# Aspects of Frustrated Magnetism


Peter Hugh Conlon

University College

University of Oxford

A thesis submitted for the degree of

*Doctor of Philosophy*

Michaelmas 2010

To my parents

# Acknowledgements

I am very grateful to my supervisor, Professor John Chalker, for his help and guidance during my time in Oxford. I have learnt a lot from our discussions.

I have enjoyed the company of my office mates Adam and Tom over the past three years and I'm grateful to them for making Room 2.7 a great place to work.

Above all, I am indebted to my family and boyfriend Tom for the love and support which keeps me going.

# Abstract

Models of magnetism show complex collective behaviour which arises from simple interactions among microscopic degrees of freedom. Upon cooling from high temperatures conventional magnets typically undergo a phase transition to a magnetically ordered phase due to microscopic interactions which favour an ordered state. In frustrated magnets however, competing microscopic interactions place non-trivial constraints on the allowed configurations at low temperature, without selecting a unique ordered state.

In this thesis, we investigate the collective behaviour of a paradigmatic frustrated magnet, the classical Heisenberg model on the pyrochlore lattice with antiferromagnetic nearest neighbour interactions. Within a self consistent Gaussian approximation, we derive analytic expressions for correlation functions which match Monte Carlo simulations extremely well at all temperatures. We study the precessional dynamics of the model and provide a comprehensive description of the dynamics by constructing an analytically tractable stochastic model by extending the self-consistent Gaussian approximation to include dynamics. We relate these results to other highly constrained models.

Real experimental systems often have features that go beyond the phenomenology afforded by the simplest models; we investigate the effects of further neighbour interactions on paramagnetic spin correlations, and propose further neighbour interactions as the mechanism underlying experimentally observed patterns of scattering in frustrated spinel compounds.

In the dynamics linearized around a ground-state, the macroscopic degeneracy of the classical model leads to modes with zero frequency. Small perturbations stabilize ordered states and lift such zero modes to finite frequency. The ordered state has two widely separated energy scales both of which affect the dynamics, the leading scale of nearest neighbour exchange coupling, and the much smaller scale of the perturbation which relieves the frustration. We investigate the interplay between these widely different energy scales in setting the mode frequencies in states ordered by weak interactions.

# Contents

# Chapter 1

# The study of magnetism in condensed matter

The physical phenomena familiar from every day life occur on a scale much larger than the scale of the constituents from which they are composed. The swell of an ocean wave may involve over $10^{32}$ water molecules but no amount of observing ocean waves would lead us to conclude that water is composed of molecules, each of which has two hydrogen atoms and one oxygen atom. Only by careful experiments probing smaller and smaller scales have we come to know the nature of the constituents of the world around us, whose collective behaviour is so easy to see. The diversity of chemical compounds are understood as arising from different combinations of atoms, the diversity of chemical elements themselves are understood in terms of electrons and the nucleus and these subatomic particles are understood within the standard model of particle physics, humankind's best understanding of the most elementary constituents of the universe we inhabit. With an understanding of how the microscopic constituents interact with each other, one might hope to deduce from elementary laws the diversity of phenomena which we experience, yet this turns out to be extremely difficult.

While understanding the microscopic constituents of matter has been a great triumph, access to a microscopic model from which to deduce the behaviour of things is a luxury which has not always been available but this does not prevent understanding. The equations of fluid mechanics which describe ocean waves were written down without any reference to water molecules; half a century before anyone knew about electrons, Rudolf Clausius formulated a version of the second law of thermodynamics. Both these examples have a universal applicability that remains even given a much greater knowledge of the microscopic constituents of matter. That fluid mechanics and thermodynamics remain useful is because they address questions whose answers do not depend on microscopic details, but rather are consequences of symmetries

and conservation laws common to a variety of possible microscopic interactions. The exact relationships between experimentally measurable quantities which come from thermodynamic reasoning are not deduced from microscopic details but transcend them.

In understanding how the macroscopic emerges from the microscopic, one must recognise that observed phenomena involve statistical averages of many microscopic constituents. Understanding how this occurs is the subject of statistical mechanics. The development of statistical mechanics established the relationship between thermodynamic functions and the microscopic descriptions of matter. In the simplest cases, the macroscopic behaviour can be understood from statistically independent behaviour of the constituents, such as in the theory of ideal gases. More generally, the constituents will interact strongly such that the behaviour of the constituents is correlated and the macroscopic phenomena encountered is much richer: it is naive to hope that the superconductivity of mercury can be understood by perturbing an ideal gas of mercury atoms.

Condensed matter physics addresses precisely these problems, to understand the phenomena which arise from a large number of constituents which interact strongly. In condensed matter systems, the properties of the constituents are often known extremely well: properties of a single electron are some of the most accurately measured quantities in experimental physics, and even the properties of single atoms are well characterised yet despite this, understanding what happens when many such constituents come together is not straightforward - crucially, this is not because of imprecisely specified models but rather deduction directly from the governing equations to the collective behaviour can be impossible. Seen in this context, the importance of simplified models of condensed matter is clear. By studying models which can be understood analytically or through simulation, one gains insight into the nature of the transformation from microscopic to macroscopic which applies to a broader class of situations than simply the specific models one studies.

Magnetism has fascinated for thousands of years and magnetic phenomena continued to be studied both for their own sake, and the insight into collective behaviour which it provides. Models of magnetism strike a balance of being simple enough to imagine and sometimes to solve, while also exhibiting the behaviour characteristic of the strongly interacting systems which are the focus of condensed matter physics. Many models have a well defined lattice structure with short range interactions which makes the system amenable to efficient computer simulations.

Magnetism is also an experimentally rich subject since there are many compounds with different crystal structures and interactions which exhibit magnetic behaviour and the experimental probes of magnetism give detailed information on magnetic behaviour. The diversity of systems gives reason to search for exotic phases of matter both in magnetic materials, and in the theoretical models used to describe them.

## 1.1 Models of magnetism

In light of the dual aims, both to understand real magnetic materials, but also to understand more general principles of strongly interacting systems, many theoretical models of magnetic phenomena exist - some are simple, others more realistic. In most materials, magnetic behaviour is due to the properties of electrons. Electrons have charge, and their orbital motion gives rise to an orbital magnetic moment. In addition, electrons have an intrinsic angular momentum, spin, which has an associated intrinsic magnetic moment. Atomic ions have many electrons which can have net angular momentum and a net magnetic moment. When in a lattice, the magnetism of each ion can interact with that of the others nearby and lead to interesting collective behaviour. The theoretical models of magnetism which we study seek to reproduce this situation: degrees of freedom, often called spin, are considered localised on lattice sites. The physics is governed by the interactions of each site with those close by.

The simplest such model is known as the *Ising model*, which was invented by Wilhelm Lenz. The degrees of freedom are 'Ising spins', variables which can take one of two values such that on a lattice site $i$, the spin $\sigma_i$ is either $+1$ or $-1$. According to the principles of statistical mechanics, configurations enter the partition sum weighted by the Boltzmann factor, $e^{-\beta E_n}$ where $E_n$ is the energy of the configuration labeled by $n$. Configurations of the Ising model have energy

$$E = \frac{1}{2} \sum_{ij} J_{ij} \sigma_i \sigma_j$$

where $J_{ij}$ is a symmetric matrix of couplings between lattice sites. Often, $J_{ij}$ is a constant for nearest neighbour sites, and zero otherwise. The Ising model is appealing in its simplicity, and because thermodynamic functions can be obtained exactly in a number of situations. Lars Onsager famously solved the square lattice Ising model allowing insight into the properties at the critical point which separates a disordered phase at high temperature from an ordered phase at low temperatures. Exact solutions are not available in higher dimensions, but the simplicity of the model makes it amenable to large scale simulations.

## 1.2 Heisenberg model

The Ising model is not a very realistic model of magnetic materials. We discuss below the *Heisenberg model*, which is a quantum mechanical model describing localised magnetic moments. As discussed in textbooks on quantum mechanics, the quantum mechanical operators for spin obey the angular momentum commutation relations:

$$[\hat{S}^\alpha, \hat{S}^\beta] = i\hbar\epsilon^{\alpha\beta\gamma}\hat{S}^\gamma\,, \tag{1.1}$$

where the Greek letters are labels for spin components. A spin degree of freedom is characterised by the spin quantum number, $S$. While a single electron always has $S = 1/2$, many electrons combined can give rise to $S = 1/2, 1, 3/2, \ldots$, and so the Heisenberg model is considered for general $S$. It is frequent also to refer to the spin length $S$, which follows from the operator equality $(\hat{S}^x)^2+(\hat{S}^y)^2+(\hat{S}^z)^2 = S(S+1)\hbar^2\mathbf{1}$. Thus $\sqrt{S(S+1)}\hbar$ is like the conserved length of the spin angular momentum vector.

The isotropic Heisenberg model is given by a quantum mechanical Hamiltonian operator

$$H = \frac{1}{2}\sum_{ij} J_{ij}\hat{\mathbf{S}}_i \cdot \hat{\mathbf{S}}_j \tag{1.2}$$

where $J_{ij}$ is a symmetric matrix of couplings between the spins at the different sites of the lattice. The couplings $J_{ij}$ are often called exchange constants because such couplings arise as the exchange integral when taking into account Coulomb repulsion between electrons in perturbation theory. One can deform the Heisenberg model by allowing the coupling strength to depend on the spin components - in this case it is anisotropic and in the limit where there are only couplings for one spin component of an $S = 1/2$ degree of freedom, one recovers the Ising model. There is however an important difference between the Ising and Heisenberg models - the Heisenberg model is specified by a Hamiltonian operator and has an intrinsic quantum dynamics. In contrast, the Ising model simply assigns energy to configurations, without generating any equations of motion. The dynamics of Ising models is imposed rather than intrinsic.

### 1.2.1 Susceptibility and mean field theory

This thesis is concerned with the Heisenberg model so we consider here some thermodynamic properties of the model, Eq. (1.2), with $N$ sites. The uniform magnetic susceptibility per spin is the coefficient of the linear part of the change in total magnetisation $\sum_i S_i^z$ in an applied external field, i.e $\chi_z = \partial m/\partial h$. In the high temperature

phase, one can show from the partition function that the susceptibility is related to the correlation function of the magnetisation.

$$\chi(T) = \frac{1}{NT}\langle(\sum_i S_i^z)^2\rangle = \frac{1}{NT}\sum_{ij}\langle S_i^z S_j^z\rangle\,. \tag{1.3}$$

Here, $\langle\cdots\rangle$ is a thermal average. The Curie-Weiss law which we now discuss relates to the high temperature behaviour of the susceptibility, which we now derive. At infinite temperature, spins are statistically independent and $\langle S_i^z S_j^z\rangle_0 = \frac{1}{3}S(S+1)\delta_{ij}$ where the $\langle\cdots\rangle_0$ indicates an average on the infinite temperature distribution. To derive the leading order corrections we expand the thermal average at high temperatures:

$$\langle S_i^z S_j^z\rangle = \frac{\frac{1}{3}S(S+1) - \beta\langle S_i^z S_j^z H\rangle_0}{1-\beta\langle H\rangle_0}$$

Since all sites are independent under the infinite temperature distribution, while the Hamiltonian only has terms which couple different sites, the only contribution at first order in $\beta$ comes from the terms in the numerator:

$$\langle S_i^z S_j^z\rangle \simeq \frac{1}{3}S(S+1) - \beta\left[\frac{1}{2}J_{ij}\langle S_i^z S_j^z \mathbf{S}_i\cdot\mathbf{S}_j\rangle + \frac{1}{2}J_{ji}\langle S_i^z S_j^z \mathbf{S}_j\cdot\mathbf{S}_i\rangle\right]\,.$$

Furthermore, separate components of spin are uncorrelated and one obtains the first order expansion of the correlation function

$$\frac{S(S+1)}{3}\left[\delta_{ij} + \frac{S(S+1)}{3}\beta J_{ij} + \mathcal{O}((\beta J)^2)\right]\,.$$

Substituting this in Eq. (1.3) leads to the Curie-Weiss law

$$\chi(T) = \frac{1}{T}\frac{S(S+1)}{3}\left[1 + \frac{\theta}{T} + \mathcal{O}\left(\frac{\theta}{T}\right)^2\right] \quad \text{where } \theta = \frac{S(S+1)}{3}J_0\,, \quad J_0 = \frac{1}{N}\sum_{ij}J_{ij}\,.$$

For nearest neighbour lattices with coordination number $z$, and interaction strength $J$ on nearest neighbour bonds, we have $J_0 = zJ$. We see that the high temperature expansion provides a direct probe of the strength of the interactions. This is made more transparent by examining the inverse susceptibility

$$\chi^{-1}(T) = \frac{3}{S(S+1)}\left[T - \theta + \cdots\right]\,. \tag{1.4}$$

By a linear extrapolation of a high temperature susceptibility obtained from experimental measurements, the intercept with the $T$-axis defines a temperature[1] $\Theta_{CW}$ from

[1] we use $\Theta_{CW}$ for the experimentally measured quantity, and $\theta$ for the theoretical one

which it is possible to infer the strength and sign of the spin interactions, assuming some knowledge of the model Hamiltonian. In general, a negative $\Theta_{CW}$ indicates antiferromagnetic interactions and positive, ferromagnetic. The presence of further neighbour exchange which may be of a different sign can complicate the inference however.

The Curie-Weiss temperature relates a measurable property of a real system to model parameters in a Hamiltonian and thus is a useful way of finding models to describe systems. It has further significance, as we now show, in that the temperature of the magnetic phase transition, $T_c$, in the mean field approximation is equal to $\theta$ as defined above.

In generalising Eq. (1.2) to include an external field $\mathbf{h}$, the only terms containing site $i$ in the Hamiltonian are

$$-\mathbf{S}_i \cdot \mathbf{h}_{\text{eff}} \quad \text{where} \quad \mathbf{h}_{\text{eff}} = \sum_{j \neq i} J_{ij} \mathbf{S}_j + \mathbf{h}$$

The field $\mathbf{h}_{\text{eff}}$ describes the interactions both with the external field and with the other spins. In the full problem, the $\mathbf{S}_j$ are of course fluctuating, but the mean-field approximation makes progress by replacing in $\mathbf{h}_{\text{eff}}$ each spin with its thermal average, which must be the total magnetisation per spin, $\mathbf{m}$. This leads to an effective field acting on each spin $J_0\mathbf{m} + \mathbf{h}$ and the interacting problem has been reduced to a self-consistent one-site problem. The self-consistent condition is that $m = m_0(\beta(J_0 m + h))$ where $m_0(\beta h)$ is the thermodynamic magnetisation obtained from solving the problem of a single spin in a field of strength $h$ - such a solution is in terms of Brillouin functions for quantum spins, [see 1, Chapter 31]. It can be shown that non-zero solutions for $m$ exist only for $T < S(S+1)J_0/3$, ie. mean field theory predicts a phase transition at

$$T_c = \frac{S(S+1)}{3} J_0 = \theta \,.$$

Mean field theory is crude, but nevertheless provides a useful starting point to understand the behaviour of interacting systems.

