\frac{15}{4}, -\frac{3}{4}, -\frac{3}{2}\right)$
$90b$	$\sqrt{123}$	$\left(\frac{7}{4}, \frac{1}{4}, -\frac{7}{2}\right)$	$95b$	$\sqrt{129}$	$\left(\frac{1}{4}, \frac{1}{4}, -4\right)$	$99c$	$\sqrt{135}$	$\left(\frac{7}{4}, \frac{5}{4}, -\frac{7}{2}\right)$
$90c$	$\sqrt{123}$	$\left(\frac{5}{4}, -\frac{5}{4}, \frac{7}{2}\right)$	$95c$	$\sqrt{129}$	$\left(\frac{11}{4}, \frac{11}{4}, -1\right)$	$99d$	$\sqrt{135}$	$\left(\frac{13}{4}, -\frac{1}{4}, -\frac{5}{2}\right)$
$90d$	$\sqrt{123}$	$\left(\frac{11}{4}, -\frac{11}{4}, \frac{1}{2}\right)$	$96a$	$\sqrt{131}$	$\left(\frac{9}{4}, -\frac{9}{4}, \frac{5}{2}\right)$	$100a'$	$\sqrt{136}$	$(3, 2, -2)$
$91a'$	$\sqrt{124}$	$\left(\frac{7}{2}, \frac{3}{2}, -1\right)$	$96b$	$\sqrt{131}$	$\left(\frac{15}{4}, \frac{1}{4}, -\frac{3}{2}\right)$	$100b'$	$\sqrt{136}$	$(4, 1, 0)$
$91b'$	$\sqrt{124}$	$\left(\frac{7}{2}, \frac{3}{2}, 1\right)$	$97a'$	$\sqrt{132}$	$\left(\frac{5}{2}, \frac{5}{2}, -2\right)$	$100c'$	$\sqrt{136}$	$(4, 0, -1)$
$91c'$	$\sqrt{124}$	$\left(\frac{7}{2}, 1, -\frac{3}{2}\right)$	$97b'$	$\sqrt{132}$	$\left(4, \frac{1}{2}, \frac{1}{2}\right)$	$100d'$	$\sqrt{136}$	$(3, 2, 2)$
$91d'$	$\sqrt{124}$	$\left(\frac{7}{2}, -1, -\frac{3}{2}\right)$	$97c'$	$\sqrt{132}$	$\left(4, \frac{1}{2}, -\frac{1}{2}\right)$	$100e'$	$\sqrt{136}$	$(3, -2, -2)$
$91e'$	$\sqrt{124}$	$\left(3, \frac{5}{2}, \frac{1}{2}\right)$	$97d'$	$\sqrt{132}$	$\left(4, -\frac{1}{2}, -\frac{1}{2}\right)$			
$91f'$	$\sqrt{124}$	$\left(3, \frac{5}{2}, -\frac{1}{2}\right)$	$97e'$	$\sqrt{132}$	$\left(\frac{7}{2}, 2, \frac{1}{2}\right)$			

Table B.5: Further neighbor symmetry classes up to 100'th nearest neighbor.

n	Δ_n	n	Δ_n	n	Δ_n	n	Δ_n
1	2.654×10^{1}	18a	1.188×10^{-2}	28b	-4.904×10^{-4}	39a	4.886×10^{-3}
2	-2.436×10^{-1}	18b	-3.619×10^{-2}	28c	-4.121×10^{-3}	39b	-1.625×10^{-3}
3a	1.194×10^{-1}	19a	-1.051×10^{-2}	29	5.285×10^{-4}	39c	1.357×10^{-3}
3b	-1.197×10^{-1}	19b	2.339×10^{-4}	30a	-8.698×10^{-4}	39d	-1.017×10^{-4}
4	-2.501×10^{-1}	19c	1.254×10^{-2}	30b	6.019×10^{-4}	39e	-8.975×10^{-4}
5	-1.38×10^{-1}	20a	8.566×10^{-3}	31a	1.624×10^{-3}	39f	-4.423×10^{-3}
6	-1.339×10^{-4}	20b	1.57×10^{-2}	31b	6.62×10^{-3}	40a	3.22×10^{-3}
7a	4.812×10^{-3}	20c	4.098×10^{-3}	32a	-3.979×10^{-3}	40b	-1.852×10^{-3}
7b	-6.796×10^{-2}	21a	1.967×10^{-2}	32b	-8.161×10^{-4}	41a	-2.161×10^{-3}
8	5.76×10^{-2}	21b	-6.877×10^{-3}	32c	5.297×10^{-3}	41b	-8.954×10^{-4}
9a	1.073×10^{-1}	21c	6.807×10^{-3}	33a	6.104×10^{-3}	41c	-1.667×10^{-3}
9b	-1.214×10^{-1}	21d	-2.014×10^{-2}	33b	-5.224×10^{-3}	42a	-1.017×10^{-3}
9c	-9.366×10^{-2}	22	1.982×10^{-2}	33c	-7.075×10^{-4}	42b	1.665×10^{-3}
10a	1.225×10^{-1}	23a	-8.445×10^{-3}	34a	1.654×10^{-3}	42c	-1.142×10^{-3}
10b	9.133×10^{-2}	23b	1.13×10^{-2}	34b	-3.055×10^{-3}	43a	-1.323×10^{-3}
11	-4.009×10^{-2}	24	-1.339×10^{-4}	34c	2.215×10^{-3}	43b	3.593×10^{-4}
12a	8.829×10^{-2}	25a	-4.11×10^{-3}	35a	4.239×10^{-3}	44a	1.258×10^{-3}
12b	-8.856×10^{-2}	25b	-8.643×10^{-3}	35b	2.427×10^{-3}	44b	-2.656×10^{-3}
13a	-3.395×10^{-2}	25c	-7.144×10^{-3}	36a	4.245×10^{-3}	44c	1.358×10^{-3}
13b	-3.775×10^{-2}	26	-5.988×10^{-3}	36b	-4.018×10^{-3}	44d	-4.959×10^{-4}
14a	5.52×10^{-3}	27a	4.096×10^{-3}	36c	-5.008×10^{-3}	45a	-3.391×10^{-4}
14b	-3.094×10^{-2}	27b	-1.066×10^{-2}	37a	-9.817×10^{-4}	45b	4.286×10^{-4}
15a	-1.972×10^{-2}	27c	1.93×10^{-3}	37b	4.006×10^{-3}	45c	-1.034×10^{-3}
15b	1.946×10^{-2}	27d	1.287×10^{-2}	37c	-9.622×10^{-4}	46a	-1.499×10^{-3}
16	-3.123×10^{-2}	27e	-1.314×10^{-2}	38a	6.187×10^{-5}	46b	1.487×10^{-3}
17	8.938×10^{-3}	28a	-3.973×10^{-3}	38b	-3.388×10^{-3}	46c	8.755×10^{-4}

Table B.6: The further-neighbor couplings appearing in the difference Δ between the magnetic dipolar interaction matrix and the projector, up to 100'th neighbor.