## 1.3 Dynamics of magnets and the classical limit

Much of the above applies equally to the Ising model which also has a thermodynamic susceptibility and undergoes a phase transition but we have already mentioned that an important difference is the presence of an intrinsic dynamics in the Heisenberg

model. In the Heisenberg picture of standard quantum mechanics, operators have a time evolution governed by the Hamiltonian, where an operator evolves as

$$S_i^\alpha(t) = e^{iHt/\hbar} S_i^\alpha(0) e^{-iHt/\hbar}\,.$$

The differential form is the Heisenberg equation of motion. For the Heisenberg model we study, the equations of motion for the operators are

$$\frac{d\hat{\mathbf{S}}_i}{dt} = \hat{\mathbf{S}}_i \times \sum_j J_{ij} \hat{\mathbf{S}}_j \tag{1.5}$$

where we have suppressed the $\hbar$. There is a natural classical limit for the Heisenberg model, obtained by replacing commutators, Eq. (1.1), by Poisson brackets. The spin degrees of freedom in the classical Heisenberg model have a phase space which is a sphere and the dynamics is governed by Poisson-bracket relations among the dynamical variables which mimics the commutation relations:

$$\left\{S^\alpha, S^\beta\right\} = \epsilon^{\alpha\beta\gamma} S^\gamma \tag{1.6}$$

In this way, the classical Hamiltonian takes the same form, Eq. (1.2), and the equations of motion are Eq. (1.5) but interpreted for classical dynamical variables rather than operators. This thesis is mostly concerned with the classical Heisenberg model with the dynamics of Eq. (1.5).

It is worth making the distinction clear between classical spins, and simply degrees of freedom which live on a sphere. When kinetic energy is quadratic in conjugate momenta, its contribution to the thermodynamics is trivial since the momenta can be integrated out of the classical partition sum. Thus, Eq. (1.2) could be considered simply as a potential energy term, which assigns energy to configuration space variables on the sphere. In fact, Eq. (1.2) is a Hamiltonian, but one which does not split into kinetic and potential terms. The dynamics generated by Eq. (1.2) are not the same as those in a rotor model for example, for which there are separate momentum variables.

The classical limit of the quantum problem is obtained for large values of $S$, the spin quantum number. Corrections to the classical limit can be taken into account in a systematic expansion in terms of the small parameter $1/S$. One way to effect this is the Holstein-Primakoff transformation, a representation of the spin operators satisfying Eq. (1.1), in terms of bosonic creation and annihilation operators $a$ which

satisfy canonical commutation relations $[a, a^\dagger] = 1$. The representation facilitates the expansion in $1/S$. Relative to a quantization axis $z$, the spin operators are :

$$\hat{S}^z = S - a^\dagger a \qquad \hat{S}^+ = (2S - a^\dagger a)^{1/2} a \qquad \hat{S}^- = a^\dagger (2S - a^\dagger a)^{1/2} \tag{1.7}$$

where $S^\pm = S^x \pm i S^y$. For large $S$, one obtains

$$S^z = S - a^\dagger a \tag{1.8}$$

$$S^+ \simeq \sqrt{2S} a \tag{1.9}$$

$$S^- \simeq \sqrt{2S} a^\dagger \tag{1.10}$$

Taking a product of two spins, the classical energy contributes to the Hamiltonian at $JS^2$ whereas quantum behaviour appears at $JS$. Correspondingly, for large $S$, there is a temperature regime for which quantum corrections may be ignored ($T \gg JS$) but is nevertheless strongly interacting ($T \ll JS^2$).

## 1.4 Role of correlation functions

In the calculation of the high-temperature susceptibility above, we used that the susceptibility is proportional to the variance of the total magnetisation - the fluctuations in equilibrium of magnetisation determine the response in the presence of an externally applied conjugate field. The relationship between equilibrium correlations and response to externally applied fields is in fact more general and will be briefly discussed below. We first introduce time-dependent correlation functions which will characterise the dynamic behaviour.

In both the classical and quantum case, one can define a time-dependent correlation function such as

$$\langle S_i^\alpha(t) S_j^\beta(t') \rangle$$

where $\langle \cdots \rangle$ denotes the thermal average. In the quantum case, this is a trace over quantum states with a thermal density matrix, while in the classical case, it means a Boltzmann-weighted average over starting points for time evolution in classical phase space.

### 1.4.1 Linear response

To measure a system, one must couple it to an external probe and in doing so, the system Hamiltonian is altered. Typically, a perturbation at position $\mathbf{x}'$ and time $t'$ will cause a response of the system at other positions $\mathbf{x}$ and later times $t > t'$.

When the perturbation is small, one expects the response to the perturbation to be linear and the relationship between the response of the system and the perturbation is given by a response function, a linear operator. The remarkable fact alluded to above is that within linear response, the response function is determined by equilibrium fluctuations. The physical intuition behind this is easy to understand: when a perturbation is small, the system does not know whether it arose from a spontaneous fluctuation in equilibrium or from an externally applied probe; the response of the system must therefore be the same in both cases. This observation lies at the heart of the fluctuation-dissipation theorem, which relates equilibrium correlation functions to response functions, such as for susceptibility. For dynamic quantities, Kubo formulae relate equilibrium current correlations to transport coefficients.

For real materials, the linear response regime can often be large, and justifies the importance of correlation functions to characterise both equilibrium behaviour and the response to perturbations.

### 1.4.2 Measuring with neutrons

Neutrons have a magnetic moment but no charge. Their electrical neutrality allows them to penetrate into the bulk of a material while the magnetic moment causes interactions with magnetic materials. The scattering of neutrons from materials is one of the most powerful experimental techniques to probe strongly correlated magnetic behaviour. The interaction of neutrons with matter is weak, and probes the linear response regime. The neutron scattering cross-section is directly proportional to the Fourier transform of the correlation function, called the dynamic structure factor.

To express the dynamic structure factor, we define the magnetisation density $M^\alpha(\mathbf{r},t) = \sum_i S_i^\alpha(t)\delta(\mathbf{r}-\mathbf{r}_i)$. The Fourier transform of $M^\alpha(\mathbf{r},t)$ is

$$M^\alpha(\mathbf{q},t) = \int d\mathbf{r} e^{i\mathbf{q}\cdot\mathbf{r}} M^\alpha(\mathbf{r},t)\,.$$

The dynamic structure factor is

$$S^{\alpha\beta}(\mathbf{q},\omega) = \int dt e^{-i\omega t}\langle M^\alpha(-\mathbf{q},0)M^\beta(\mathbf{q},t)\rangle$$

which in general carries spin indices - when there is full spin rotation symmetry, $S^{\alpha\beta}(\mathbf{q},\omega)$ is proportional to $\delta_{\alpha\beta}$. Many of the results of this thesis are in terms of predictions for the dynamic structure factor, since this has direct experimental relevance.

## 1.5 Universality and hydrodynamics in magnets

The ingredients of a thermodynamic description of matter are symmetries, since these give rise to conservation laws in dynamics. The conserved quantities specify the macrostates of thermodynamics. Callen [2] has said 'thermodynamics is the study of the restrictions on the possible properties of matter that follow from the symmetry properties of the fundamental laws of physics'. The symmetries of fundamental laws necessarily appear in microscopic models of matter and different phases of matter are characterised by the symmetries of their microscopic models which they respect. A phase transition occurs as a function of a microscopic parameter or temperature at the boundary between two phases of matter. When a phase has a lower symmetry than the microscopic model, one talks about a spontaneously broken symmetry. When a phase has a spontaneously broken symmetry, a new thermodynamic variable is required to describe thermodynamic states in addition to the conserved quantities which follow from the symmetries of the microscopic laws. In characterising phases of matter by the symmetries of the model, it is clear that this characterisation is in some ways insensitive to the microscopic details in as far as they are not determined by symmetry considerations.

While thermodynamics is a theory of time-independent phenomena, the leading correction is hydrodynamics, the theory of the slow, long-wavelength behaviour. The variables which enter a hydrodynamic description are the thermodynamic variables, but allowing for variation in time and space. Since many experimental probes are coarse in time and space, they necessarily measure slow, long-wavelength behaviour, which is the hydrodynamic regime whose laws are governed by symmetry considerations. In studying unrealistic microscopic models, we recognise that the hydrodynamic laws will be shared by other more complicated models which respect the same symmetries. The parameters of the laws will indeed change with the microscopic model, and may be important in deciding whether a material is suitable or not for use as a heat-shield, or for making transmission wires, but the essence of particular behaviours is given by symmetries.

The original example of hydrodynamics is fluid mechanics, universal laws for describing fluid behaviour, but the same ideas apply to spin systems [3]. In Heisenberg magnets, there is a global spin rotation symmetry which guarantees the conservation of magnetisation. In ordered states, spin rotation symmetry is spontaneously broken. The ordered state has hydrodynamic modes with linear dispersion $\omega \propto k$

in antiferromagnets, while in ferromagnets there are modes with quadratic dispersion $\omega \propto k^2$. The latter case arises because the order parameter in a ferromagnet is a microscopically conserved quantity, while in an antiferromagnet, the staggered magnetisation order parameter is not microscopically conserved, and enters as a separate thermodynamic variable. The coefficients in these dispersion relations depend on microscopic details while the relations themselves are universal, independent of microscopic details.

A description of universal phenomena cannot depend on the microscopic details and so such a description cannot be formulated on any particular lattice but must be in the continuum, characterised only by a spatial dimensionality. The study of variables carrying a continuous label is the topic of field theory. In this thesis as we investigate particular models, we include field theory descriptions to include related models which share the same universal features.

# Chapter 2

# Geometrical frustration and the Coulomb-phase description

In the previous chapter, we outlined the origin of magnetism in condensed matter and the motivation for its study. We also explained how much of the collective behaviour is encoded in correlation functions, which often are probed directly in experiments. The most striking feature of conventional magnets is that they undergo a phase transition to long range order, and this occurs at a critical temperature set by the local interaction energy. Near this temperature, which in mean field theory is $T_c = |\Theta_{CW}|$, the appropriate description is in terms of a field theory for coarse grained microscopic variables. The critical field theories are universal because many different microscopic models can have the same description near their critical points.

In this chapter we introduce the notion of frustrated magnets in general and strongly geometrically frustrated magnets in particular. In contrast to conventional magnets, their most striking feature is the absence of any phase transition at the scale of the strongest interactions. In its place, frustrated magnets undergo a crossover from a weakly interacting paramagnetic regime at high temperatures to a strongly interacting regime at low temperatures, which nevertheless remains paramagnetic. The strong interactions put strong constraints on allowed low energy configurations, but in doing so can leave a macroscopic number of collective degrees of freedom approximately unconstrained. Villain [4] has called this low temperature regime *cooperative paramagnetism*. Generically, the approximately unconstrained degrees of freedom will have a energy width $\Delta \ll J$, which leads at sufficiently low temperatures, $T \sim \Delta$, to selection of a unique ground state. Such selection isn't guaranteed however, and simple classical models can have macroscopic ground state degeneracy. Research in this area is directed both at what exotic ground states are selected, and the nature

of the correlations in the cooperative paramagnet. This thesis is mainly concerned with the latter.

A universal description of frustrated systems might seem to be lacking without an order parameter whose local average can be promoted to a field. Nevertheless, we will explain how at low temperatures, certain types of frustrated magnet can be understood to be in a 'Coulomb phase' described by an associated continuum field theory. This continuum description unifies geometrically frustrated magnets with other highly constrained models such as dimer coverings.

## 2.1 Geometrically frustration in two and three dimensions

For the sake of having a definition, *frustration* occurs when a system has competing interactions whose energy cannot be simultaneously minimised. In magnets, there are a number of ways to achieve this situation. One route is by the imposition of randomness on the strength and sign of the spin interactions. This results in a free energy landscape with a large number of local minima, with long timescales for moves between minima. At low temperatures, the spins freeze into a local minimum of the energy, even though it may not be the global minimum. The random interactions mean the frozen state is unlike the uniform or periodic patterns into which conventional magnets order. Such systems are known as spin glasses and there is an extensive literature on them [5].

In contrast to frustration due to randomness, this thesis is concerned with situations where the frustration arises directly from the lattice geometry. To distinguish from the former case, Ramirez [6] coined the term *geometrically frustrated* to describe such systems. What sort of models on what sort of lattices have geometrical frustration? Clearly, purely ferromagnetic models cannot be frustrated[1] since the ground-state with all spins aligned is unique up to global rotations - there are no competing interactions in simple ferromagnets and we are led to consider antiferromagnetic interactions. Antiferromagnetic models on a bipartite lattice[2] are really ferromagnetic in disguise, at least in classical models. This is because the classical Hamiltonian has an exact symmetry associated with inverting all spins on a particular sublattice while simultaneously changing the sign of the interaction. This correspondence does not

[1] although sometimes systems with ferromagnetic interactions and strong crystal fields map onto antiferromagnetic models, as occurs in spin ice[7].

[2] A lattice is bipartite if its sites can be partitioned into two sets, $A$ and $B$ say, such that $A$ sites are only connected to $B$ sites and *vice versa*.

carry over to the quantum case where the operator nature of the spins invalidates the argument given and nor is it true in the classical dynamics which is derived from the spin commutation relations. As we mentioned in Chapter 1, hydrodynamic spin wave modes have linear dispersion in antiferromagnets, but quadratic dispersion in ferromagnets.

Lattices which are not bipartite are guaranteed to have loops with an odd number of sites, which generically leads to frustrated antiferromagnetic interactions. We demonstrate this point with triangles, the shortest such frustrated loop. The spin Hamiltonian for three mutually interacting spins is

$$H = J(\mathbf{S}_1 \cdot \mathbf{S}_2 + \mathbf{S}_2 \cdot \mathbf{S}_3 + \mathbf{S}_3 \cdot \mathbf{S}_1) = \frac{J}{2}(\mathbf{S}_1 + \mathbf{S}_2 + \mathbf{S}_3)^2 - \frac{3}{2}JS^2 \,.$$

Clearly the minimum energy is $-3JS^2/2$, which is only half the $-3JS^2$ which would have been achieved if all the bonds were satisfied. The minimum of energy is achieved by a coplanar spin state, where the spins are mutually at 120°. For 2-component spins, there are two topologically distinct ways of doing this, but for 3 and higher, all ground states are connected by global spin rotations. For Ising spins, the 120° state is impossible, and the minimum energy is $-JS^2$. There is a 3-fold degeneracy in choosing the unsatisfied bond, which is in addition to the degeneracy due to a global spin flip.

Applying these ideas on a bigger scale, frustrated lattices will be those with triangular motifs. In Fig. 2.1(Left) we show the two dimensional triangular lattice. Wannier [8] obtained the thermodynamic functions for the antiferromagnetic Ising model (AFMIM) on the triangular lattice exactly. The system remains paramagnetic for all $T > 0$, and there is a finite entropy density at $T = 0$, consistent with a macroscopic ground state degeneracy associated with the freedom to choose a single unsatisfied bond on every triangle. The degeneracy is not a trivial one, as one would find in a model with independent spins; Stephenson [9] calculated the correlation functions for the AFMIM, and demonstrated algebraic correlations at $T = 0$, associated with a zero temperature critical point. Behaviour such as this is one reason for the interest frustrated magnetism has generated. For two component spins and higher on the triangular lattice, there is long range order at zero temperature in the 120° state discussed above, extended across the lattice. In this case, the only degeneracies are global, so although there is frustration, there is not a macroscopic degeneracy of ground states.

On the right in Fig. 2.1 is the kagome lattice. Whereas the triangular lattice is composed of edge-sharing triangles, the kagome lattice is composed of corner-sharing

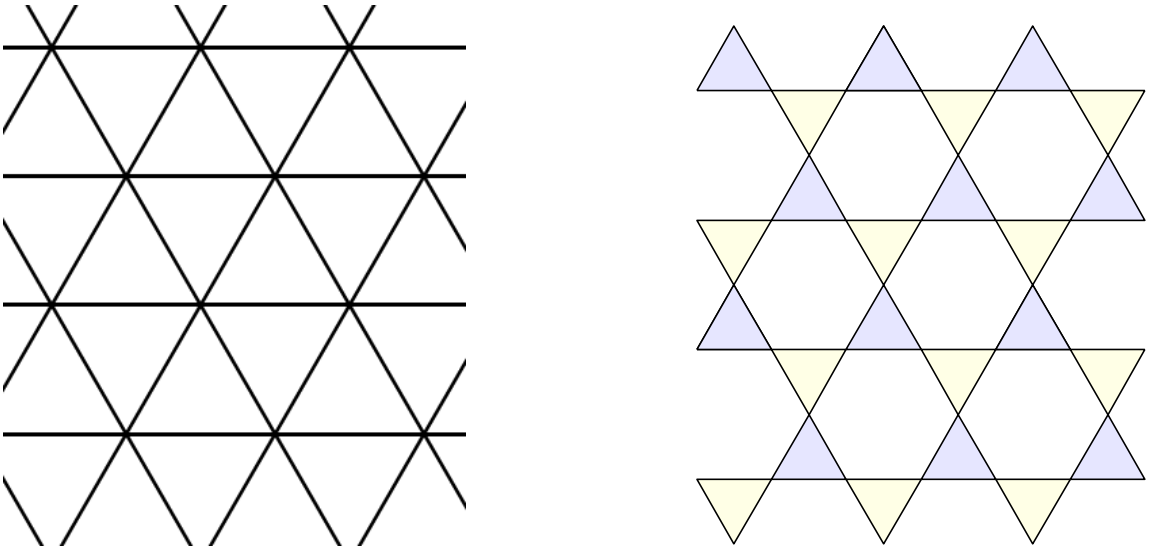

Figure 2.1: Left: the triangular lattice, of edge-sharing triangles. Right: the kagome lattice of corner-sharing triangles. Kagome sites are at the midpoints of the bonds of the honeycomb lattice.

triangles. There is a much larger potential degeneracy, since information propagates from triangle to triangle via only one spin. For two component (XY) spins for example, ground states map onto those of a 3 state Potts model [10], and so a macroscopic, although discrete, degeneracy occurs. For three component spins, the XY state is also a ground state but starting from it a non-trivial continuous degeneracy of ground states then appears, since in any region bounded entirely by collinear spins, the interior spins can undergo a rotation about the axis defined by the boundary spins at no energy cost. The rearrangment is entirely local, indicating a macroscopic ground state degeneracy. This is our first example of a system where ground states have a continuous degeneracy, beyond trivial global rotations. For continuous classical degeneracy like this, we take 'macroscopic degeneracy of ground states' to mean the dimension of the manifold of groundstates is extensive, whereas a discrete degeneracy is macroscopic if the number of degenerate states is extensive.

The equilibrium behaviour of the three-component (Heisenberg) kagome model is more subtle than the preceding analysis would suggest. For continuous degrees of freedom in general, zero energy states are weighted by a measure associated with the phase space volume (or entropy) of low energy fluctuations around the states. The effect of the entropic weighting can lead to selection of ordered states at sufficiently low temperature, known as order-by-disorder. Chalker *et al.* [11] showed that Heisenberg spins on the kagome lattice order in a coplanar state by the order-by-disorder

mechanism while Moessner and Chalker [12] established that order-by-disorder is less important as the number of spins mutually interacting in a frustrated unit increases.

In three dimensions, the equivalent of the triangle is the tetrahedron, providing two 3-dimensional lattices analogous to the 2-dimensional ones already discussed: the face-centred cubic (fcc) lattice can be thought of as edge-sharing tetrahedra, while the lattice of corner-sharing tetrahedra in Fig. 2.2 is known as the pyrochlore lattice. The figure also shows the checkerboard lattice, a two dimensional analogue obtained by imposition of periodic boundary conditions in the $z$-direction. The checkerboard is also known as planar pyrochlore, or the square lattice with crossings. The Heisenberg model on the pyrochlore lattice is the main focus of this thesis. A mode counting argument establishes that it has a macroscopic degeneracy [13]: for a lattice of corner sharing frustrated units, assume there are $2N$ frustrated units, each with $q$ spins, each with $n$ components. Due to spin length constraints and the fact that each spin is shared with two units, the total number of degrees of freedom is $D = Nq(n-1)$. By the same argument we used for the single triangle, groundstates will have the sum of the spins vanish on each frustrated unit. This introduces $F = 2Nn$ constraints. Naively, the groundstate has $D - F = N(q(n-1) - 2n)$ degrees of freedom. For Heisenberg spins ($n = 3$) on the pyrochlore lattice ($q = 4$), one expects a macroscopic groundstate degeneracy. We will say more about the pyrochlore Heisenberg model in Chapter 3. For other models, Anderson [14] appreciated that Ising spins on the pyrochlore lattice would have a large degeneracy in the ground state and finite ground state entropy. Spin ices are a class of materials which map onto the pyrochlore Ising antiferromagnet; their properties are described in Ref. 7.

In summary, frustrated systems can have a classical groundstate degeneracy which is much larger than that required by symmetry alone. This extra degeneracy, which is sometimes called *accidental*, means that classical frustrated systems have significant entropy close to zero temperature.

## 2.2 Experimental signatures of frustration

As we have intimated, one of the properties of frustrated systems is to avoid phase transitions at the temperature set by the scale of the microscopic interactions. Macroscopic probes such as susceptibility are often used to identify phase transitions, since thermodynamic functions show non-analytic behaviour at a transition. As explained by Ramirez [6], one macroscopic signature of frustration is the behaviour of the inverse susceptibility. The left panel of Fig. 2.3 shows the typical situation in a frustrated

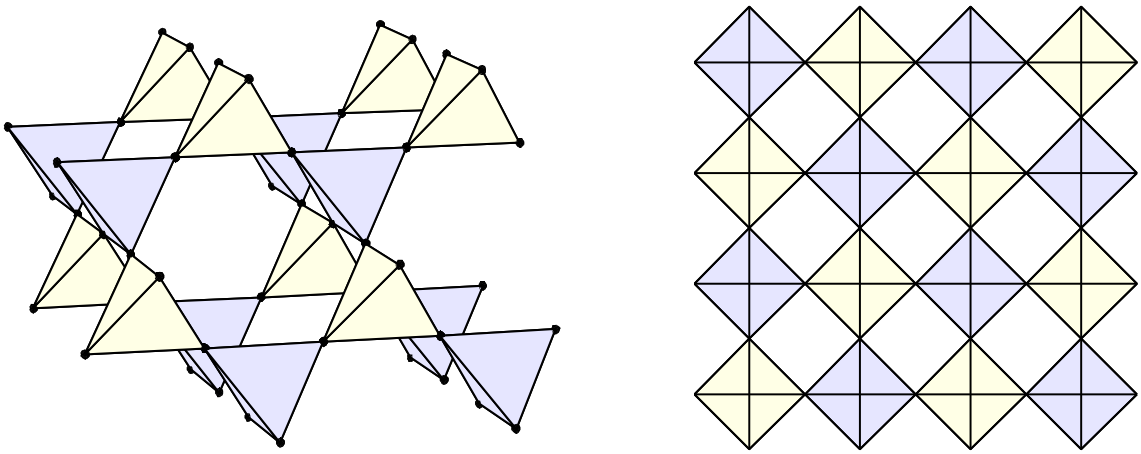

Figure 2.2: Left: section of pyrochlore lattice. The section repeated infinitely generates the pyrochlore lattice. (111) planes in the pyrochlore lattice form kagome lattices. Right: the same image as on the left viewed from above, also the checkerboard lattice. The sites are at the corners of the shaded squares.

magnet. Ordering occurs at a $T_N$ which is notably less than $|\Theta_{CW}|$, the approximate temperature at which, were the substance ferromagnetic, there would have been a transition, as in the discussion in Chapter 1. At the empirical level at least, this motivates a measure of frustration for antiferromagnets $f = -\Theta_{CW}/T_c$ where $T_c$ is the temperature of any eventual cooperative ordering transition [6]. An example is $ZnCr_2O_4$, which approximates a Heisenberg pyrochlore system, and has $f = 24$, indicating strong frustration.

Conventional phase transitions are associated with a diverging correlation length, which is manifest in scattering by sharp features in reciprocal space. Another experimental signature of strong frustration is that such correlations generally do not develop. The right panel of Fig. 2.3 shows a cartoon of typical results from powder scattering on frustrated systems. It appears to show only short range correlations.

## 2.3 Bipartite lattices and the Coulomb-phase description

As mentioned at the start of this chapter, there is a class of geometrically frustrated magnets which have an attractive continuum description. The models with this property share the following construction, described in Ref. 15 and implied in Refs. 16, 17. Some of the specific details in the earlier discussion reappear here in a generalised way.

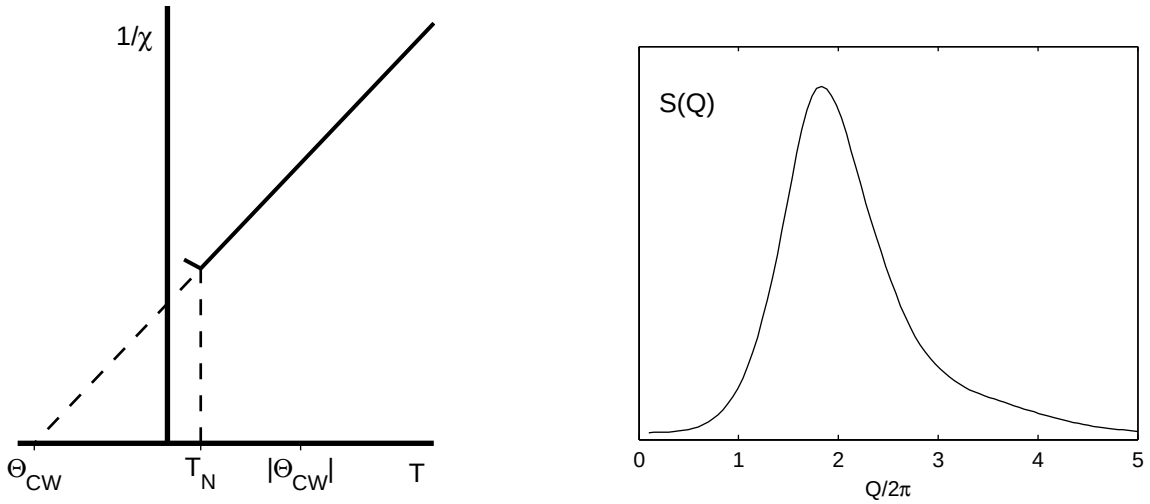


Figure 2.3: The experimental signatures of frustrated magnetic systems. Left (after Ref. 6): the inverse susceptibility follows linear behaviour to well below $T = |\Theta_{CW}|$. Right: typical behaviour of a powder diffraction sample at low temperatures, scattering seems to show only short range correlations

One starts with a bipartite lattice $B$ with $2N$ sites, and constant coordination number $q$, which we call the *parent lattice*, and then one defines the *medial lattice* $L$ such that the sites of $L$ are on the midpoints of the bonds of $B$. In this way $L$ has $Nq$ sites, each associated with two sites of $B$. Two sites on $L$ are adjacent if and only if they are associated with the same site from $B$. This procedure ensures that $L$ is a lattice of corner-sharing units, each centered on the sites of $B$. This construction is shown explicitly in Fig. 2.4 where $B$ is the diamond lattice, and $L$ is the pyrochlore lattice.

The antiferromagnetic Heisenberg Hamiltonian for a spin model with nearest neighbour interactions on a medial lattice can be written as

$$H = \frac{J}{2} \sum_{\alpha} (|\mathbf{L}_\alpha|^2 - qS^2) \,,$$

where $\alpha$ labels sites of the parent lattice, and $\mathbf{L}_\alpha = \sum_{i \in \alpha} \mathbf{S}_i$, is the total spin on the frustrated unit centered on $\alpha$. One immediate consequence of this representation is that it is possible to read off the constraint on groundstates of such models. In groundstates, the total spin on each frustrated unit is zero. Only for Ising models with odd $q$ is it not possible to satisfy this on a single frustrated unit. The representation above suggests that the excitations above the ground state are associated with particular sites of the parent lattice, rather than any particular unsatisfied bond of the medial lattice.

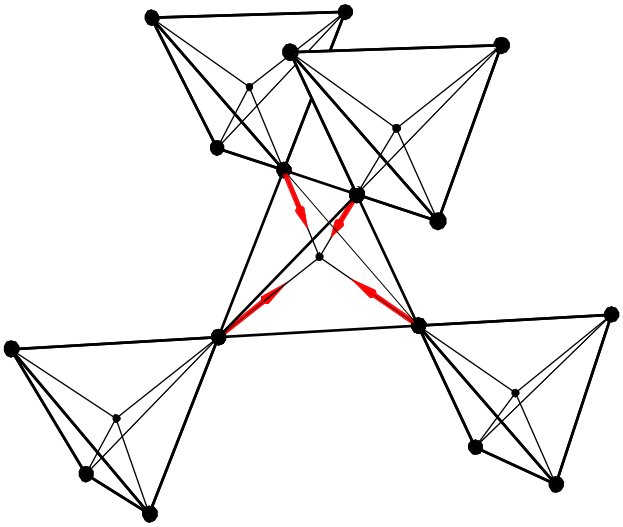

Figure 2.4: The parent lattice of pyrochlore is the diamond lattice, which is bipartite. The bonds of the diamond lattice can be oriented from one sublattice to the other in the way indicated by the red arrows.

Quite generally, any variable which lives on the bonds of a lattice can be interpreted as an instantaneous current, subject to a choice of orientation for the bonds. Since the sites of the medial lattice are the bonds of the parent lattice, we can interpret the spin degrees of freedom as a current and for $n$-component spins, the current has $n$ flavours. The benefit of this identification comes for bipartite parent lattices, where bonds can always be oriented from one sublattice to another. The arrows in Fig. 2.4 show this orientation for the diamond lattice. In the general case the groundstate constraint of the spin model above translates to demanding that the total current entering or leaving each site of the parent lattice vanish, that is, that each flavour of lattice current has zero lattice-divergence. For this reason, we refer to the lattice-currents as flux variables. Ground state spin configurations map onto the configurations of a divergence-free lattice vector-field.

The flux variables can be promoted to a continuum field by invoking the geometry of the lattice. Each oriented bond of the parent lattice can be represented by a real space vector $\mathbf{e}_i$, where $i$ labels the particular bond. The $n$ flavours of flux field, here indexed by $l$, are constructed as

$$\mathbf{B}_i^l = S_i^l \mathbf{e}_i \,,$$

where bold-type represents a spatial vector of the same dimension as the lattice. Coarse graining leads to a set of continuum vector fields $\mathbf{B}^l(\mathbf{r})$ while the lattice di-

vergence constraint is promoted to the continuum version, $\nabla \cdot \mathbf{B}(\mathbf{r}) = 0$ in three dimensions.

Since the microscopic spin length constraint ensures the microscopic fluxes have finite variance, we can appeal to the central limit theorem to conclude that the coarse grained flux should be Gaussian distributed. That is, the contribution to the free energy functional, arising entirely from the groundstate entropy, is

$$\beta H = \frac{1}{2} \int d\mathbf{r} K \sum_l |\mathbf{B}^l|^2 \text{ subject to } \nabla \cdot \mathbf{B}^l = 0 \,. \tag{2.1}$$

for some value of $K$, known as the stiffness [18, 16, 17]. Its microscopic equivalent is the inverse variance of the flux through a single frustrated unit. Since this form resembles the field energy in electromagnetism, systems described by this free energy are in a 'Coulomb phase'.