n	Δ_n	n	Δ_n	n	Δ_n	n	Δ_n
$47a$	2.508×10^{-3}	$55b$	5.248×10^{-4}	$63a$	-7.299×10^{-4}	$72a$	-4.676×10^{-4}
$47b$	-2.776×10^{-3}	$55c$	1.47×10^{-3}	$63b$	9.96×10^{-4}	$72b$	1.045×10^{-3}
$48a$	2.874×10^{-3}	$56a$	1.635×10^{-3}	$63c$	1.019×10^{-4}	$72c$	-3.789×10^{-4}
$48b$	-6.947×10^{-4}	$56b$	-1.423×10^{-3}	$63d$	7.448×10^{-4}	$72d$	4.252×10^{-4}
$49a$	1.682×10^{-4}	$56c$	2.582×10^{-4}	$64a$	-2.392×10^{-4}	$72e$	1.557×10^{-4}
$49b$	-1.013×10^{-3}	$56d$	3.707×10^{-4}	$64b$	1.049×10^{-4}	$72f$	6.383×10^{-4}
$49c$	4.493×10^{-4}	$56e$	-1.223×10^{-3}	$64c$	-3.926×10^{-4}	$73a$	8.584×10^{-4}
$49d$	1.371×10^{-3}	$56f$	6.392×10^{-4}	$65a$	6.245×10^{-4}	$73b$	-6.663×10^{-4}
$50a$	1.174×10^{-3}	$56g$	-1.587×10^{-3}	$65b$	-1.312×10^{-3}	$73c$	-4.352×10^{-4}
$50b$	4.667×10^{-4}	$57a$	5.562×10^{-4}	$65c$	-4.724×10^{-4}	$73d$	3.676×10^{-4}
$50c$	-7.345×10^{-4}	$57b$	-6.259×10^{-4}	$66a$	-2.528×10^{-4}	$74a$	7.183×10^{-4}
$50d$	-1.839×10^{-3}	$58a$	-6.211×10^{-4}	$66b$	-3.328×10^{-4}	$74b$	-3.596×10^{-4}
$50e$	-1.045×10^{-3}	$58b$	1.163×10^{-3}	$66c$	4.151×10^{-4}	$74c$	9.177×10^{-5}
$51a$	1.789×10^{-3}	$59a$	-5.867×10^{-4}	$67a$	-3.413×10^{-4}	$74d$	6.354×10^{-4}
$51b$	-6.624×10^{-4}	$59b$	3.189×10^{-4}	$67b$	9.979×10^{-4}	$74e$	-6.581×10^{-4}
52	5.006×10^{-4}	$60a$	-8.264×10^{-4}	$67c$	3.701×10^{-4}	$74f$	3.829×10^{-5}
$53a$	2.286×10^{-3}	$60b$	-1.697×10^{-4}	$68a$	4.92×10^{-4}	$74g$	-9.033×10^{-4}
$53b$	-1.339×10^{-4}	$60c$	-5.303×10^{-4}	$68b$	-6.76×10^{-4}	$74h$	-6.342×10^{-4}
$53c$	-8.923×10^{-4}	$60d$	4.441×10^{-5}	$68c$	1.429×10^{-4}	$75a$	1.702×10^{-4}
$53d$	-9.647×10^{-4}	$61a$	-6.772×10^{-4}	$68d$	-4.946×10^{-4}	$75b$	-1.603×10^{-4}
$54a$	4.139×10^{-4}	$61b$	-4.596×10^{-4}	$69a$	2.635×10^{-5}	$76a$	1.193×10^{-4}
$54b$	-1.407×10^{-3}	$61c$	-7.41×10^{-5}	$69b$	1.691×10^{-4}	$76b$	-3.599×10^{-4}
$54c$	-1.029×10^{-3}	$62a$	1.432×10^{-3}	$69c$	-2.889×10^{-4}	$76c$	-2.908×10^{-4}
$54d$	-6.525×10^{-4}	$62b$	9.515×10^{-5}	70	-2.049×10^{-4}	$76d$	1.275×10^{-4}
$54e$	2.21×10^{-4}	$62c$	-7.392×10^{-4}	$71a$	8.783×10^{-5}	$76e$	6.269×10^{-4}
$55a$	1.327×10^{-3}	$62d$	-1.324×10^{-3}	$71b$	-7.993×10^{-4}		

Table B.7: The further-neighbor couplings appearing in the difference Δ between the magnetic dipolar interaction matrix and the projector, up to 100'th neighbor.

n	Δ_n	n	Δ_n	n	Δ_n	n	Δ_n
$77a$	2.284×10^{-4}	$84a$	-3.79×10^{-4}	$91b$	-5.792×10^{-4}	$97f$	1.701×10^{-4}
$77b$	-4.963×10^{-4}	$84b$	1.729×10^{-4}	$91c$	1.416×10^{-4}	$97g$	-1.276×10^{-4}
$78a$	3.384×10^{-4}	$84c$	1.765×10^{-4}	$91d$	-3.407×10^{-4}	$97h$	-3.389×10^{-4}
$78b$	1.848×10^{-4}	$85a$	-2.846×10^{-4}	$91e$	3.038×10^{-7}	$97i$	-3.905×10^{-4}
$79a$	-7.526×10^{-5}	$85b$	3.858×10^{-4}	$91f$	4.499×10^{-4}	$97j$	-1.596×10^{-4}
$79b$	-1.425×10^{-4}	$85c$	3.692×10^{-5}	$91g$	-4.275×10^{-4}	$98a$	6.175×10^{-5}
$80a$	4.328×10^{-4}	$85d$	-1.618×10^{-4}	$91h$	-5.584×10^{-4}	$98b$	-2.021×10^{-4}
$80b$	-6.973×10^{-4}	$85e$	1.817×10^{-4}	$92a$	-4.237×10^{-5}	$98c$	-1.695×10^{-4}
$80c$	1.565×10^{-5}	$86a$	-2.039×10^{-4}	$92b$	2.772×10^{-4}	$98d$	1.345×10^{-4}
$80d$	1.637×10^{-4}	$86b$	-6.394×10^{-5}	$92c$	-2.091×10^{-4}	$98e$	2.384×10^{-4}
$80e$	-1.769×10^{-5}	$87a$	-3.104×10^{-4}	$92d$	2.613×10^{-4}	$99a$	2.913×10^{-4}
$80f$	-6.98×10^{-4}	$87b$	6.547×10^{-4}	$93a$	-4.946×10^{-5}	$99b$	-1.402×10^{-4}
$80g$	-7.032×10^{-4}	$87c$	-2.245×10^{-4}	$93b$	4.989×10^{-4}	$99c$	8.873×10^{-5}
$80h$	-2.627×10^{-4}	$87d$	-1.571×10^{-4}	$93c$	2.574×10^{-4}	$99d$	-5.913×10^{-5}
$80i$	8.306×10^{-4}	$87e$	4.046×10^{-4}	$93d$	-9.905×10^{-5}	$100a$	1.286×10^{-4}
$80j$	-5.518×10^{-4}	$88a$	-2.029×10^{-4}	94	-1.339×10^{-4}	$100b$	-2.615×10^{-4}
$81a$	6.045×10^{-4}	$88b$	-1.229×10^{-4}	$95a$	-1.69×10^{-4}	$100c$	-6.288×10^{-6}
$81b$	-3.099×10^{-4}	$89a$	-4.406×10^{-5}	$95b$	-3.291×10^{-4}	$100d$	-3.915×10^{-4}
$81c$	4.159×10^{-5}	$89b$	-2.082×10^{-5}	$95c$	-3.194×10^{-5}	$100e$	-4.015×10^{-4}
$82a$	4.596×10^{-4}	$89c$	-1.177×10^{-4}	$96a$	3.968×10^{-4}	$101a$	4.01×10^{-4}
$82b$	-4.71×10^{-6}	$89d$	1.043×10^{-4}	$96b$	-2.382×10^{-5}	$101b$	-1.101×10^{-4}
$82c$	-1.276×10^{-4}	$90a$	-2.289×10^{-4}	$97a$	1.739×10^{-4}	$101c$	4.145×10^{-4}
$83a$	6.14×10^{-4}	$90b$	1.387×10^{-4}	$97b$	-3.84×10^{-4}	$101d$	-2.786×10^{-5}
$83b$	-4.253×10^{-4}	$90c$	2.588×10^{-4}	$97c$	-7.098×10^{-5}		
$83c$	-5.963×10^{-5}	$90d$	-2.064×10^{-5}	$97d$	-9.536×10^{-6}		
$83d$	-6.649×10^{-4}	$91a$	2.425×10^{-4}	$97e$	-2.393×10^{-4}		

Table B.8: The further-neighbor couplings appearing in the difference Δ between the magnetic dipolar interaction matrix and the projector, up to 100'th neighbor.