The essence of the Coulomb-phase description is that the microscopic degrees of freedom map onto a flux, and there is a microscopic constraint which can be interpreted as a conservation law for that flux. Put in these terms, it is no surprise that the Coulomb phase idea extends beyond magnetic models. Youngblood *et al.* [19] applied similar ideas to study correlations in ferroelectrics. Huse *et al.* [18] established that Eq. (2.1) describes the 3-dimensional close-packed cubic dimer model. In two dimensions, models which admit a height representation and are in the rough phase can be imagined as two dimensional versions; such models include dimer models on bipartite lattices [20].

The local constraint on flux is an energetic constraint in frustrated magnets, which can be broken by finite temperature and so it is important to consider defects, i.e. regions which violate $\nabla \cdot \mathbf{B} = 0$. For discrete spin systems, defects are discrete, and exponentially suppressed at low temperatures. For continuous systems, in contrast, one expects everywhere a non zero density of $\nabla \cdot \mathbf{B} \sim \sqrt{T}$. On the lattice scale, the divergence of the flux is equivalent to the staggered magnetisation of the frustrated units, and this staggered magnetisation should enter hydrodynamic descriptions at finite temperature. In addition, the total magnetisation is a typically a conserved quantity and will also enter the hydrodynamic description.

## 2.4 Correlations in the Coulomb-phase

We turn now to the correlations in the Coulomb phase we have discussed. One way to enforce the divergence constraint strictly is to solve the constraint by writing $\mathbf{B} = \nabla \times \mathbf{A}$, but we have already explained that defects at finite temperature are important

in systems with continuous degrees of freedom, and such a gauge representation does not deal well with finite temperature. Instead, we enforce the constraint energetically, as occurs in frustrated magnets, where the nearest neighbour exchange interaction is responsible for the constraint. We take a rectangular volume with side-lengths $L_x$, $L_y$ and $L_z$, and volume $V = L_x L_y L_z$. Assuming periodic boundary conditions, we consider the probability weight on configurations which comes from the coarse grained free energy functional:

$$P(\mathbf{B}) \propto \exp\left[-\frac{1}{2}\int_V d^3x K\mathbf{B}^2 + \beta J a^2 (\nabla\cdot\mathbf{B})^2\right] . \tag{2.2}$$

The coefficients in the exponent combine to provide a natural length scale $\xi = a\sqrt{\beta J/K}$, which will turn out to be a correlation length in what follows. Going to Fourier space according to the prescription

$$B_i(\mathbf{x}) = \frac{1}{V}\sum_{\mathbf{q}} B_i(\mathbf{q}) e^{i\mathbf{q}.\mathbf{x}} \quad ; \quad B_i(\mathbf{q}) = \int d^3x B_i(\mathbf{x}) e^{-i\mathbf{q}.\mathbf{x}} \quad ; \quad \mathbf{q} = (\frac{2\pi}{L_x}n_x, \frac{2\pi}{L_y}n_y, \frac{2\pi}{L_z}n_z) \tag{2.3}$$

the exponent becomes

$$-\frac{1}{2}\sum_{\mathbf{q}} B_i(\mathbf{q})\frac{K}{V}\delta_{\mathbf{q}-\mathbf{q}'}(\delta_{ij} + \xi^2 q_i q_j) B_j^*(\mathbf{q}') \tag{2.4}$$

and so the $q$-space correlation function for the $\mathbf{B}$ field is obtained by matrix inversion to give

$$\langle B_i(\mathbf{q}) B_j^*(\mathbf{q}')\rangle = V\delta_{\mathbf{q},\mathbf{q}'}\frac{1}{K}\left[\delta_{ij} - \frac{q_i q_j}{q^2 + \xi^{-2}}\right] . \tag{2.5}$$

### 2.4.1 Pinch points

The correlations described by (2.5) have the feature of being non-analytic at $q = 0$ when the correlation length $\xi$ is infinite. Referring back to the definition of $\xi$, this occurs in the zero temperature limit when the divergence constraint is strictly enforced. The non-analyticity at $\mathbf{q} = 0$ is known as a pinch point due to the visual appearance of these correlations, shown in Fig. 2.5. Probing systems by scattering generically measures reciprocal space correlations, and pinch points were first noticed experimentally in scattering on ferroelectric compounds. Youngblood *et al.* [19] explained the importance of the flux conservation law in understanding the singularity. The non-analyticity at $\mathbf{q} = 0$ is not a divergence, so only in wave-vector resolved scattering is it likely to be visible. At finite temperatures, the width of the pinch point is the inverse correlation length.

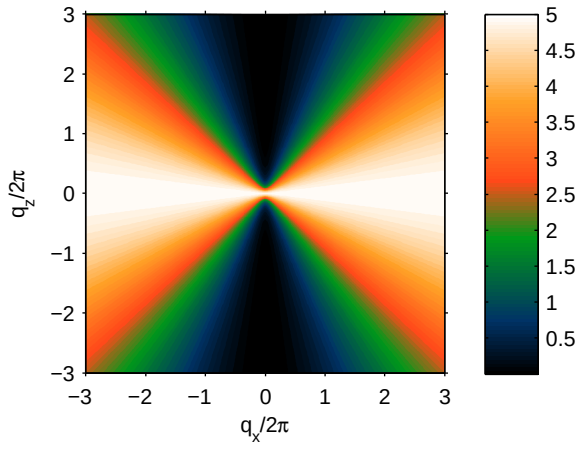


Figure 2.5: A pinch point with a narrow but non-zero width, corresponding to a large but finite correlation length. The correlation function $\langle B_z(\mathbf{q})B_z^*(\mathbf{q})\rangle$ is shown.

### 2.4.2 Real space correlations

Sharp features in reciprocal space correspond to long-range features in real space. To develop intuition for Coulomb phase systems, we will consider the finite-volume real-space correlations first. The total flux passing through a plane can be defined, along with its Fourier transform, as

$$\Phi_x(x) = \iint dydz B_x(\mathbf{x}) \quad ; \quad \Phi_x(q_x) = \int_0^{L_x} dx \Phi_x(x) e^{-iq_x x} = \lim_{q_y,q_z\to 0} B_x(\mathbf{q})$$

with equivalent definitions for flux in the $y$ and $z$ directions. It follows that the $\Phi$ correlation functions can be inferred directly from those of the $B$-field, Eq. (2.5), by taking the limit as transverse wavevectors approach 0.

$$\langle \Phi_x(q_x)\Phi_x(-q_x)\rangle = \lim_{q_y,q_z\to 0}\langle B_x(\mathbf{q})B_x(-\mathbf{q})\rangle = \frac{V}{K}\left[1 - \frac{q_x^2}{q_x^2+\xi^{-2}}\right] = \frac{V}{K}\left[\frac{\xi^{-2}}{q_x^2+\xi^{-2}}\right]$$

The inverse transform is a discrete sum in reciprocal space, which gives

$$\begin{aligned}\langle \Phi_x(x)\Phi_x(0)\rangle &= \frac{L_y L_z}{K L_x}\sum_{q_x}\frac{\xi^{-2}}{q_x^2+\xi^{-2}}e^{iq_x x} \\ &= \frac{L_y L_z}{K}\frac{1}{2\xi}\left(\frac{e^{-(2x-L_x)/2\xi}+e^{-(L_x-2x)/2\xi}}{e^{L_x/2\xi}-e^{-L_x/2\xi}}\right) \quad \text{for } 0\le x < L_x \end{aligned} \tag{2.6}$$

The behaviour of the resulting function is not obvious on inspection, but the following limits can be taken and interpreted:

$$\begin{aligned}
\text{Zero temperature limit: } \xi \gg L_x \qquad & \langle \Phi_x(x)\Phi_x(0)\rangle \sim \frac{1}{K}\frac{L_y L_z}{L_x} \\
\text{Finite } T\text{, large volume limit: } L_x \gg \xi \qquad & \langle \Phi_x(x)\Phi_x(0)\rangle \sim \frac{1}{K}\frac{L_y L_z}{2\xi} e^{-x/\xi} \\
\text{Mean square flux: } x = 0 \qquad & \langle \Phi_x(0)\Phi_x(0)\rangle = \frac{1}{K}\frac{L_y L_z}{2\xi}\coth(L_x/2\xi)
\end{aligned}$$

According to these results, in the zero temperature limit at finite volume, the fluxes through parallel planes become completely correlated, with $\Phi \sim \mathcal{O}(\sqrt{L_y L_z/L_x})$ - this is a manifestation of the flux conservation law. However at finite $T$ in sufficiently large systems, the flux correlations decay exponentially with distance. Finally, we consider the mean square flux through a plane. This is minimised in the zero temperature limit, but grows as the correlation length becomes shorter. Mathematically, this is easy to understand: As the correlation length increases, the pinch point broadens but its maximum remains the same; the total spectral weight in the pinch point, essentially the strength of the local flux fluctuations, must therefore increase.

The Helmholtz decompositon guarantees that we can express $\mathbf{B}$ as $\mathbf{B} = \mathbf{B}_\perp + \mathbf{B}_\parallel$ where $\mathbf{B}_\perp = \nabla \times \mathbf{A}$ and $\mathbf{B}_\parallel = \nabla\phi$. Using this decomposition, we can repeat the above calculation but keeping the transverse and longitudinal contributions separate:

$$\langle B_i(\mathbf{q}) B_j^*(\mathbf{q}')\rangle = \frac{V}{\kappa}\delta_{\mathbf{q}-\mathbf{q}'}\left[\left(\delta_{ij} - \frac{q_i q_j}{q^2}\right)_\perp + \left(\frac{q_i q_j}{q^2} - \frac{q_i q_j}{q^2+\xi^{-2}}\right)_\parallel\right]$$

In this equation, it is understood that $\frac{q_i q_j}{q^2}|_{\mathbf{q}=0} = 0$, which follows from considering the $\mathbf{q} = 0$ case separately. Real space correlations come from discrete fourier inversion

$$\langle \Phi_x(x)\Phi_x(0)\rangle = \frac{L_y L_z}{\kappa L_x}\left[1_\perp + \left(\frac{L_x}{2\xi}\left(\frac{e^{-(2x-L_x)/2\xi} + e^{-(L_x-2x)/2\xi}}{e^{L_x/2\xi} - e^{-L_x/2\xi}}\right) - 1\right)_\parallel\right] \quad \text{for } 0 \le x < L_x$$

This function is clearly identical to Eq. (2.6), but the first term on the right hand side is the contribution from the solenoidal part of the field configuration. We take the ratio of the mean square flux contribution ($x = 0$) from the longitudinal part to that from the transverse part and obtain $\frac{L_x}{2\xi}\coth(L_x/2\xi) - 1$. As expected, this ratio is zero for $\xi \to \infty$, the $T = 0$ limit, where all flux is solenoidal. But at any finite $T$, it asymptotes with $\frac{L_x}{2\xi} - 1$ in the limit where the system length gets much larger than the correlation length. We conclude that the proportion of flux through planes which is attributable to the solenoidal part of the field configuration is arbitrarily small for large systems.

In performing the Helmholtz decomposition on a typical finite temperature configuration with a correlation length $\xi$, both the longitudinal and solenoidal parts will have long range correlations, in the way given above, which exactly cancel beyond the scale $\xi$. For the reason that these long range correlations are an artificial consequence of the decomposition, we will avoid using the Helmholtz decomposition to describe the physical field $\mathbf{B}$, even despite the appeal of the electromagnetism analogy obtained by writing $\mathbf{B} \simeq \nabla \times \mathbf{A}$.

### 2.4.3 Dipolar correlations

More generally, we are interested in infinite systems. In this case, the correlations in real space may be obtained analytically by inverse Fourier transform of Eq. (2.5).

$$\langle B_i(\mathbf{x}) B_j(\mathbf{0}) \rangle = \frac{e^{-r/\xi}}{4\pi K} \frac{(3 + 3(r/\xi) + (r/\xi)^2) x_i x_j - r^2(1 + r/\xi)\delta_{ij}}{r^5} .$$

Correlations decay exponentially beyond a correlation length $\xi = a\sqrt{\beta J/K}$. In the zero temperature limit, the correlations are dipolar:

$$\langle B_i(\mathbf{x}) B_j(\mathbf{0}) \rangle_{\xi=\infty} = \frac{1}{4\pi K} \partial_i \partial_j \frac{1}{|\mathbf{x}|} = \frac{1}{4\pi K} \frac{3x_i x_j - r^2 \delta_{ij}}{r^5}$$

so called due to having the functional form of a dipole-dipole interaction. Since microscopically, the flux degrees of freedom were spin degrees of freedom on a frustrated lattice, this result is remarkable in showing that the spins have power-law correlations [16][17]. Exponential decay of correlations is the norm and power-law correlations normally appear as a feature of critical behaviour leading to an ordered state. Here, power law correlations appear at zero temperature in three dimensional classical systems which are not critical in the conventional sense.

## 2.5 Band structure of the adjacency matrix of frustrated lattices

To conclude this chapter, we discuss briefly how the frustrated nature of a lattice shows up in the spectrum of the adjacency matrix of the lattice. Quite generally, the energy of a nearest neighbour model can be expanded as

$$H = \frac{1}{2} J \sum_{\mathbf{q}} A(\mathbf{q}) |S(\mathbf{q})|^2$$

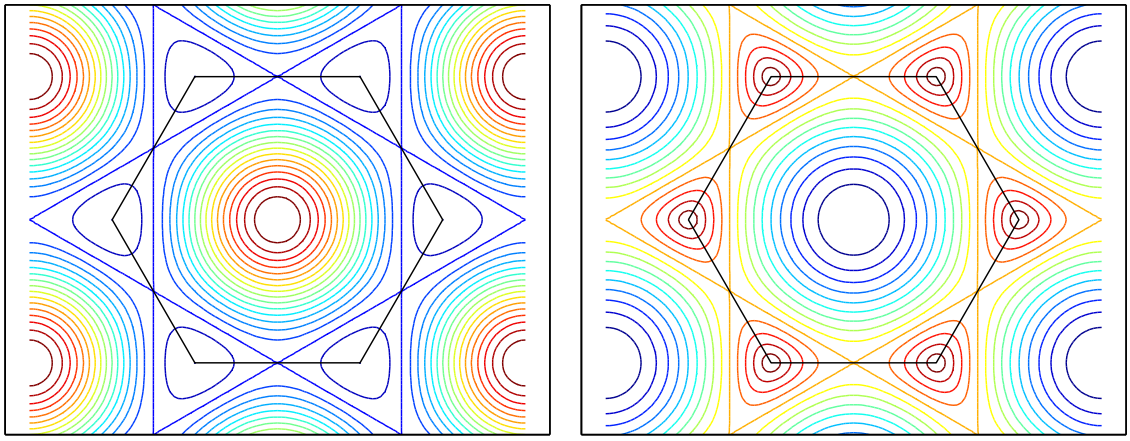

Figure 2.6: Left: spectrum of triangular lattice adjacency matrix. The maximum value is 6 at zone center and the minimum is -3 with contours every 0.5. The frustration is evident in the shallow minimum around the zone boundary. Right: second band of the spectrum of the kagome adjacency matrix. The minimum is in the zone center, at -2, which touches the flat band, the maximum at 1 is at the zone boundary, with contours every 0.2. The maxima are known as Dirac points due to their linear dispersion. This behaviour is inherited from the parent lattice of kagome, the honeycomb lattice, as discussed in the main text.

where $A(\mathbf{q})$ are the eigenvalues of the adjacency matrix of the lattice and $\mathbf{q}$ labels the eigenvalue (which may include a band index). The spin length constraint restricts the allowed Fourier coefficients, especially for Ising spins. Groundstate degeneracy can arise in two ways: it may be that the minimum of $A(q)$ is not unique, such as occurs with a flat band at the lowest energy; alternatively the spin length constraint may make it impossible to only use Fourier modes from the minimum of the dispersion, and the degeneracy then arises from the number of allowed combinations of Fourier wave vectors at some higher energy. The triangular lattice Ising model is an example of the latter case, since its ground state has energy $-J$ per spin, rather than $-3J/2$. Fig. 2.6 shows the dispersion, $A(\mathbf{q})$ of the triangular lattice. From minimum to maximum is 9, but the dispersion is much flatter in the region $-3 < A(\mathbf{q}) < -2$. In a visual way, this indicates the frustration inherent in the lattice, in the sense of having many low lying modes close in energy.