Appendix C

The 1/N Expansion

As discussed in Chapter 2, the large-N approximation reproduces the correlations of the SCGA in the limit $N \to \infty$, and thus can serve as a justification for the use of the SCGA. On the other hand, the large-N method does not lend itself to anisotropic interactions directly, because it relies on an isotropic model in order to decouple to decouple the N components (and in any case it is unclear what is meant by "anisotropy" for spins with $N > 3$ components). The benefit of the large-N approach, however, is that it allows for a controlled perturbative expansion away from the $N = \infty$ limit in powers of $1/N$, which can be considered to be corrections to the SCGA. In this appendix we work out in detail the $1/N$ expansion, which is used in the main text.

C.1 Correlation Functions

Our primary interest is to compute the $1/N$ correction to the spin-spin correlation functions. To do so we first introduce an auxiliary field term to the partition function Eq. (2.16) in order to write it as

$$Z_{\text{Large-}N}[\boldsymbol{h}] = \int \mathcal{D}\mu\mathcal{D}\boldsymbol{S} \exp\left(\frac{S_N^2}{2}\sum_i \mu_i - \frac{1}{2}\sum_{ij}\mathcal{G}_{ij}^{-1}\boldsymbol{S}_i\cdot\boldsymbol{S}_j + \sum_i \boldsymbol{h}_i\cdot\boldsymbol{S}_j\right), \tag{C.1}$$

where $\mathcal{G}_{ij}^{-1} = \mu_i\delta_{ij} + \beta\mathcal{J}_{ij}$, such that the spin-spin correlation functions are computed from the logarithm of the partition function (i.e. βF where F is the free energy),

$$\langle S_i^\alpha S_j^\beta\rangle \equiv \frac{\partial}{\partial h_i^\alpha}\frac{\partial}{\partial h_j^\beta}\ln Z_{\text{Large-}N}[\boldsymbol{h}]\Big|_{\boldsymbol{h}=\boldsymbol{0}}. \tag{C.2}$$

We will expand the partition function diagrammatically as a series in powers of $1/N$. As is standard for expansions, the partition function generates all disconnected diagrams, and its logarithm generates all connected irreducible diagrams.

C.2 Saddle Point Expansion

We complete the square in the exponent by making a linear change of variables $\boldsymbol{\xi}_i \equiv \boldsymbol{S}_i - \sum_j \mathcal{G}_{ij}\boldsymbol{h}_j$, with which the partition function becomes

$$Z_{\text{Large-}N}[\boldsymbol{h}] = \int \mathcal{D}\mu\mathcal{D}\boldsymbol{\xi}\, \exp\left(\frac{S_N^2}{2}\sum_i \mu_i - \frac{1}{2}\sum_{ij}\mathcal{G}_{ij}^{-1}\boldsymbol{\xi}_i\cdot\boldsymbol{\xi}_j + \frac{1}{2}\sum_{ij}\mathcal{G}_{ij}\boldsymbol{h}_i\cdot\boldsymbol{h}_j\right). \tag{C.3}$$

Integrating over the Gaussian $\boldsymbol{\xi}$ variables we obtain

$$Z_{\text{Large-}N}[\boldsymbol{h}] = \int_{\lambda-i\infty}^{\lambda+i\infty}\prod_i \mathrm{d}\mu_i\, \exp\left(\frac{N}{2}\left[\sum_i \mu_i\frac{S^2}{N_t} + \operatorname{Tr}\ln[\mathcal{G}]\right] + \frac{1}{2}\sum_{ij}\mathcal{G}_{ij}\boldsymbol{h}_i\cdot\boldsymbol{h}_j\right). \tag{C.4}$$

We now expand around the saddle point as $\mu_i = \lambda - i\psi_i/\sqrt{N}$, where λ is determined by the saddle point condition and ψ is real.[1] We then expand $\mathcal{G}^{-1}$ as

$$\mathcal{G}_{ij}^{-1} = (\lambda\delta_{ij} + \beta\mathcal{J}_{ij}) - \frac{1}{\sqrt{N}}i\psi_i\delta_{ij} \equiv [\mathcal{G}_0^{-1}]_{ij} - i\Psi_{ij}, \tag{C.5}$$

[1] A simplified one-dimensional example of the calculation follows: consider an integral of the form $\int_{-\infty}^{\infty}\mathrm{d}x\, e^{Nf(x)}$, where f has a maximum at x_0 (ignoring convergence issues of the integral at infinity). The maximum of f dominates the integral for large N. Taylor expanding f about x_0, denoting its derivatives by $[\partial f/\partial x]_{x_0} \equiv f^{(k)}(x_0) \equiv f_0^{(k)}$, the linear term vanishes and we obtain

$$\int_{-\infty}^{\infty}\mathrm{d}x\, e^{Nf(x)} = e^{Nf(x_0)}\int_{-\infty}^{\infty}\mathrm{d}x\, \exp\left(-N\left|f_0^{(2)}\right|(x-x_0)^2/2\right)\exp\left[N\sum_{k=3}^{\infty}\frac{1}{3!}f_0^{(k)}(x-x_0)^k\right],$$

where we note that the second derivative is negative since we are expanding about a maximum. Expanding $x = x_0 + y/\sqrt{N}$, i.e. changing variables to $y = \sqrt{N}(x-x_0)$, we obtain

$$\int_{-\infty}^{\infty}\mathrm{d}x\, e^{Nf(x)} = \frac{1}{\sqrt{N}}e^{Nf(x_0)}\int_{-\infty}^{\infty}\mathrm{d}y\, e^{-y^2/2\left|f_0^{(2)}\right|^{-1}}\sum_{n=0}^{\infty}\frac{1}{n!}\left(\sum_{k=3}^{\infty}\frac{1}{k!}\frac{f_0^{(k)}}{\sqrt{N}^{k-2}}y^k\right)^n$$

The higher order terms can then be integrated against the Gaussian of variance $|f_0^{(2)}|^{-1}$. This perturbative series in powers of $1/\sqrt{N}$ may in general fail to converge but it will at least be asymptotic, meaning that its first few terms contain valuable data.

where we define

$$\Psi_{ij} \equiv \psi_i \delta_{ij}/\sqrt{N}. \tag{C.6}$$

The matrix $\mathcal{G}_0$ is the spin-spin correlation matrix evaluated at the saddle point, and can be thought of as the "non-interacting propagator". We can then expand $\ln[\mathcal{G}]$ as

$$\ln \mathcal{G} = \ln[(\mathcal{G}_0^{-1} - i\Psi)^{-1}] = \ln[(\mathbb{1} - i\mathcal{G}_0\Psi)^{-1}\mathcal{G}_0] = \ln \mathcal{G}_0 + \ln[(\mathbb{1} - i\mathcal{G}_0\Psi)^{-1}], \tag{C.7}$$

and

$$\sum_i \mu_i = N_{\text{spins}}\lambda - \sum_i \frac{1}{\sqrt{N}} i\psi_i = N_{\text{spins}}\lambda - i \operatorname{Tr} \Psi. \tag{C.8}$$

Subbing Eqs. (C.7) and (C.8) into Eq. (C.4), we obtain

$$Z_{\text{Large-}N}[\boldsymbol{h}] = Z_\infty \int_{-\infty}^{\infty} \prod_i \mathrm{d}\psi_i \, \exp\left(\frac{N}{2}\left[-\frac{S^2}{N_t} i \operatorname{Tr} \Psi + \operatorname{Tr} \ln[(\mathbb{1} - i\mathcal{G}_0\Psi)^{-1}]\right] + \frac{1}{2}\sum_{ij} \mathcal{G}_{ij} \boldsymbol{h}_i \cdot \boldsymbol{h}_j\right), \tag{C.9}$$

where

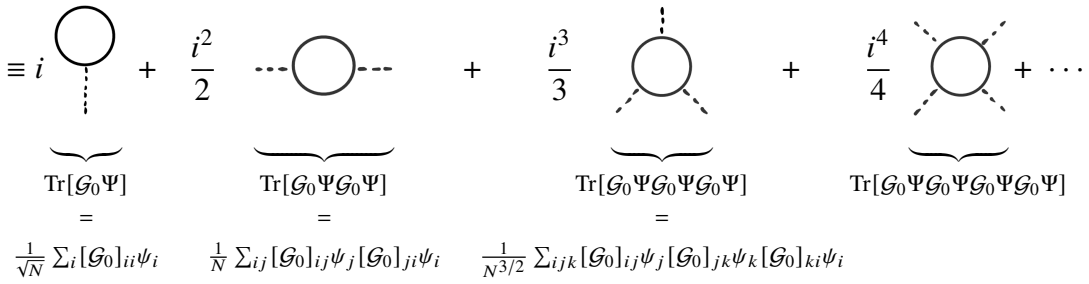
$$Z_\infty \equiv \exp\left(\frac{N}{2}\left[N_{\text{spins}}\lambda \frac{S^2}{N_t} + \operatorname{Tr} \ln \mathcal{G}_0\right]\right) \tag{C.10}$$

is the partition function evaluated at the saddle point. The next task is to expand the logarithm using the expansion

$$\ln\left(\frac{1}{1-x}\right) = \sum_{n=1}^{\infty} \frac{x^n}{n} \tag{C.11}$$

with $x = i\mathcal{G}_0\Psi$. We will do so diagrammatically by representing each $[\mathcal{G}_0]_{ij}$ as a solid line from i to j, and each $\Psi_{ij} = \delta_{ij}\psi_j/\sqrt{N}$ by a dashed line emanating from a single point i, to write the following expansion:

$$\operatorname{Tr} \ln[(\mathbb{1} - i\mathcal{G}_0\Psi)^{-1}] = \sum_{n=1}^{\infty} \frac{1}{n} \operatorname{Tr}[(i\mathcal{G}_0\Psi)^n]$$



(C.12)

In these diagrams, each solid line segment represents a $\mathcal{G}_0$, each external dashed line segment represents one ψ_i, and each three-point vertex carries a weight $1/\sqrt{N}$. Consider the tadpole diagram in the first term, whose value is given by



$$= \mathrm{Tr}[i\mathcal{G}_0\Psi] = \frac{1}{\sqrt{N}}\sum_{ij}[\mathcal{G}_0]_{ij}\psi_i\delta_{ji}, \tag{C.13}$$

$$= \frac{1}{\sqrt{N}}\sum_{i}[\mathcal{G}_0]_{ii}\psi_i, \tag{C.14}$$

$$= \frac{S^2}{N_t}\sum_{i}\frac{1}{\sqrt{N}}\psi_i, \tag{C.15}$$

$$= \frac{S^2}{N_t}\,\mathrm{Tr}\,\Psi, \tag{C.16}$$

which exactly cancels the linear term in Eq. (C.4) (this is to be expected since we are expanding about an extremum, i.e. the term linear in ψ cancels). The second term is the most important one, as it defines the propagator for ψ,

$$\frac{1}{2}\text{---}\bigcirc\text{---} = \frac{1}{2}\,\mathrm{Tr}[\mathcal{G}_0\Psi\mathcal{G}_0\Psi] = \frac{1}{2N}\sum_{ij}[\mathcal{G}_0]_{ij}\psi_j[\mathcal{G}_0]_{ji}\psi_i = \frac{1}{N}\sum_{ij}\psi_i\left[\frac{1}{2}\mathcal{G}_0\circ\mathcal{G}_0^{\mathrm{T}}\right]_{ij}\psi_j. \tag{C.17}$$

In the final expression, $\circ$ is the Hadamard or elementwise matrix product, $[A\circ B]_{ij}\equiv A_{ij}B_{ij}$. Note that the Hadamard product is *not* basis-independent, and this expression must be interpreted in the position basis where it is defined. We denote this matrix by $\mathcal{G}^{-1}$,

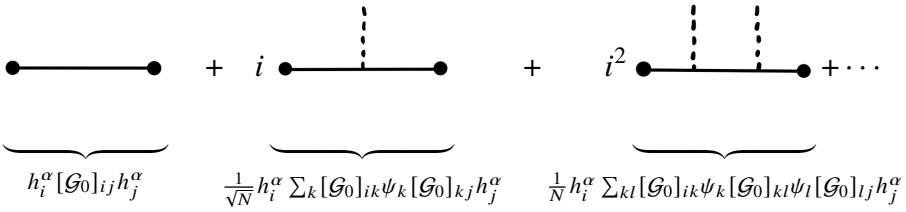
$$\mathcal{G}^{-1}_{ij} \equiv \frac{1}{2}\left[\mathcal{G}_0\circ\mathcal{G}_0^{\mathrm{T}}\right]_{ij} = \frac{1}{2}[\mathcal{G}_0]^2_{ij}, \tag{C.18}$$

where $\mathcal{G}$ is the covariance matrix (propagator) for ψ at the Gaussian level, and we utilized that $\mathcal{G}_0$ is a symmetric matrix since $\mathcal{J}$ is. Note that the $1/N$ in Eq. (C.17) cancels with the N in the exponent of Eq. (C.4). The partition function now reads

$$\frac{Z_{\text{Large-}N}[\boldsymbol{h}]}{Z_\infty} = \int\mathcal{D}\psi\,\exp\left(-\frac{1}{2}\sum_{ij}\psi_i\mathcal{G}^{-1}_{ij}\psi_j\right)\exp\left(\frac{N}{2}\sum_{n=3}^{\infty}\frac{1}{n}\,\mathrm{Tr}[(i\mathcal{G}_0\Psi)^n]\right)\exp\left(\frac{1}{2}\sum_{ij}\mathcal{G}_{ij}\boldsymbol{h}_i\cdot\boldsymbol{h}_j\right) \tag{C.19}$$

where $\mathcal{G}_{ij}$ is a function of ψ. To complete the diagrammatic expansion, we now expand $\mathcal{G}$ as

$$\mathcal{G} = (\mathcal{G}_0 - i\Psi)^{-1} = (\mathbb{1} - i\mathcal{G}_0\Psi)^{-1}\mathcal{G}_0 = \sum_{n=0}^{\infty}(i\mathcal{G}_0\Psi)^n\mathcal{G}_0, \tag{C.20}$$

$\mathcal{G}_{ij}h_i^\alpha h_j^\alpha =$  (C.21)

In this expression each solid line segment again represents a $\mathcal{G}_0$, each dashed line a ψ_i, each 3-point (two solid and one dashed line) vertex carries weight $1/\sqrt{N}$, and the external "bubbles" represent the auxiliary fields h_i^α, with an implicit sum over α. Recall that we had to assume for this expansion to make sense that the interaction matrix is isotropic, so there are no propagators connecting h_i^α and h_j^β for $\alpha \neq \beta$.