For many-component spins, macroscopic groundstate degeneracy is associated with flat bands. Flat bands are a property of frustrated lattices of the corner sharing type. The construction in terms of a parent lattice and a medial lattice have another name in the mathematical literature. There, a physicist's lattice is an infinite connected graph, say $G$. The *line graph* of $G$, denoted $L(G)$, is the graph obtained by

associating vertices with the edges of $G$; the vertices of $L(G)$ are adjacent if and only if their corresponding edges share a vertex of $G$. This construction is identical to the one outlined above for corner-sharing frustrated units. Using results from Ref. 21 we can make statements about the spectrum of $L(B)$, when $B$ is a bipartite lattice with constant coordination number $q$. We first define the Laplacian[3] on any graph $G$ with coordination number $z$ as $\Delta(G) = A(G) - zI$. Then the spectrum of $\Delta(B)$ is negative with minimum eigenvalue $-2q$, and symmetric about $-q$ due to bipartiteness. The spectrum of $\Delta(L(B))$ is identical to $\Delta(B)$, except for the addition of flat bands at eigenvalue $-2q$. In total, a fraction $(q-2)/q$ of the bands are flat.

So flat bands emerge naturally from the construction of corner sharing units. In addition, the dispersion of the non-flat bands is inherited unchanged from the parent lattice - the interpretation is that the defects or excitations in such models effectively live on the parent lattice. To give an example, the band structure of kagome is the same as that of honeycomb, its parent lattice, with the addition of 1 flat band. The honeycomb lattice is famous for its Dirac points at the zone boundary and the same features exist in the kagome spectrum. Fig. 2.6 (right) shows the second band of kagome, which is identical to the honeycomb dispersion. The dispersionful bands touch the flat bands at $\mathbf{q} = 0$.

[3]Other definitions use a different sign convention, but this choice coincides with the finite difference approximation to $\nabla^2$.

# Chapter 3

# The statistical mechanics and thermodynamics of the pyrochlore Heisenberg antiferromagnet

In this chapter, we turn to the classical pyrochlore Heisenberg antiferromagnet with only nearest neighbour interactions. The pyrochlore lattice is of the type which admits the Coulomb phase description since its sites are the midpoints of the bipartite diamond lattice as in Fig. 2.4. The model was first addressed by Villain [4] who demonstrated explicitly the construction of an infinite family of groundstates with closed loops of antiferromagnetically aligned spins. In the subsequent language of the Coulomb phase, these ground states describe a tiling of the lattice with closed loops of flux, thereby suppressing long wavelength flux fluctuations. Mean field theory [22] formally predicts a 'transition' temperature, but there is no associated ordering wavevector since the dispersion is minimized by any eigenvector in the flat bands, which are present due to the frustrated nature of the lattice. The mean field transition is of course an artifact of mean field theory, and simulations on the model confirm there is no phase transition at any temperature [23][13]. Specifically, the order-by-disorder mechanism which selects collinear states in the face-centred cubic Heisenberg antiferromagnet [24] does not cause ordering in the pyrochlore system. Selection of collinear states would give rise to quartic modes which alter the heat capacity away from the classical equipartition result. Simulations reveal a low temperature heat capacity of $\sim 0.75$ per spin consistent with there being no soft modes, and thus no selection of collinear states [13].

To go beyond thermodynamic measurements, one is generally interested in the spin correlation functions as explained in Chapter 1. By far the most successful approximation scheme to obtain correlation functions for this model is the infinite-

component model, obtained from the large-$n$ limit for $n$-component spins interacting on the same lattice. The correlations in the large $n$-limit are described by a Gaussian probability density, and unsurprisingly the same distribution appears as the leading order term in various approximation schemes. The spherical model of Berlin and Kac [25] is equivalent as demonstrated by Stanley [26], and Garanin [27] describes an elaborate diagram scheme with the same infinite component limit. Henley [17] arrives at the same result by what he calls a maximum-likelihood approach, similar to that presented below. We will use this approximation to develop a comprehensive understanding of the static properties of the the pyrochlore Heisenberg antiferromagnet.

While the application of the large-$n$ method to pyrochlore correlations is not new [16, 17, 28], we include analytical results at finite temperature which have not previously appeared, and recognise the physical meaning of the Lagrange multiplier field entering the large-$n$ calculation as the stiffness parameter in the Coulomb phase description of the pyrochlore antiferromagnet. Furthermore, the method is presented in a way amenable to a simple dynamic generalisation which is presented in Chapter 4.

## 3.1 The self-consistent Gaussian approximation

We begin with the nearest neighbour Hamiltonian on the pyrochlore lattice. Whereas the adjacency matrix $A_{ij}$ describes nearest neighbour couplings, we introduce the interaction matrix $V_{ij} = A_{ij} + 2\delta_{ij}$ which includes an on-site self-spin interaction. Due to the spin length constraint, the effect is simply to change the zero of energy. Consequently, the model is

$$H = \frac{1}{2}\sum_{ij} JV_{ij}\mathbf{S}_i \cdot \mathbf{S}_j \,.$$

The minimum energy is zero by construction, since the matrix $V_{ij}$ has non-negative eigenvalues. Full thermodynamic information is contained in the partition function, and correlation functions can be obtained by its logarithmic derivatives with respect to external inhomogeneous fields. The partition function for the classical Heisenberg model is

$$\mathcal{Z} = \int \prod_i \delta(\mathbf{S}_i^2 - 1) d\mathbf{S}_i \exp\left[-\beta \frac{1}{2}\sum_{ij} JV_{ij}\mathbf{S}_i \cdot \mathbf{S}_j\right] .$$

The non-linear constraint on the spin lengths makes calculating the partition function impossible. Approximation schemes roughly divide into those which render $\mathcal{Z}$ tractable by simplifying the interactions in the exponent, and those which simplify the

measure on the degrees of freedom. The simplest mean-field methods trade-in fluctuations as the price for performing exact integrals over the measure on the spins. The self-consistent Gaussian approximation (SCGA) which we pursue chooses instead to relax the spin length constraint as the price for retaining interactions between fluctuating variables. Here, we motivate this approximation heuristically, and simply point out that it is rigorous in the limit of infinite component spins, with $1/n$ as the expansion parameter [26] [see 29, for the pyrochlore problem]. Some further discussion is in Appendix B.

In a Heisenberg spin system, the different spin components are certainly not statistically independent, but they are uncorrelated[1]. At the level of two spin correlations functions, they may as well be independent. Motivated by this, and with a determination only to calculate two point functions, we approximate the partition function by taking each spin component as independent, and approximate the spin length constraint by a Gaussian weight.

$$\int \prod_i d\mathbf{S}_i \exp\left[-\beta \frac{1}{2} \sum_{ij} JV_{ij} \mathbf{S}_i \cdot \mathbf{S}_j\right] \to \left(\int_{-\infty}^{\infty} \prod_i ds_i e^{-\frac{1}{2}\lambda s_i^2} \exp\left[-\beta \frac{1}{2} \sum_{ij} JV_{ij} s_i s_j\right]\right)^3$$

The parameter $\lambda$ is set in a self-consistent fashion by demanding that the local constraint we aimed to replace is now enforced on average. It makes sense to consider just a single spin component from here on. The true model remains paramagnetic at all temperatures, so there is no symmetry breaking in spin space. Thus the probability weight of the configurations of real spin variables, $s_i$, is given by

$$P[\{s_i\}] \propto \exp\left[-\frac{1}{2} \sum_{ij} s_i (\lambda \delta_{ij} + \beta J V_{ij}) s_j\right] . \tag{3.1}$$

For a single component of the full model, we have $\langle (S_i^\alpha)^2 \rangle = 1/3$ on all sites. Therefore the self-consistent condition for the soft spin variables is $\langle s_i^2 \rangle = 1/3$ on all sites. Due to pyrochlore lattice symmetry, all sites are equivalent so a single $\lambda$ is sufficient. This is the basic content of the SCGA which can be applied to any isotropic spin model.

This probability weight defines a multivariate Gaussian distribution and moments on the distribution can be calculated using the standard procedures of Gaussian integration. The spin correlation function (also known as the covariance matrix in the context of multivariate Gaussians) is

$$\langle s_i s_j \rangle = [\lambda \mathsf{I} + \beta J \mathsf{V}]^{-1}_{ij} . \tag{3.2}$$

[1]Random variables $A$ and $B$ with mean zero are uncorrelated if $\langle AB \rangle = 0$

By using the equivalence of all sites on the lattice, the self consistency equation which fixes $\lambda$ can be written as sum over all sites, leading to

$$\frac{1}{3} = \frac{1}{4N}\mathrm{Tr}\left[\lambda \mathsf{I} + \beta J \mathsf{V}\right]^{-1} \tag{3.3}$$

where $4N$ is the total number of spin sites. Taken together, Eqs. 3.2 and 3.3 define the solution of the SCGA. The representation is clearly rather compact, and to understand the structure of these equations we need to understand the structure of the interaction matrix.

### 3.1.1 Lattice conventions

To make what follows unambiguous, it is necessary to establish some conventions for the labeling of the sublattices. We take throughout a conventional cubic unit cell of side 1, and use $a$ to refer to the pyrochlore nearest neighbour distance, so that $a^2 = 1/8$. The pyrochlore lattice can be imagined as a face centered cubic (fcc) lattice decorated with a tetrahedron at each fcc site. Using the following choice of fcc basis vectors

$$\mathbf{a}_1 = \frac{1}{2}(\mathbf{e}_y + \mathbf{e}_z)\ , \quad \mathbf{a}_2 = \frac{1}{2}(\mathbf{e}_z + \mathbf{e}_x)\ , \quad \mathbf{a}_3 = \frac{1}{2}(\mathbf{e}_x + \mathbf{e}_y)\ , \tag{3.4}$$

then (up to an arbitrary displacement) the corners of the tetrahedra appear at positions

$$\mathbf{c}_\mu \in \left\{\mathbf{0}\ , \frac{\mathbf{a}_1}{2}\ , \frac{\mathbf{a}_2}{2}\ , \frac{\mathbf{a}_3}{2}\right\} \tag{3.5}$$

where $\mu$ labels the four sublattices and runs from $1-4$. The sublattice labels are shown in Fig. 3.1.

Variables which live on the lattice are labeled by an fcc lattice vector $\mathbf{r}$ and a sublattice index $\mu$. Fourier transforms of such variables are defined as

$$\begin{aligned} S_\mu(\mathbf{r}) &= \frac{1}{\sqrt{N}} \sum_{\mathbf{q}\in \mathrm{B.Z}} S_\mu(\mathbf{q}) e^{i\mathbf{q}.(\mathbf{r}+\mathbf{c}_\mu)} \\ S_\mu(\mathbf{q}) &= \frac{1}{\sqrt{N}} \sum_{\mathbf{r}} S_\mu(\mathbf{r}) e^{-i\mathbf{q}.(\mathbf{r}+\mathbf{c}_\mu)} \end{aligned} \tag{3.6}$$

where $\mathbf{r}$ runs over $N$ fcc lattice sites. Although not necessary mathematically, the inclusion of the phase factors $\exp(i\mathbf{q}\cdot\mathbf{c}_\mu)$ which distinguish the sublattices is important for comparing to scattering experiments.

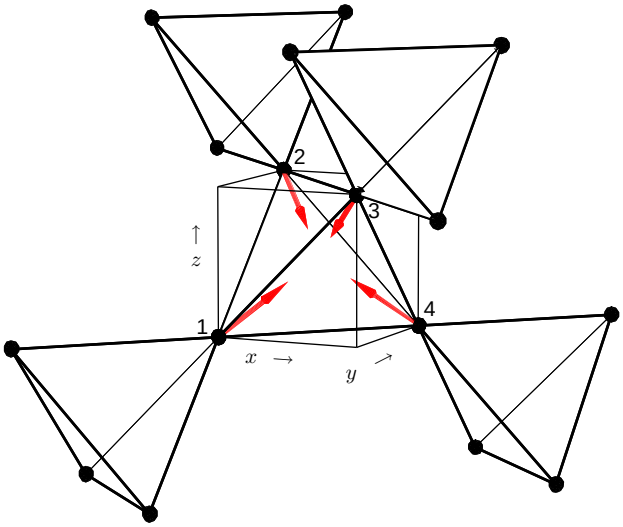


Figure 3.1: Sublattice labels for the pyrochlore lattice.

### 3.1.2 Eigenstructure of the interaction matrix

The analysis of Eqs. 3.2 and 3.3 is simplified by exploiting the translational invariance of the lattice and going to a Fourier representation using the above definitions. Under this transformation, the interaction matrix is block-diagonalised by Fourier transform to $4 \times 4$ blocks. Labeled by wavevector $\mathbf{q}$, they have the form:

$$\mathsf{V}(\mathbf{q}) = 2 \begin{pmatrix} 1 & c_{yz} & c_{xz} & c_{xy} \\ c_{yz} & 1 & c_{\overline{xy}} & c_{\overline{xz}} \\ c_{xz} & c_{\overline{xy}} & 1 & c_{\overline{yz}} \\ c_{xy} & c_{\overline{xz}} & c_{\overline{yz}} & 1 \end{pmatrix} \tag{3.7}$$

where we have borrowed the notation of Ref. [16] with $c_{ab} = \cos\left(\frac{q_a+q_b}{4}\right)$ and $c_{\overline{ab}} = \cos\left(\frac{q_a-q_b}{4}\right)$. The four eigenvalues of this matrix are $0, 0, 4 \pm 2\sqrt{1+Q}$ with $Q = c_{xy}^2 + c_{\overline{xy}}^2 + c_{yz}^2 + c_{\overline{yz}}^2 + c_{xz}^2 + c_{\overline{xz}}^2 - 3$ [22]. There are two flat bands at zero energy, and there are two bands with dispersion. In accordance with the discussion in Sec. 2.5, the bands with dispersion are inherited from the diamond lattice, and since pyrochlore has frustrated units of $q = 4$ spins, exactly half $(1 - 2/q)$ of the bands are flat. The band structure is represented visually in Fig. 3.2. Expanded at small $\mathbf{q}$, the interaction matrix has eigenvalues $0, 0, a^2q^2, 8 - a^2q^2$, and this small $\mathbf{q}$ quadratic behaviour is visible in Fig. 3.2.