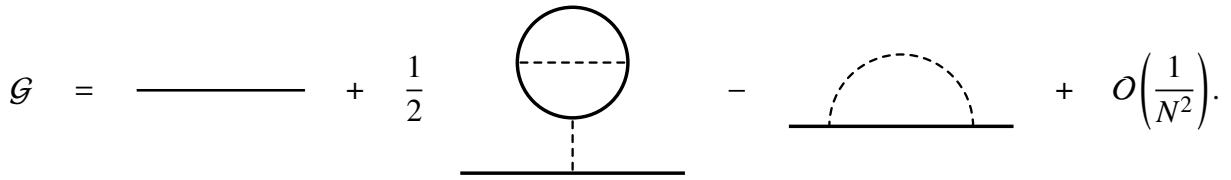
Diagrammatically, the partition function can now be formally expanded as the infinite sum of all disconnected diagrams,

$$Z_{\text{Large-}N}[\boldsymbol{h}] = Z_\infty \int \mathcal{D}\psi \exp\left(-\frac{1}{2}\sum_{ij}\psi_i\mathcal{G}_{ij}^{-1}\psi_j\right)$$
$$\times \sum_{v=0}^{\infty}\frac{1}{v!}\frac{1}{2^v}\left(\frac{i^3}{3}N \; + \frac{i^4}{4}N \; + \cdots\right)^v$$
$$\times \sum_{p=0}^{\infty}\frac{1}{p!}\frac{1}{2^p}\left(\sum_\alpha\left[\; + i \; + i^2 \; + \cdots\right]\right)^p .$$
(C.22)

Note that each vertex is multiplied by a factor N, making the k-point vertex order $(1/\sqrt{N})^{k-2}$, while the propagators with k dashed lines are of order $(1/\sqrt{N})^k$. The integration over ψ is performed by applying Wick's theorem, contracting all dashed lines in all possible ways in each diagram. Diagrams with an odd number of dashed lines before integration vanish, since their corresponding integrands are odd functions of the ψ_i. Since each dashed line comes with a factor of i this means that all of the non-zero terms are real, and diagrams with $2k$ dashed lines (before applying Wick's theorem, k dashed lines after applying Wick's theorem) carry an overall sign

$(-1)^k$. In the final diagrams obtained after connecting dashed lines, the order in $1/N$ is the number of dashed lines minus the number of solid circles.

The normalized spin-spin correlation functions are computed from Eq. (C.2). The partition function is the sum of all disconnected diagrams, while its logarithm is the sum of all connected diagrams. Diagrams with two external sources, originating from the $p = 1$ term in the last line, contribute to the two-spin correlators. The p'th term contributes to correlations of $2p$ spins. The zeroth order propagator is $\mathcal{G}_0$, represented by the solid line. The order $1/N$ correction contains two contributions, one coming the 1-point propagator combined with the 3-point vertex ($p = 1, v = 1$) and one connecting the two external ψ's in the 2-point propagator ($p = 1, v = 0$), given by the following diagrams:



(C.23)

The first term is the "bare propagator" $\mathcal{G}_0$. The second term acts as a correction to the saddle point value of λ, with the $1/2$ coming from the $1/2^v$ in Eq. (C.22) with $v = 1$. The factor of $1/3$ is canceled by the three ways to connect the three-point vertex to the 1-point propagator. The third term is a "radiative correction" which will give us momentum-dependent corrections to the correlation functions. Note the negative sign coming from the i^2 in Eq. (C.21). We note that $\mathrm{Tr}\,\mathcal{G} = \mathrm{Tr}\,\mathcal{G}_0$, i.e. when we trace the two diagrams at order $1/N$ they cancel each other.

C.3 Diagrammatic Corrections to the Correlator

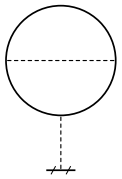
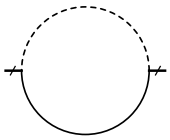
What we actually need to compute is the "self energy", which amounts to resumming a subset of diagrams which form a geometric series. Other than the first term (the bare propagator), every remaining term has two external bare propagators, and the series has the form of Dyson's equation

$$\mathcal{G} = \mathcal{G}_0 + \mathcal{G}_0\Sigma\mathcal{G}_0 + \mathcal{G}_0\Sigma\mathcal{G}_0\Sigma\mathcal{G}_0 + \cdots = \mathcal{G}_0 + \mathcal{G}_0\Sigma\mathcal{G}. \tag{C.24}$$

Solving for $\mathcal{G}$, we obtain

$$\mathcal{G} = (\mathbb{1} - \mathcal{G}_0\Sigma)^{-1}\mathcal{G}_0 = (\mathcal{G}_0^{-1} - \Sigma)^{-1}. \tag{C.25}$$

We thus have to order $1/N$,

$$\Sigma \;=\; \frac{1}{2}$$  $$-$$  $$+\; O\left(\frac{1}{N^2}\right), \tag{C.26}$$

where we have removed the external propagators on the diagrams in Eq. (C.23). The diagrams above can be easily interpreted in the position basis. Each 3-point vertex corresponds to a lattice site i, each solid line to a $[\mathcal{G}_0]_{ij}$, and each dashed line to a $\mathcal{G}_{ij}$. All internal vertices (not containing an exterior "amputated" solid line) is summed over. Therefore Eq. (C.26) reads

$$N\Sigma_{ij} = \delta_{ij}\frac{1}{2}\left(\sum_{k,l,m}\mathcal{G}_{ik}[\mathcal{G}_0]_{kl}[\mathcal{G}_0]_{lm}\mathcal{G}_{lm}[\mathcal{G}_0]_{mk}\right) - [\mathcal{G}_0]_{ij}\mathcal{G}_{ij} + O\left(N^{-1}\right). \tag{C.27}$$

We can easily write down the same diagrams in the Fourier basis: the vertices then correspond to sublattice positions and we need only sum over all internal sublattices and momenta, and conserve momentum at each vertex. Therefore we have[2]

$$N\Sigma_{\mu\nu}(\boldsymbol{q}) = \delta_{\mu\nu}\frac{1}{2}\frac{1}{(L^3)^2}\sum_{\boldsymbol{Q},\boldsymbol{Q}'}\mathcal{G}_{\mu\rho}(\boldsymbol{0})[\mathcal{G}_0(\boldsymbol{Q})]_{\rho\sigma}[\mathcal{G}_0(\boldsymbol{Q}-\boldsymbol{Q}')]_{\sigma\lambda}\mathcal{G}_{\sigma\lambda}(\boldsymbol{Q}')[\mathcal{G}_0(\boldsymbol{Q})]_{\lambda\rho}$$
$$-\frac{1}{L^3}\sum_{\boldsymbol{Q}}[\mathcal{G}_0(\boldsymbol{q}+\boldsymbol{Q})]_{\mu\nu}\mathcal{G}(-\boldsymbol{Q})_{\mu\nu}. \tag{C.28}$$

The auxiliary field propagator is defined by Eq. (C.18), with its inverse given by

$$\mathcal{G}^{-1}_{\mu\nu}(\boldsymbol{q}) = \frac{1}{2}\frac{1}{L^3}\sum_{\boldsymbol{Q}}[\mathcal{G}_0(\boldsymbol{q}+\boldsymbol{Q})]_{\mu\nu}[\mathcal{G}_0(-\boldsymbol{Q})]_{\mu\nu}. \tag{C.29}$$

In principle the expansion can be continued, but the scaling of the computation of each diagram at order N^{-k} will be $(L^3)^{k+1}$, and the number of diagrams also grows exponentially.

[2]For this calculation we should work with the Fourier basis without sublattice phase factors so that $\boldsymbol{G}_0(\boldsymbol{q})$ is periodic in the FBZ. Each sum is then over the L^3 momenta in the FBZ.