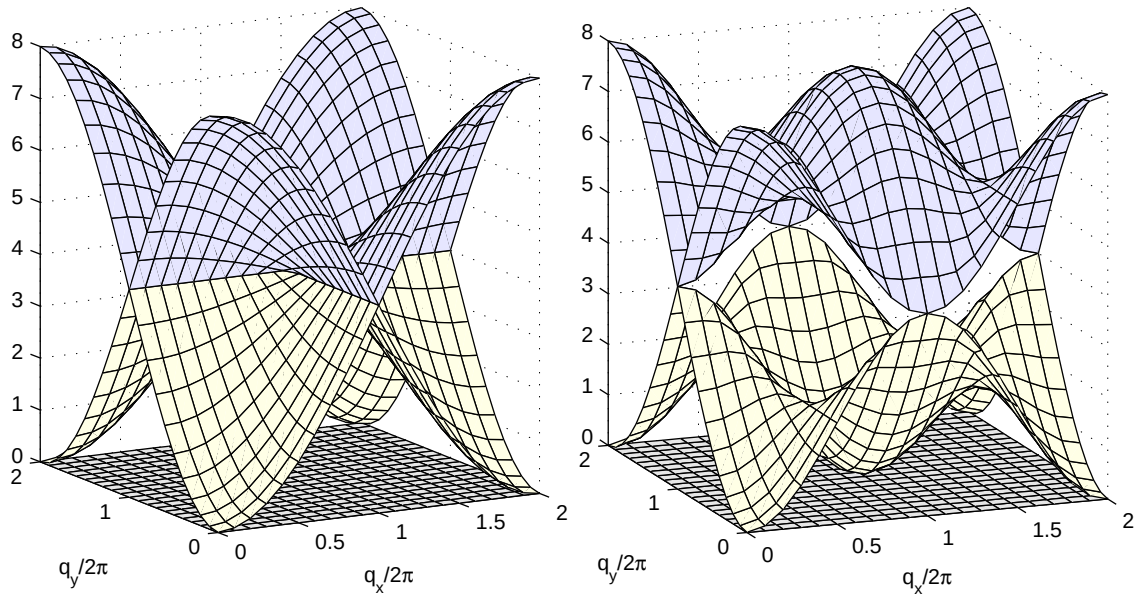


Figure 3.2: Eigenbands of the pyrochlore lattice interaction matrix in the $(q_x, q_y, 0)$ plane (left) and the $(q_x, q_y, q_x)$ plane (right). There are two completely flat bands at 0, and two bands with dispersion.

### 3.1.3 Solution of self-consistency equation

The self-consistency condition for $\lambda$, Eq. 3.3, can be expressed in the eigenbasis of the interaction matrix as

$$\frac{1}{3} = \frac{1}{4N} \sum_{\mathbf{q},\alpha} \frac{1}{\lambda + \beta J v_\alpha(\mathbf{q})} \tag{3.8}$$

where $\alpha$ is band index, and $v_\alpha(\mathbf{q})$ are the eigenvalues of the interaction matrix. At high temperatures, the interaction term is weak and the solution is $\lambda = 3 - 2\beta J + \mathcal{O}(\beta J)^2$.

At low temperatures, $\beta J \to \infty$, only the two flat bands with $v_\alpha(\mathbf{q}) = 0$ contribute, and $\lambda = 3/2 + \mathcal{O}(T/J)$. A thermodynamic phase transition would be predicted if the denominator in Eq. (3.8) ever touches zero. Since this never happens, there is no phase transition in the model. No analytic solution for $\lambda$ is possible at a general temperature but it is not difficult to solve the equation numerically, as described in Appendix A. Fig. 3.3 shows the numerical solution for the nearest neighbour model over a wide range of temperatures.

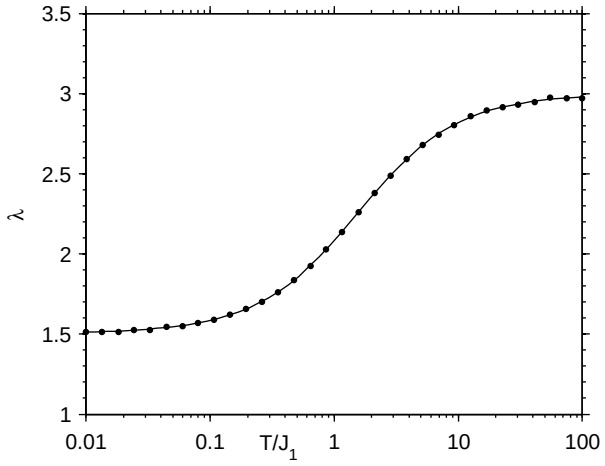


Figure 3.3: Solid line: solution of self-consistency equation, Eq. (3.8). The cross over into the strongly correlated paramagnet occurs at $T \simeq J$, but there is no phase transition. The Curie-Weiss temperature is $2J$. Points: measurements from Monte Carlo simulations on 2048 spins of the stiffness of the Coulomb phase, plotted according to Eq. (3.13)

### 3.1.4 Spin correlations

To obtain the spin correlation functions, one simply writes Eq. (3.2) in the Fourier basis. The sublattice spin correlations are

$$\langle s_\mu(\mathbf{q}) s^*_\nu(\mathbf{q})\rangle = [\lambda \mathsf{I} + \beta J \mathsf{V}(\mathbf{q})]^{-1}_{\mu\nu}$$

but more useful is the structure factor that is probed by scattering experiments. As explained in Chapter 1, the structure factor is the correlation function of the linear combination $\sum_{\mu=1}^{4} s_\mu(\mathbf{q})$, and thus is the sum of all sixteen sublattice correlation functions. For the nearest neighbour model, the spin structure factor can be calculated analytically. Reusing the notation $c_{ab} = \cos\left(\frac{q_a+q_b}{4}\right)$ and $c_{\overline{ab}} = \cos\left(\frac{q_a-q_b}{4}\right)$ and further defining $c_{(ab)} \equiv c_{ab} + c_{\overline{ab}}$, the static structure factor at any temperature is

$$S(\mathbf{q})_{\rm nn} = \frac{1}{\lambda}\left[2 + \frac{(2 - c_{(xy)} - c_{(yz)} - c_{(zy)})t\lambda + \frac{1}{2}(t\lambda)^2 - 4\left[c_{(xy)}s_z^2 + c_{(yz)}s_x^2 + c_{(zx)}s_y^2\right]}{3 - Q + 2t\lambda + \frac{1}{4}(t\lambda)^2}\right]; \tag{3.9}$$

where $t = T/J_1$ and $\lambda$ is the solution of Eq. (3.8) from Fig. 3.3. Up to an overall rescaling due to the choice of spin length, Eq. (3.9) reduces in the limit $t \to 0$ to the zero temperature form provided by Isakov *et al.* [16].

The correlations described by Eq. (3.9) have a striking visual appearance shown in Fig. 3.4. While a conventional magnet would show sharp Bragg peaks in an ordered

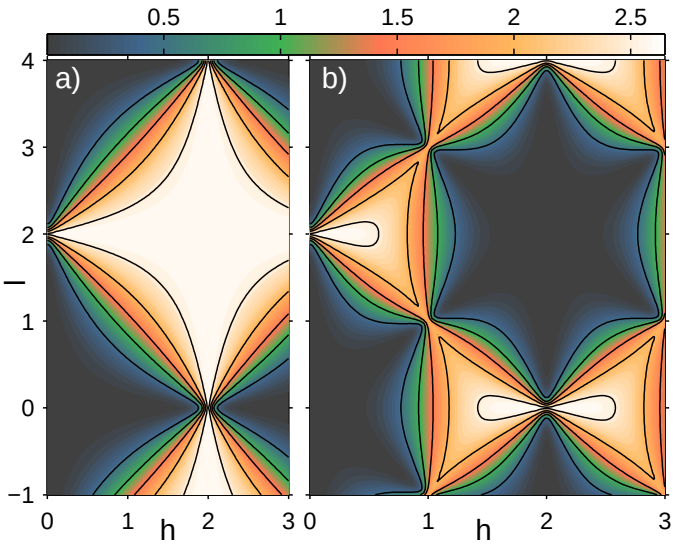


Figure 3.4: Static structure factor, Eq. (3.9), in a) $(h0l)$ and b) $(hhl)$ planes in reciprocal space at $\beta J = 100$. The sharp features at reciprocal lattice vectors are pinch points, broadened slightly by finite temperature.

phase and featureless diffuse scattering in the paramagnetic phase, here the pyrochlore antiferromagnet has sharp features in diffuse scattering. They were first noticed by Zinkin *et al.* [30] in Monte Carlo simulations for the Heisenberg antiferromagnet. The understanding of their origin is through the Coulomb phase description of Chapter 2, which we derive microscopically in the following section, and they are commonly referred to as pinch points. Along the high symmetry lines in reciprocal space which pass through the pinch points, there is zero intensity at $T = 0$, which follows from the groundstate constraint [13]. Henley [17] refers to these as *nodal lines*, and any spectral weight on nodal lines is due to contributions from excited states.

We emphasise that although Fig. 3.4 was generated from Eq. 3.9, Monte Carlo simulations give essentially identical results. In Fig. 3.5, we have plotted a comparison between Monte Carlo simulations and the SCGA at $\beta J = 10$.

### 3.1.5 $\lambda$ is the stiffness in the Coulomb-phase description

To understand the origin of the pinch points, it is necessary to extract a theory for the long wavelength degrees of freedom from the probability weight on spins in Eq. (3.1). Since the theory is Gaussian, modes at different wavevector are already decoupled, so the long wavelength behaviour is governed by the same action restricted to long wavelengths. The appropriate long wavelength degrees of freedom to consider are the

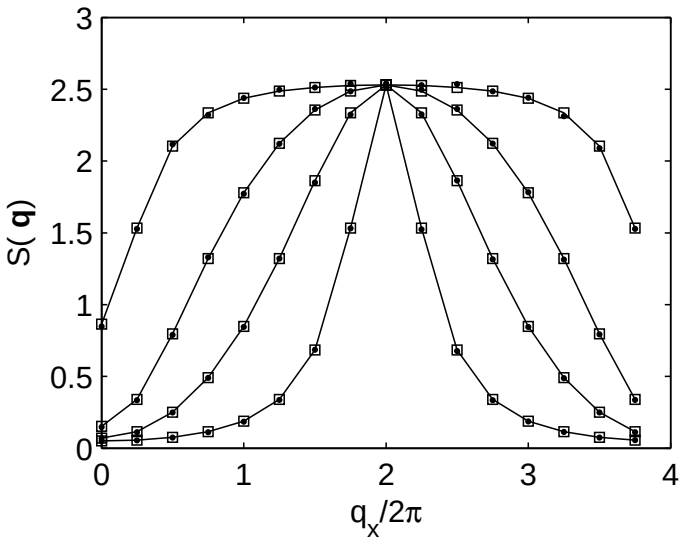


Figure 3.5: Comparison between Monte Carlo and the SCGA at finite temperature $\beta J = 10$. Lines with square symbols are from the SCGA, while points are from Monte Carlo simulations on 2048 spins. Each line is a cut at constant $q_z$ with $q_y = 0$.

flux variables discussed in Chapter 2. For the pyrochlore lattice, the long wavelength flux variables can be obtained by taking linear combinations of the spins on the four sublattices. With reference to the arrows in Fig. 3.1, we use the following orthogonal change of basis to define the flux fields:

$$\begin{pmatrix} M^\alpha(\mathbf{q}) \\ B_x^\alpha(\mathbf{q}) \\ B_y^\alpha(\mathbf{q}) \\ B_z^\alpha(\mathbf{q}) \end{pmatrix} = \frac{1}{2} \begin{pmatrix} 1 & 1 & 1 & 1 \\ 1 & 1 & -1 & -1 \\ 1 & -1 & 1 & -1 \\ 1 & -1 & -1 & 1 \end{pmatrix} \begin{pmatrix} S_1^\alpha(\mathbf{q}) \\ S_2^\alpha(\mathbf{q}) \\ S_3^\alpha(\mathbf{q}) \\ S_4^\alpha(\mathbf{q}) \end{pmatrix} .$$

This corresponds to choosing vectors of length $|\mathbf{e}_i| = \sqrt{3/4}$ in the mapping $\mathbf{B}_i^l = S_i^l \mathbf{e}_i$ described in Chapter 2. At $\mathbf{q} = 0$, the three spatial components of $\mathbf{B}$ are the three staggered magnetization order parameters which would describe a transition to $\mathbf{q} = 0$ Néel order as referred to in Refs. 31, 32. The experimentally accessible quantity is the structure factor, $S(\mathbf{q}) = 4\langle |M(\mathbf{q})|^2 \rangle$.

Applied to the interaction matrix, the change of basis gives:

$$V(\mathbf{q}) = \begin{pmatrix} 4 - G_M & -2s_y s_z & -2s_x s_z & -2s_x s_y \\ -2s_y s_z & 4 - G_x & 2s_x s_y & 2s_x s_z \\ -2s_x s_z & 2s_x s_y & 4 - G_y & 2s_y s_z \\ -2s_x s_y & 2s_x s_z & 2s_y s_z & 4 - G_z \end{pmatrix}$$

where $s_a \equiv \sin\left(\frac{q_a}{4}\right)$ and where the other parameters are:

$$\begin{aligned} G_M &= 2 - c_{(yz)} - c_{(xz)} - c_{(xy)} \overset{q\to 0}{\simeq} -4 + a^2q^2 \\ G_x &= 2 - c_{(yz)} + c_{(xz)} + c_{(xy)} \overset{q\to 0}{\simeq} 4 - a^2q_x^2 \\ G_y &= 2 + c_{(yz)} - c_{(xz)} + c_{(xy)} \overset{q\to 0}{\simeq} 4 - a^2q_y^2 \\ G_z &= 2 + c_{(yz)} + c_{(xz)} - c_{(xy)} \overset{q\to 0}{\simeq} 4 - a^2q_z^2 \,. \end{aligned} \tag{3.10}$$

The off diagonal terms have a small $\mathbf{q}$ expansion with $2s_xs_y \simeq a^2q_xq_y$. The interaction matrix at small $q$ is

$$V(\mathbf{q}) \overset{q\to 0}{\simeq} \begin{pmatrix} 8 - a^2q^2 & -a^2q_yq_z & -a^2q_xq_z & -a^2q_xq_y \\ -a^2q_yq_z & a^2q_x^2 & a^2q_xq_y & a^2q_xq_z \\ -a^2q_xq_z & a^2q_xq_y & a^2q_y^2 & a^2q_yq_z \\ -a^2q_xq_y & a^2q_xq_z & a^2q_yq_z & a^2q_z^2 \end{pmatrix} \tag{3.11}$$

Substituting this long wavelength form into Eq. (3.1) yields the long wavelength description for the flux fields, and the total magnetisation. To obtain a theory of the flux fields alone, we should integrate out $M(\mathbf{q})$, which will generate further couplings between the components of $\mathbf{B}$. In practice, the new couplings enter at $\mathcal{O}(q^4)$ and so can be ignored at sufficiently small $\mathbf{q}$. Therefore, the long wavelength theory can be read off from Eq. (3.11) by restricting the matrix to the $3 \times 3$ subblock which acts among the $\mathbf{B}$ field components, yielding:

$$\beta E = \frac{1}{2}\sum_{\mathbf{q}} \lambda|\mathbf{B}|^2 + \beta J_1 a^2 |\mathbf{q}\cdot\mathbf{B}|^2 \tag{3.12}$$

One should recognise this as the same as the Coulomb phase action discussed in Chapter 2 since the $q$-dependent terms are simply the Fourier representation of $\beta J a^2 (\nabla\cdot\mathbf{B})^2$. By following the microscopic derivation the coarse graining procedure is well defined and moreover, the stiffness parameter $K$ is the Lagrange multiplier $\lambda$ which was introduced to control the spin length in the Gaussian approximation. A microscopic constraint on spin length, equivalently the maximum flux on a bond in the flux description, is necessary for the coarse grained flux to be finite, having fluctuations of finite variance. Thus it is unsurprising that the parameter which controls the microscopic spin length should reappear as the stiffness in the long wavelength action.

In general, the strength of fluctuations of a coarse grained variable can be quite different from its microscopic counterpart due to the effect of correlations. It turns out that in our case, the strength of fluctuations of the coarse-grained flux is not dissimilar

to that expected from purely independent microscopic fluctuations. To put this on a quantitative footing, we compare the strength of the coarse-grained and microscopic flux fluctuations. We begin microscopically by considering a single tetrahedron. We can define variables $(M, \mathbf{B})$ using the same relations to the sublattices as in Eq. (3.1.5) and it follows from the definitions that

$$(M^\alpha)^2 + (\mathbf{B}^\alpha)^2 = \sum_\mu (S_\mu^\alpha)^2 \,.$$

In ground states $M^\alpha = 0$, and symmetries of the pyrochlore lattice relate the thermodynamic averages of the three components of $\mathbf{B}^\alpha$. Thus providing we are still in the paramagnetic phase, we have

$$\langle (B_x^\alpha)^2 \rangle = \frac{1}{3} \sum_\mu \langle (S_\mu^\alpha)^2 \rangle = \frac{4}{3} \langle (S^\alpha)^2 \rangle$$

This result is correct for whichever $O(n)$ family model we put on the pyrochlore lattice, providing it doesn't undergo a phase transition, thus excluding $n = 2$ only.

The long wavelength fluctuations define the stiffness $K$ by

$$\langle B_i^\alpha(\mathbf{q}) B_j^\beta(\mathbf{q}) \rangle_{\mathbf{q}=0} = \delta_{ij}\delta_{\alpha\beta}/K \tag{3.13}$$

indicating a variance of $1/K$. Within the SCGA, we have $K = \lambda$ and at zero temperature, $1/\lambda = 2\langle (S^\alpha)^2 \rangle$. The SCGA result is supported by simulations as shown in Fig. 3.3.

By expressing the $\mathbf{q} = 0$ flux component in Eq. (3.13) in real space and using the SCGA result for the stiffness, one obtains

$$2\langle (S^\alpha)^2 \rangle = \frac{4}{3} \langle (S^\alpha)^2 \rangle + \sum_{\mathbf{r}} \langle B_x^\alpha(\mathbf{0}) B_x^\alpha(\mathbf{r}) \rangle$$

which states that two-thirds of the variance of the coarse grained flux variables is explained by the microscopic variance, while one third is due to correlations. This is consistent with correlations being short ranged and liquid like to the extent permitted by the divergence-free constraint.

## 3.2 Locations of pinch points in structure factor

The experimental probe in scattering experiments is the structure factor $S(\mathbf{q}) = 4\langle |M(\mathbf{q})|^2 \rangle$. However the long-wavelength physics is in the field $\mathbf{B}$ near $\mathbf{q} = 0$. For completeness, we determine the places in reciprocal space where signatures of the

behaviour of $\mathbf{B}$ are visible in scattering. For $\mathbf{K}$ a reciprocal lattice vector we have the sublattice Fourier transforms behaving as

$$S_\mu(\mathbf{q}+\mathbf{K}) = e^{i\mathbf{K}.\mathbf{c}_\mu} S_\mu(\mathbf{q}) \; . \tag{3.14}$$

Taking $\mathbf{b}_i$ as the reciprocal basis to the $\mathbf{a}_j$ defined in 3.1.1, then a reciprocal lattice vector $\mathbf{K} = n_1\mathbf{b}_1 + n_2\mathbf{b}_1 + n_3\mathbf{b}_1$ is

$$\mathbf{K} = 2\pi[(n_2+n_3-n_1)\hat{\mathbf{x}} + (n_3+n_1-n_2)\hat{\mathbf{y}} + (n_1+n_2-n_3)\hat{\mathbf{z}}]$$

Bearing in mind the definitions of $\mathbf{c}_\mu$, the phase factors for the different sublattices in Eq. (3.14) are $1, (-1)^{n_1}, (-1)^{n_2}, (-1)^{n_3}$. It follows that

$$M(\mathbf{q}+\mathbf{K}) = \frac{1}{4}[v_0 M(\mathbf{q}) + \mathbf{v}\cdot\mathbf{B}(\mathbf{q})]$$

where

$$v_0 = 1 + (-1)^{n_1} + (-1)^{n_2} + (-1)^{n_3}$$

and

$$\mathbf{v} = \begin{pmatrix} 1+(-1)^{n_1}-(-1)^{n_2}-(-1)^{n_3} \\ 1-(-1)^{n_1}+(-1)^{n_2}-(-1)^{n_3} \\ 1-(-1)^{n_1}-(-1)^{n_2}+(-1)^{n_3} \end{pmatrix} \; .$$

If all $n_i$ are even, the structure factor in the vicinity of the reciprocal lattice point probes only the total magnetization, such as the point (222) in Fig. 3.4. If only one $n_i$ is even, it probes $\mathbf{v}\cdot\mathbf{B}$, such as the pinch point at (200), and in the two other cases, there are contributions from both $M(\mathbf{q})$ and $\mathbf{v}\cdot\mathbf{B}(\mathbf{q})$, such as the pinch point at (111).

## 3.3 Thermodynamics

The SCGA gives excellent agreement with simulations for two point correlation functions, and so one expects that any thermodynamic quantity which can be expressed in terms of the two point function will be well approximated by the SCGA. Canals and Garanin [33] report results for thermodynamic quantities for the checkerboard (Fig. 2.2) antiferromagnet, calculated within the SCGA. As we have emphasised earlier, the SCGA retains correlations while treating the spin length constraint as soft. A complementary approach to thermodynamic quantities is a single-unit approximation, where the spin-length constraint is hard, but correlations between frustrated units are assumed small and ignored which allows for an exact treatment of the thermodynamics of a single frustrated unit [34].

As described in Chapter 1, the susceptibility per spin is obtained from the fluctuations of the total magnetisation

$$\chi^z = \frac{\beta}{4N} \langle \left[ \sum_{\mu,\mathbf{q}} S^z_\mu(\mathbf{q}) \right]^2 \rangle \,,$$

and calculated within the self-consistent Gaussian framework, this is:

$$\chi = \frac{\beta}{\lambda + 8\beta J} \,. \tag{3.15}$$

The denominator is simply the eigenvalue of the quadratic form in Eq. (3.1) associated with the uniform eigenvector. At high temperatures, the solution to the self-consistency equation is $\lambda = 3 - 2\beta J + \mathcal{O}((\beta J)^2)$, and substituting this in the equation above yields the correct Curie-Weiss behaviour

$$\chi = \frac{1}{3T} \left[ 1 - \frac{\theta}{T} + \ldots \right]$$

with $\theta = 2J \equiv zJ/3$. Conversely, at very low temperatures, the susceptibility is

$$\chi = \frac{1}{8J} \left[ 1 - \frac{3T}{16J} + \ldots \right]$$

whose zero temperature limit is exact. The SCGA result interpolates between exact behaviour at $T = 0$ and $T = \infty$, as does the single-unit approximation [34]. Fig. 3.6 shows the susceptibility per spin measured from Monte Carlo simulations with the predictions from both the single-unit approximation and the SCGA. At high temperatures $T \gtrsim J/2$ the SCGA gives a better account of the simulation data whereas at low temperatures, $T \lesssim J/2$, the single-unit approximation is superior. The difference between the two approximations is maximum at $T \simeq 0.4J$, where they differ by less than 5%.

Other thermodynamic quantities can be calculated in the SCGA also, although some care must be taken: the implicit temperature dependence that $\lambda$ acquires through the self-consistency equation invalidates the usual thermodynamic relations, for example the temperature derivative of the internal energy is no longer proportional to the variance of the energy. Nevertheless, the internal energy itself is a sum of two point functions, and should be well approximated by the SCGA. For three component spins, this gives

$$U = \frac{3}{2} \sum_{\alpha,\mathbf{q}} \frac{J v_\alpha(\mathbf{q})}{\lambda + \beta J v_\alpha(\mathbf{q})} \,.$$

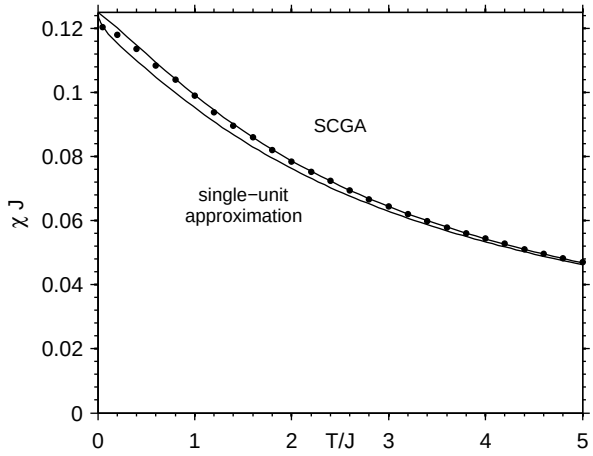


Figure 3.6: Susceptibility per spin of the nearest neighbour pyrochlore antiferromagnet Monte Carlo simulations on 16384 spins (points) and the single-unit approximation (lower line) and the SCGA (upper line).

Were it not for the presence of $\lambda$ in this equation, this would simply be the equipartition result of $\frac{1}{2}T$ per quadratic degree of freedom. The non-harmonic nature of the degrees of freedom associated with the non-linear spin-length constraint is contained in $\lambda$. With the help of the self-consistency equation, Eq. (3.8), the internal energy per spin has a particularly simple form:

$$u = \frac{3-\lambda}{2}T \tag{3.16}$$

At low temperatures when $\lambda \to 3/2$, the internal energy is $u = 3T/4$ which corresponds to a heat capacity per spin of $c = \frac{3}{4}$. This is in agreement with earlier simulations, and this is consistent with there being no soft (i.e quartic) modes in groundstates, which is evidence that the order-by-disorder mechanism does not affect pyrochlore groundstates [13]. Order-by-disorder is of course absent in the SCGA where all modes are quadratic. Fig. 3.7 shows the internal energy per spin from the SCGA and from Monte Carlo simulations. The single-unit approximation of Ref. 34 has similarly excellent agreement for the internal energy. Spurred on by Fig. 3.6, one might ask what the results from a Gaussian single-unit approximation would look like. For a probability weight

$$P(s_1, s_2, s_3, s_4) \propto \exp[-\frac{1}{2}\mu \sum_i s_i^2 - \frac{1}{2}\beta J(\sum_i s_i)^2]$$

the solution of the self-consistency equation which guarantees $\langle s_i^2 \rangle = 1/3$ is

$$\mu = \frac{1}{2}(3 - 4\beta J + \sqrt{9 + 12\beta J + 16(\beta J)^2})$$

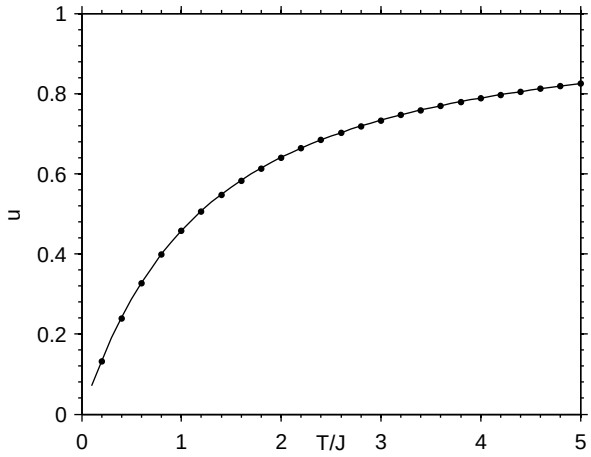


Figure 3.7: Internal energy per spin of the nearest neighbour pyrochlore antiferromagnet from the SCGA (line) and Monte Carlo simulations on 131072 spins (points).

and for three component spins, the internal energy per spin is then

$$u_{\rm su} = \frac{3-\mu}{2} T \,.$$

Since each spin is involved in two units, in comparing to the full model we take the per-spin energy as twice the above. Rather remarkably, the discrepancy between the energy calculated from the single-unit approximation and that from the full model is very small. By equating the two expressions for the internal energy as $u_{\rm SCGA} = 2u_{\rm su}$, one arrives at an approximate form for the solution of the self-consistency equation with

$$\lambda \simeq 2\mu - 3 = \sqrt{9 + 12\beta + 16\beta^2} - 4\beta \tag{3.17}$$

We have verified numerically that this approximation is extremely good: the maximum relative error is 0.3%. Using Eq. (3.17) in Eqs. (3.15) and (3.16) yields analytic expressions which are good approximations to pyrochlore thermodynamic functions across the whole temperature range. The approximate solution, Eq. (3.17), avoids the need to solve the self-consistency equation numerically, and with the analytic expression for the structure factor, Eq. (3.9), gives an analytic expression for pyrochlore correlations at all temperatures which is in excellent agreement with Monte Carlo simulations, differing by no more than a few per cent.

## 3.4 Summary

The self-consistent Gaussian approximation (SCGA) provides an excellent tool to examine the thermodynamics and correlations of the nearest neighbour pyrochlore antiferromagnet. The static structure factor can be expressed analytically at all temperatures, Eq. (3.9). The Coulomb phase description can be derived from the microscopic theory, and the stiffness parameter understood as a consequence of the microscopic spin length constraint, which is right since both effects are entropic. The approximation is easily generalised to include other interactions beyond nearest neighbour, which we do in Chapter 6. Furthermore, while the SCGA requires a numerical solution of the self-consistency equation, an approximate solution to the self-consistency equation is obtained by comparison to a Gaussian model of a single frustrated unit.

# Chapter 4

# Dynamics of the pyrochlore Heisenberg antiferromagnet

In the previous chapter, we discussed the nature of the strongly correlated low temperature paramagnetic regime of the pyrochlore nearest neighbour antiferromagnet, and showed how features of its static behaviour can be understood analytically in the framework of the self consistent Gaussian approximation. In this chapter we examine the dynamics of the low temperature paramagnet. The aim is to establish a much more comprehensive description of cooperative paramagnets with precessional dynamics than has been available previously.

In the pyrochlore Heisenberg model with precessional dynamics at low temperatures, the short-time behaviour can be viewed in terms of harmonic spinwave fluctuations in the vicinity of a specific ground state, while over longer times the background spin configuration itself changes. This second component to the motion results in decay of the spin autocorrelation function at long times, with a decay rate shown to be linear in $T$ using simulations and phenomenological arguments [12, 13].

The topic is interesting from several perspectives. First, dynamics in the paramagnetic phase of unfrustrated antiferromagnets is dominated by spin diffusion [35, 36, 37], and one would like to know the nature of the diffusive modes in the cooperative paramagnet. Second, in view of the dipolar static correlations, it is natural to ask about the wavevector dependence of the dynamical structure factor, which can be measured in experiments on single-crystal samples. Specifically, one can ask if short wavelength ground state degrees of freedom relax faster than the long wavelength ones. Little was previously known about this: the autocorrelation function of Ref [12, 13] is expressed as an integral over all wavevectors so one cannot infer about the detail of the dynamics near pinch points. Thirdly, the mapping to the Coulomb phase provides a route to compare with dynamics in other Coulomb-phase systems

and behaviour in the Heisenberg model offers a counterpoint to that in spin ice, which is represented by the Ising pyrochlore antiferromagnet with dynamics controlled by the motion of monopole excitations [38].