Appendix D

Derivations for N=3 Spins

In this appendix we provide the derivations for results utilized in Chapter 6. In particular, we provide the full derivation of the classical low temperature expansion to obtain the Hessian for small fluctuations about a classically ordered ground state. We then perform the classical linear spin wave expansion of the equations of motion for spin precession, to rephrase it as a generalized eigenvalue problem.

D.1 Classical Low Temperature Expansion

We begin from a generic spin Hamiltonian

$$H = \frac{1}{2} \sum_{ij} \sum_{\alpha\beta} S_i^\alpha \mathcal{J}_{ij}^{\alpha\beta} S_j^\beta. \tag{D.1}$$

Let $\boldsymbol{Z}_i$ denote the ground state configuration which we wish to expand about. We then separate each $\boldsymbol{S}_i$ into its components along and orthogonal to the ordering axis,

$$\boldsymbol{S}_i \equiv \boldsymbol{S}_i^\perp + (\boldsymbol{S}_i \cdot \boldsymbol{Z}_i)\boldsymbol{Z}_i \equiv \boldsymbol{S}_i^\perp + S_i^z \boldsymbol{Z}_i.$$

Putting this into Eq. (D.1), we obtain

$$\begin{aligned} H &= \frac{1}{2} \sum_{ij} \sum_{\alpha\beta} \left[(S_i^\perp)^\alpha + S_i^z Z_i^\alpha \right] \mathcal{J}_{ij}^{\alpha\beta} \left[(S_j^\perp)^\beta + S_j^z Z_j^\beta \right], \\ &\simeq \frac{1}{2} \sum_{ij} \sum_{\alpha\beta} \left[(S_i^\perp)^\alpha \mathcal{J}_{ij}^{\alpha\beta} (S_j^\perp)^\beta + S_i^z S_j^z Z_i^\alpha \mathcal{J}_{ij}^{\alpha\beta} Z_j^\beta \right], \end{aligned} \tag{D.2}$$

where we drop any terms which are odd in $\boldsymbol{S}_i^\perp$ since we are expanding about a minimum of H to quadratic order. Enforcing the spin length constraint $|\boldsymbol{S}_i|^2 = |\boldsymbol{S}_i^\perp|^2 + (S_i^z)^2 = 1$, we can expand S_i^z as

$$S_i^z = 1 - \frac{1}{2}|\boldsymbol{S}_i^\perp|^2 + O\Big(|\boldsymbol{S}_i^\perp|^4\Big). \tag{D.3}$$

Substituting into Eq. (D.2), we obtain

$$H = E_0 + \frac{1}{2}\sum_{ij}\sum_{\alpha\beta}(S_i^\perp)^\alpha \mathcal{H}_{ij}^{\alpha\beta}(S_j^\perp)^\beta + O\Big(|\boldsymbol{S}_i^\perp|^3\Big), \tag{D.4}$$

where

$$E_0 \equiv \frac{1}{2}\sum_{ij,\alpha\beta} Z_i^\alpha \mathcal{J}_{ij}^{\alpha\beta} Z_j^\beta, \tag{D.5}$$

is the ground state energy, and

$$\mathcal{H}_{ij}^{\alpha\beta} \equiv \mathcal{J}_{ij}^{\alpha\beta} + \Delta_i \delta_{ij}\delta^{\alpha\beta}, \tag{D.6}$$

is the Hessian in the global basis, with

$$\Delta_i \equiv -\sum_j \sum_{\alpha\beta} Z_i^\alpha \mathcal{J}_{ij}^{\alpha\beta} Z_j^\beta, \tag{D.7}$$

where we have used that $\mathcal{J}$ is symmetric with respect simultaneously swapping $i \leftrightarrow j$ and $\alpha \leftrightarrow \beta$. Δ_i is the mean field energy cost of reducing the ordered moment S_i^z. The effective Hamiltonian for the fluctuating $\boldsymbol{S}_i^\perp$ is then given by

$$H_{\text{eff}} = \frac{1}{2}\sum_{ij}\sum_{\alpha\beta}(S_i^\perp)^\alpha \mathcal{H}_{ij}^{\alpha\beta}(S_j^\perp)^\beta. \tag{D.8}$$

A convenient compact representation of the Hessian is the following; let

$$P_{ij}^{\alpha\beta} = \delta_{ij}(\delta_{\alpha\beta} - Z_i^\alpha Z_i^\beta) \tag{D.9}$$

be the matrix which projects out the spin components along the ordering axes $\boldsymbol{Z}_i$, then the Hessian can be written as

$$\mathcal{H}_{ij}^{\alpha\beta} = [P\mathcal{J}P]_{ij}^{\alpha\beta} + \Delta_i P_{ij}^{\alpha\beta}. \tag{D.10}$$

In this form, it is clear that $\mathcal{H}$ annihilates any vector S with components of the form $S_i^\alpha = c_i Z_i^\alpha$. This means that 1/3 of the eigenvalues of $\mathcal{H}$ are zero, corresponding to the fact that we projected out the ordered spin components since the Hessian only describes the fluctuations of the transverse components. To obtain the physically meaningful portion of the Hessian we should restrict to the

subspace of transverse spin components. Let $\{\boldsymbol{\eta}_i^x, \boldsymbol{\eta}_i^y\}$ denote an orthonormal basis spanning the local plane orthogonal to $\boldsymbol{Z}_i$. Then the redundant modes can be removed by defining

$$\tilde{\mathcal{H}}_{ij}^{ab} \equiv \sum_{\alpha\beta} (\eta_i^a)^\alpha \mathcal{H}_{ij}^{\alpha\beta} (\eta_j^b)^\beta, \tag{D.11}$$

$$= \sum_{\alpha\beta} \left[(\eta_i^a)^\alpha \mathcal{J}_{ij}^{\alpha\beta} (\eta_j^b)^\beta \right] \Delta_i \delta_{ij} \delta^{ab} \tag{D.12}$$

where we use a tilde to denote reduced matrices which only act on transverse spin components. If the ground state is a stable minimum (i.e. there are no continuous zero modes) then the eigenvalues of $\tilde{\mathcal{H}}$ are all positive.

D.1.1 Classical Low Temperature Expansion for the Easy Axis Model

The easy axis model defined by Eq. (6.34) has interaction matrix (cf. Eq. (D.1), and note the factor of 1/2)

$$\mathcal{J}_{ij}^{\alpha\beta} = J\mathcal{A}_{ij}\delta^{\alpha\beta} - D\delta_{ij}\hat{z}_i^\alpha \hat{z}_j^\beta, \tag{D.13}$$

where $\hat{\boldsymbol{z}}_i$ is the local [111] vector (we take the convention that all $\boldsymbol{z}_i$ point from diamond sublattice A to diamond sublattice B), and $\mathcal{A}_{ij}$ is the adjacency matrix,

$$\mathcal{A}_{ij} \equiv \begin{cases} 1 & i, j \text{ nearest neighbors,} \\ 0 & \text{otherwise.} \end{cases}$$

The AIAO ground state has ordering vectors $\boldsymbol{Z}_i = \hat{\boldsymbol{z}}_i$. We can compute Δ_i from Eq. (D.7),

$$\Delta_i = -\sum_j \sum_{\alpha\beta} \hat{z}_i^\alpha (J\mathcal{A}_{ij}\delta^{\alpha\beta} - D\delta_{ij}\hat{z}_i^\alpha \hat{z}_j^\beta)\hat{z}_j^\beta \,. \tag{D.14}$$

Using that fact that $\hat{\boldsymbol{z}}_i \cdot \hat{\boldsymbol{z}}_j = -1/3$ for nearest-neighbors, and that there are six neighbors per site, we obtain

$$\Delta_i = 2J + D, \tag{D.15}$$

independent of i, and the Hessian Eq. (D.6) is

$$\mathcal{H}_{ij}^{\alpha\beta} = J(\mathcal{A}_{ij} + 2\delta_{ij})\delta^{\alpha\beta} + D\delta_{ij}\delta^{\alpha\beta}.$$

Note appearance of $\mathcal{A}_{ij} + 2\delta_{ij}$, discussed extensively in Chapter 3 for $N = 1$ component spins. Plugging this expression into the effective Hamiltonian Eq. (D.8), we obtain

$$H_{\text{eff}} = \frac{J}{2} \sum_t \Big| \sum_{i \in t} \boldsymbol{S}_i^\perp \Big|^2 + \frac{D}{2} \sum_i |\boldsymbol{S}_i^\perp|^2.$$

Note that the simplicity of this Hamiltonian is slightly deceiving, because we have expressed it in the global basis where we must additionally impose the anisotropic constraint that the $\boldsymbol{S}_i^\perp$ lie in the local xy-planes. This anisotropy is reflected in the complex band structure of the Hessian in Fig. 6.5.

D.2 Classical Spin Waves

The semi-classical dynamics for the classical spin variables $\boldsymbol{S}_i$ is governed by the torque equation,

$$\frac{\mathrm{d}}{\mathrm{d}t}\boldsymbol{S}_i = -\boldsymbol{S}_i \times \boldsymbol{B}_i, \tag{D.16}$$

where $\boldsymbol{B}_i$ is the local molecular field at site i, which is determined from the Hamiltonian by writing it in the form

$$H = -\sum_i \boldsymbol{S}_i \cdot \boldsymbol{B}_i, \tag{D.17}$$

implying that

$$\boldsymbol{B}_i \equiv -\frac{\partial H}{\partial \boldsymbol{S}_i}. \tag{D.18}$$

For the generic bilinear Hamiltonian Eq. (6.8), this means that

$$B_i^\alpha = -\sum_{j,\beta} \mathcal{J}_{ij}^{\alpha\beta} S_j^\beta. \tag{D.19}$$

Utilizing the same notation as in Appendix D.1, we expand $\boldsymbol{S}_i \equiv \boldsymbol{S}_i^\perp + S_i^z \boldsymbol{Z}_i$ in Eq. (D.16), from which we obtain

$$\begin{aligned} -\frac{\mathrm{d}}{\mathrm{d}t}(\boldsymbol{S}_i^\perp + S_i^z \boldsymbol{Z}_i) = \boldsymbol{S}_i^\perp \times \boldsymbol{B}_i^\perp + \boldsymbol{S}_i^\perp \times \boldsymbol{B}_i^z \\ + S_i^z \boldsymbol{Z}_i \times \boldsymbol{B}_i^\perp + S_i^z \boldsymbol{Z}_i \times \boldsymbol{B}_i^z, \end{aligned} \tag{D.20}$$

where

$$(B_i^\perp)^\alpha \equiv -\sum_{j,\beta} \mathcal{J}_{ij}^{\alpha\beta} (S_j^\perp)^\beta, \tag{D.21}$$

$$(B_i^z)^\alpha \equiv -\sum_{j,\beta} \mathcal{J}_{ij}^{\alpha\beta} (S_j^z Z_j^\beta), \tag{D.22}$$

are the contributions to the molecular field at site i from the fluctuating and ordered components of neighboring spins, respectively.

There are four terms appearing on the right hand side of Eq. (D.20). The first is quadratic in the fluctuating $S^\perp$ components, containing terms of the form $(S_i^\perp)^\alpha (S_j^\perp)^\beta$, giving rise to interactions between spin fluctuations at different sites and thus spin wave interactions. We ignore this term in the linear, non-interacting approximation. The last term is a constant independent of the $S^\perp$ components, and is zero in the ground state. This can be seen from equation Eq. (D.17), which implies that in the ground state (where all $\boldsymbol{S}_i = \boldsymbol{Z}_i$) each $\boldsymbol{Z}_i$ is *parallel* to the molecular field of its neighbors, $\boldsymbol{B}_i^z$ — thus their cross product is zero. The remaining two terms are linear in the $S^\perp$ components.

On the left hand side there are two terms. In the linear approximation, we ignore the fluctuations of the amplitude of the ordered moment S_i^z (the so-called "Higgs mode"). Indeed by expanding $S_i^z = 1 - |\boldsymbol{S}_i^\perp|^2/2 + \cdots$ we can see that this term is quadratic in $S^\perp$ components. For this same reason, we can set all of the $S_i^z = 1$ in the remaining two terms on the right hand side. We are left with the approximation

$$\frac{\mathrm{d}}{\mathrm{d}t}\boldsymbol{S}_i^\perp \simeq -\boldsymbol{S}_i^\perp \times \boldsymbol{B}_i^z - \boldsymbol{Z}_i \times \boldsymbol{B}_i^\perp, \tag{D.23}$$

where $\boldsymbol{B}_i^z$ is given by Eq. (D.22) with $S_j^z = 1$.

Eq. (D.23) is written in terms of 3-component vectors $\boldsymbol{S}_i^\perp$, which are each confined to lie in their local xy-planes. We now introduce a local orthonormal basis $\{\boldsymbol{\eta}_i^x, \boldsymbol{\eta}_i^y\}$ for the local xy-plane for each i, and read off the dynamics of the local $(S_i^\perp)^a \equiv \boldsymbol{S}_i^\perp \cdot \boldsymbol{\eta}_i^a$ components by dotting each side of Eq. (D.23) with $\boldsymbol{\eta}_i^a$,

$$\frac{\mathrm{d}}{\mathrm{d}t}(\boldsymbol{\eta}_i^a \cdot \boldsymbol{S}_i^\perp) = -\boldsymbol{\eta}_i^a \cdot (\boldsymbol{S}_i^\perp \times \boldsymbol{B}_i^z) - \boldsymbol{\eta}_i^a \cdot (\boldsymbol{Z}_i \times \boldsymbol{B}_i^\perp), \tag{D.24}$$

$$\frac{\mathrm{d}}{\mathrm{d}t}(S_i^\perp)^a = -(\boldsymbol{\eta}_i^a \times \boldsymbol{S}_i^\perp) \cdot \boldsymbol{B}_i^z - (\boldsymbol{\eta}_i^a \times \boldsymbol{Z}_i) \cdot \boldsymbol{B}_i^\perp. \tag{D.25}$$

where in the second line we cycled the triple product. We will now work with the two terms on the right hand side of Eq. (D.25) separately. In the first term, we expand $\boldsymbol{S}_i^\perp = \sum_b (S_i^\perp)^b \boldsymbol{\eta}_i^b$,

$$\boldsymbol{\eta}_i^a \times \boldsymbol{S}_i^\perp = \sum_b (\boldsymbol{\eta}_i^a \times \boldsymbol{\eta}_i^b)(S_i^\perp)^b \equiv \sum_b (\epsilon^{ab} \boldsymbol{Z}_i)(S_i^\perp)^b, \tag{D.26}$$

where $\epsilon^{xy} = -\epsilon^{yx} = 1$ and $\epsilon^{xx} = \epsilon^{yy} = 0$. In the second term, we use

$$\boldsymbol{\eta}_i^a \times \boldsymbol{Z}_i = -\sum_b \epsilon^{ab} \boldsymbol{\eta}_i^b. \tag{D.27}$$

Plugging Eqs. (D.26) and (D.27) back into Eq. (D.25), we have

$$\frac{\mathrm{d}}{\mathrm{d}t}(S_i^\perp)^a = \sum_b \epsilon^{ab}\Big[-(\boldsymbol{Z}_i \cdot \boldsymbol{B}_i^z)(S_i^\perp)^b + \boldsymbol{\eta}_i^b \cdot \boldsymbol{B}_i^\perp \Big]. \tag{D.28}$$

In the first term, using Eq. (D.22) with $S_j^z \approx 1$ we have

$$\boldsymbol{Z}_i \cdot \boldsymbol{B}_i^z \simeq -\sum_j \sum_{\alpha\beta} Z_i^\alpha \mathcal{J}_{ij}^{\alpha\beta} Z_j^\beta \equiv \Delta_i, \tag{D.29}$$

where Δ_i was previously defined in Eq. (D.7). In the second term, using Eq. (D.21) and expanding $\boldsymbol{S}_j^\perp = \sum_c (\boldsymbol{S}_j^\perp)^c \boldsymbol{\eta}_i^c$, we obtain

$$\begin{aligned} \boldsymbol{\eta}_i^b \cdot \boldsymbol{B}_i^\perp &= -\sum_j \sum_c \left(\sum_{\alpha\beta} (\eta_i^b)^\alpha \mathcal{J}_{ij}^{\alpha\beta} (\eta_j^c)^\beta \right) (S_j^\perp)^c, \\ &\equiv -\sum_{j,c} \tilde{\mathcal{J}}_{ij}^{bc} (S_j^\perp)^c. \end{aligned} \tag{D.30}$$

Finally, plugging Eqs. (D.29) and (D.30) back into Eq. (D.28), we have

$$\frac{\mathrm{d}}{\mathrm{d}t}(S_i^\perp)^a = -\sum_b \epsilon^{ab}\Big[\Delta_i (S_i^\perp)^b + \sum_{j,c} \tilde{\mathcal{J}}_{ij}^{bc} (S_j^\perp)^c \Big]. \tag{D.31}$$

Let $\tilde{\Delta}_{ij}^{ab} \equiv \Delta_i \delta_{ij} \delta^{ab}$, to obtain

$$\frac{\mathrm{d}}{\mathrm{d}t}(S_i^\perp)^a = -\sum_{j,bc} \epsilon^{ab}\Big[\tilde{\Delta}_{ij}^{bc} + \tilde{\mathcal{J}}_{ij}^{bc} \Big] (S_j^\perp)^c, \tag{D.32}$$

$$\equiv -\sum_{j,bc} \epsilon^{ab} \tilde{\mathcal{H}}_{ij}^{bc} (S_j^\perp)^c, \tag{D.33}$$

where $\tilde{\mathcal{H}}$ is the Hessian matrix defined in Eq. (D.12). Lastly, let

$$\Gamma_{ij}^{ab} \equiv -\delta_{ij} \epsilon^{ab}, \tag{D.34}$$

in terms of which we obtain

$$\frac{\mathrm{d}}{\mathrm{d}t}(S_i^\perp)^a = -\sum_{jk,bc} \Gamma_{ij}^{ab} \tilde{\mathcal{H}}_{jk}^{bc} (S_j^\perp)^c. \tag{D.35}$$

We seek periodic solutions of the form

$$(S_i^{\perp})^a(t) = \psi_i^a e^{-\mathrm{i}\omega t}, \tag{D.36}$$

in terms of which Eq. (D.35) takes the form of an eigenvalue equation

$$\Gamma\tilde{\mathcal{H}}\psi = i\omega\psi, \tag{D.37}$$

or equivalently the generalized eigenvalue equation,

$$(\tilde{\mathcal{H}} + i\omega\Gamma)\psi = 0, \tag{D.38}$$

where we used $\Gamma^2 = -\mathbb{1}$.

D.3 HHD Norm

Given a configuration of the spin components, described by a vector $S = (S_1^x, S_1^y, S_1^z, S_2^x, S_2^y, \cdots)^{\mathrm{T}}$, we can apply the HHD projectors to it to obtain the rotational, irrotational, and harmonic components components (cf. Eqs. (6.29) to (6.31)),

$$S = \mathcal{P}^{\mathrm{rot}}S + \mathcal{P}^{\mathrm{irr}}S + \mathcal{P}^{\mathrm{har}}S \tag{D.39}$$

We introduce a metric to characterize how rotational or irrotational the normal modes are,

$$\mathbb{P}(S) \equiv \frac{|\mathcal{P}^{\mathrm{irr}}S|^2}{|S|^2} \tag{D.40}$$

By definition, $\mathbb{P}(S) = 1$ when the configuration is fully irrotational, and is zero for rotational configurations (with zero spin-divergence everywhere). In general it may be somewhere between, with coexistence of rotational and irrotational components. The harmonic component is subextensive, only contributing to the $\boldsymbol{q} = \boldsymbol{0}$ Fourier component, which by definition has zero rotational and irrotational content.

We are interested in describing the HHD of the dynamical normal modes,

$$S^{\perp}_{\boldsymbol{q},m}(t) \equiv \psi_{\boldsymbol{q},m}\, e^{+i\omega_n(\boldsymbol{q})t}, \tag{D.41}$$

which are periodically oscillating plane waves whose intra-unit-cell behavior is described by ψ. To do so, we slightly generalize the definition of $\mathcal{P}$ for these modes as the average irrotational projection squared over a full period $T = 2\pi/\omega_n(\boldsymbol{q})$,

$$\mathcal{P}_{\boldsymbol{q},m} \equiv \frac{\int_0^T \mathrm{d}t\, |\mathcal{P}^{\mathrm{irr}}S^{\perp}_{\boldsymbol{q},m}(t)|^2}{\int_0^T \mathrm{d}t\, |S^{\perp}_{\boldsymbol{q},m}(t)|^2} = \frac{|\mathcal{P}^{\mathrm{irr}}\psi_{\boldsymbol{q},m}|^2}{|\psi_{\boldsymbol{q},m}|^2}. \tag{D.42}$$

To demonstrate the second equality, we first consider the denominator

$$\frac{1}{T}\int_0^T \mathrm{d}t|S_{\boldsymbol{q},m}(t)|^2, \tag{D.43}$$

$$= \frac{1}{T}\int_0^T \mathrm{d}t|\psi_{\boldsymbol{q},m}(t)e^{i\omega_n(\boldsymbol{q})t}|^2, \tag{D.44}$$

$$= \frac{1}{T}\int_0^T \mathrm{d}t[\mathrm{Re}(\psi e^{i\omega t})^2 + \mathrm{Im}(\psi e^{i\omega t})^2], \tag{D.45}$$

$$= \frac{1}{T}\int_0^T \mathrm{d}t[(\mathrm{Re}(\psi)\cos(\omega t) - \mathrm{Im}(\psi)\sin(\omega t))^2 + (\mathrm{Re}(\psi)\sin(\omega t) + \mathrm{Im}(\psi)\cos(\omega t))^2] \tag{D.46}$$

Using the fact that $\int_0^T \cos(\omega t)\sin(\omega t) = 0$, and the identity $\cos^2 + \sin^2 = 1$, this reduces to

$$\frac{1}{T}\int_0^T \mathrm{d}t|S^{\perp}_{\boldsymbol{q},m}(t)|^2 = \mathrm{Re}(\psi_{\boldsymbol{q},m})^2 + \mathrm{Im}(\psi_{\boldsymbol{q},m})^2 = |\psi_{\boldsymbol{q},m}|^2.$$

The above argument works exactly the same replacing everywhere ψ by $\mathcal{P}^{\mathrm{irr}}\psi$, since $\mathcal{P}^{\mathrm{irr}}$ has real matrix elements, thus demonstrating the second equality in Eq. (D.42).