The results of our work in this chapter are supported by molecular dynamics simulations, the details of which are in Appendix A. Most of the work reported in this chapter was published in Ref. 39.

## 4.1 Precessional equations of motion

As discussed in Chapter 1, the natural equations of motion for classical spins are those for the quantum spin operators in the Heisenberg picture. The Hamiltonian, written as a sum over frustrated units is

$$H = \frac{1}{2} J \sum_{\alpha} \mathbf{L}_{\alpha}^2 \tag{4.1}$$

and the spin equations of motion are those discussed by Moessner and Chalker [13],

$$\frac{d\mathbf{S}_i}{dt} = -J\mathbf{S}_i \times (\mathbf{L}_{\alpha} + \mathbf{L}_{\beta}) \tag{4.2}$$

where $\alpha$ and $\beta$ label the two frustrated units in which $\mathbf{S}_i$ participates. In ground states when every frustrated unit satisfies $\mathbf{L}_{\alpha} = 0$, there is no dynamics. At small temperatures though, the local exchange field $\mathbf{H}_i = -J(\mathbf{L}_{\alpha} + \mathbf{L}_{\beta})$ will be fluctuating with strength $\langle |\mathbf{H}|^2 \rangle \sim JT$. At sufficiently low temperature, the short time dynamics is that of excited harmonic spin wave modes, whereas at longer times, anharmonic couplings to ground state degrees of freedom will cause motion in the groundstate degrees of freedom. Moessner and Chalker [13] show that taking the local exchange field as stochastic leads to a spin autocorrelation function that decays with a rate which is linearly proportional to temperature, but independent of the coupling constant $J$.

The precessional dynamics conserves both the energy and the total spin. The latter is conserved since the Hamiltonian is invariant under global rotation of all the spins. One can cast the conservation law explicitly, by writing the equation of motion for each spin as the divergence of a current. We define the spin current along the bond $i \to j$, as $\mathbf{J}_{ij} = JA_{ij}\mathbf{S}_i \times \mathbf{S}_j$ where $A_{ij}$ is the adjacency matrix. The expression for the spin current is antisymmetric, as it must be to represent a current. The equation of motion is then

$$\frac{d\mathbf{S}_i}{dt} + \sum_j \mathbf{J}_{ij} = 0\,. \tag{4.3}$$

By analogy with continuity equations in the continuum, a term $\sum_j v_{ij}$ should be thought of as the divergence of the lattice vector-field $v$, evaluated at site $i$. At the level of the frustrated units we have

$$\frac{d\mathbf{L}_\alpha}{dt} + \sum_\beta J\mathbf{S}_{\alpha\beta} \times (\mathbf{L}_\beta - \mathbf{L}_\alpha) = 0\,,$$

where $\mathbf{S}_{\alpha\beta}$ indicates the spin that is shared by units $\alpha$ and $\beta$. The final term on the left-hand side does not contribute since it generates $\mathbf{L}_\alpha \times \mathbf{L}_\alpha$ after the sum, but its presence ensures that we can also associate a spin current from $\alpha \to \beta$ on each bond of the diamond lattice of $J\mathbf{S}_{\alpha\beta} \times (\mathbf{L}_\beta - \mathbf{L}_\alpha)$.

The precessional equations of motion are hard to solve, even when linearised. While an ordered groundstate allows the use of Fourier transforms to simplify the linear spin wave problem, the nature of pyrochlore groundstates makes the problem much harder. As discussed in 1.4.2, the dynamic structure factor is probed directly in inelastic neutron scattering experiments and is what we would like to calculate. To do this, we build on the SCGA of the previous chapter, by endowing it with appropriate dynamics. In dealing with the full lattice but with soft spins, our approach is complementary to the earlier calculation of the autocorrelation function in Ref. 13, where a precessional stochastic equation was applied to a single spin.

## 4.2 Constructing a stochastic dynamics for the SCGA

The SCGA model gives an excellent account of static spin correlations, and except for the solution of the self-consistency equation, is analytically tractable. Motivated by this success, it seems fruitful to find a generalisation to include dynamic behaviour and in this section we describe how to do this. Our approach has similarities to the stochastic models of Hohenberg and Halperin [40], from which it is helpful to borrow nomenclature. By *model A*, we mean a relaxational stochastic dynamics for a non-conserved field. By *model B*, we mean a relaxational stochastic dynamics, where the drift and noise terms respect a conservation law for a field. While Ref. 40 considered continuum theories, our starting point is a microscopic lattice model.

The shortest timescale in the problem is that set by the exchange constant, $\tau_{SW} \sim J^{-1}$ where $SW$ refers to spin wave, but at longer times one expects to be able to treat the short time behaviour as stochastic noise and so approximate the equations of motion as Langevin equations. While there are many choices of dynamics which can reproduce the SCGA distribution, in order to best approximate the equations of motion (4.1), we demand a dynamics which is local and conserves the total spin.

The energy conservation of the Hamiltonian dynamics is replaced by working at the equivalent temperature. One imagines that the fast frequency modes behave as a heat bath for the low frequency modes.

We proceed quite generally by deriving the stochastic equations of motion appropriate for a lattice soft spin model with conserved total spin. The static distribution we consider is $P[\{s_i\}] \propto \exp[-\beta E[\{s_i\}]$, and the lattice is specified by an adjacency matrix, $A_{ij}$. The spin conservation law is enforced by starting with a continuity equation

$$\partial_t s_i + \sum_j J_{ij} = 0.$$

where $J_{ij}$, the spin current from $i$ to $j$, is $J_{ij} = -J_{ji}$ (not to be confused with the exchange constant $J$). This should be compared with the analogous expression for precessional dynamics, Eq. 4.3. The generalised force acting at each site is $-\frac{\partial E}{\partial s_i}$, and we assume that this force is responsible for driving the spin currents. The overall timescale for dynamic processes is set by $\Gamma$, which is the only free parameter. The dynamics we have described drives the system to minimise $E$; the thermal ensemble is maintained by introducing independent and identically distributed Gaussian noise, $\xi_{ij}$, on the currents, with a strength set by the fluctuation-dissipation relation. These assumptions lead to

$$J_{ij} = -\Gamma \sum_l \nabla_{ij,l} \frac{\partial E}{\partial s_l} + A_{ij}\xi_{ij}\,.$$

where we have defined a lattice gradient operator as $\nabla_{ij,l} \equiv A_{ij}(\delta_{jl} - \delta_{il})$. The iid current noise has correlation:

$$\langle \xi_{ij}(t)\xi_{kl}(t')\rangle = 2T\Gamma(\delta_{ik}\delta_{jl} - \delta_{il}\delta_{kj})\delta(t-t')\,.$$

Substituting the current into the continuity equation gives the equation of motion for the spins

$$\frac{ds_i}{dt} = \Gamma \sum_{jl} \nabla_{ij,l} \frac{\partial E}{\partial s_l} + \eta_i(t) \quad \text{where } \eta_i = -\sum_j A_{ij}\xi_{ij}$$

It is useful to introduce the lattice Laplacian operator, $\Delta_{ij} = \sum_j \nabla_{ij,l}$, understood in analogy with the continuum version $\Delta\phi = \nabla \cdot \nabla\phi$. For a lattice with constant coordination number $z$, the definition is $\Delta_{ij} = A_{ij} - z\delta_{ij}$. Although the bond noise was independent, the site noise is correlated to conserve the total spin. The full specification of our Langevin model is

$$\frac{ds_i}{dt} = \Gamma \sum_l \Delta_{il} \frac{\partial E}{\partial s_l} + \eta_i(t)\,. \tag{4.4}$$

with noise correlation function

$$\langle \eta_i(t)\eta_j(t')\rangle = -2T\Gamma\Delta_{il}\delta_{lj}\delta(t-t')$$

These equations should be recognised as a lattice version of *model B* for the dynamics of a conserved order parameter [40].

### 4.2.1 The pyrochlore antiferromagnet

The stochastic dynamics is tractable if the energy function is quadratic. To apply the model to the pyrochlore antiferromagnet, we take the energy function for the SCGA from Chapter 3:

$$E = \frac{1}{2}\sum_{ij} s_i(\lambda T\delta_{ij} + JV_{ij})s_j \ .$$

where $V_{ij} = A_{ij} + 2\delta_{ij}$. The stochastic equation of motion, Eq. (4.4), is

$$\frac{ds_i}{dt} = \Gamma\sum_{lm}\Delta_{il}(\lambda T\delta_{lm} + JV_{lm})s_m + \eta_i \ .$$

Considered as matrices, both the lattice Laplacian $\Delta$ and the interaction $V$ share identical eigenvectors. By making an orthogonal transformation to the eigenbasis of the interaction, denoted by tildes, the initial problem is diagonalised to the ordinary stochastic differential equation

$$\frac{d\tilde{s}_a}{dt} = -\Gamma(8-v_a)(\lambda T + Jv_a)\tilde{s}_a + \tilde{\eta}_a \ .$$

where $\langle\tilde{\eta}_a(t)\tilde{\eta}_b(t')\rangle = 2T\Gamma(8-v_a)\delta_{ab}\delta(t-t')$. The label $a$ indexes the eigenvectors of the interaction matrix, with eigenvalues $v_a$.

The ordinary stochastic differential equation is straightforward to solve [eg 41, Chap 7]. The dynamic correlation function is, for $t > 0$,

$$\langle\tilde{s}_a(t)\tilde{s}_b(0)\rangle = \delta_{ab}\frac{1}{\lambda+\beta Jv_a}e^{-\Gamma(8-v_a)(\lambda T+Jv_a)t} \ .$$

For the pyrochlore antiferromagnet, the diagonal basis is labeled by wavevector $\mathbf{q}$ and a band index $\alpha$ leading to

$$\langle\tilde{s}_\alpha(\mathbf{q},t)\tilde{s}^*_\beta(\mathbf{q},0)\rangle = \frac{\delta_{\alpha\beta}}{\lambda+\beta Jv_\alpha}e^{-\Gamma(8-v_\alpha)(\lambda T+Jv_\alpha)t} \ .$$

As with the static structure factor in Chapter 3, the dynamic correlation function contains contributions from all 4 bands, but unlike the static case, it is necessary to

separate the contributions since each band comes with its own relaxation rate. Doing so gives

$$S(\mathbf{q},t) \equiv \langle s(\mathbf{q},t)s^*(\mathbf{q},0)\rangle = \sum_{\alpha=1}^{4} g_\alpha(\mathbf{q})\langle \tilde{s}_\alpha(\mathbf{q},t)\tilde{s}^*_\alpha(\mathbf{q},0)\rangle \tag{4.5}$$

where the $g_\alpha(\mathbf{q})$ satisfy $\sum_{\alpha=1}^{4} g_\alpha(\mathbf{q}) = 4$, since they are the squared projection coefficients onto the diagonal basis of the vector (1,1,1,1). For equal time correlations, $t = 0$, this is just a decomposition of Eq. (3.9) into contributions from the four bands. After Fourier transform in time, the dynamic structure factor of this model is

$$S(\mathbf{q},\omega) = \sum_{\alpha=1}^{4} g_\alpha(\mathbf{q}) \frac{1}{\lambda + \beta J v_\alpha} \frac{2\Gamma_\alpha(\mathbf{q})}{\omega^2 + \Gamma_\alpha(\mathbf{q})^2} \,, \tag{4.6}$$

where the band dependent rate is $\Gamma_\alpha(\mathbf{q}) = \Gamma(8 - v_\alpha)(\lambda T + J v_\alpha)$. Although written rather compactly, it is worth commenting on the physical content in Eq. (4.6): within the stochastic SCGA, the dynamic structure factor has contributions from the four bands of the lattice; each band contribution has a structure factor $g_\alpha(\mathbf{q})$, a thermal occupation factor $[\lambda + \beta J v_\alpha(\mathbf{q})]^{-1}$ and a Lorentzian in frequency of width $\Gamma_\alpha(\mathbf{q})$.

To complete the derivation, we state the analytic forms of the ingredients to Eq. (4.6). As earlier, it is helpful to define $c_{ab} = \cos\left(\frac{q_a+q_b}{4}\right)$ and $c_{\overline{ab}} = \cos\left(\frac{q_a-q_b}{4}\right)$ and further define $s_a^2 \equiv \sin^2\left(\frac{q_a}{4}\right)$ and $c_{(ab)} \equiv c_{ab} + c_{\overline{ab}}$. Then with $Q = c_{xy}^2 + c_{\overline{xy}}^2 + c_{yz}^2 + c_{\overline{yz}}^2 + c_{xz}^2 + c_{\overline{xz}}^2 - 3$, the eigenvalues of the interaction matrix are $v_{1,2} = 0$ and $v_{3,4} = 4 \mp 2\sqrt{1+Q}$. The $g_\alpha(\mathbf{q})$ are, for the degenerate flat bands

$$g \equiv g_1(\mathbf{q}) + g_2(\mathbf{q}) = 2 - \frac{4}{3-Q}\left[c_{(xy)}s_z^2 + c_{(yz)}s_x^2 + c_{(zx)}s_y^2\right]$$

and for the dispersionful bands

$$g_{\frac{3}{4}}(\mathbf{q}) = 2 - \frac{1}{2}g\left(1 \pm \frac{2}{\sqrt{1+Q}}\right) \pm \frac{(2 - c_{(yz)} - c_{(xz)} - c_{(xy)}}{\sqrt{1+Q}}$$

which indeed satisfy $g_1 + g_2 + g_3 + g_4 = 4$.

At this stage, we have an analytically solved model for dynamic correlations in the pyrochlore antiferromagnet based on a stochastic dynamics which conserves the total spin. The approximation has only one free parameter, the rate $\Gamma$. In what follows, we discuss the physics contained in this model and compare it to computer simulations of precessional dynamics. Short time behaviour where the precessional nature of the dynamics is essential is not described by the stochastic SCGA but long-time behaviour in precessional dynamics is captured extremely well by our solvable stochastic model. We defer the discussion of the short time behaviour to the end of the chapter.

### 4.2.2 Diffusion of the conserved magnetisation

Total magnetisation is a conserved quantity, and in accordance with the hydrodynamic theory discussed in Sec. 1.5, we expect a diffusive mode in the vicinity of $\mathbf{q} = 0$ as the most important consequence of this conservation law.

Referring to Eq. (4.6), in the vicinity of $\mathbf{q} = 0$ only the fourth band contributes significantly. Substituting in the eigenvalue for the fourth band, small $\mathbf{q}$ limit, the dynamic structure factor has the characteristic diffusion form

$$S(\mathbf{q}, \omega) \simeq 4T\chi \frac{2Dq^2}{\omega^2 + (Dq^2)^2} \ .$$

where $\chi$ is the low temperature susceptibility per spin as discussed in Sec. 3.3, and where $D = \Gamma[8J + \lambda T]a^2$. This is the diffusion constant for the conserved spin density. When temperature is sufficiently low, $D = 8\Gamma J a^2$, effectively independent of temperature. Within hydrodynamics, the diffusion constant is a ratio of a transport coefficient to the susceptibility. Here, we have $D = \Gamma a^2/\chi$ suggesting a spin transport coefficient of $\Gamma a^2$.

We can always rewrite the diffusion form in a scaling way as

$$\beta q^2 S(\mathbf{q}, \omega) = 4\chi \frac{2D}{(\omega/q^2)^2 + D^2} \tag{4.7}$$

If $D$ is truly independent of temperature, then simulation data at multiple wavevectors and temperatures should collapse onto the same curve, Eq. (4.7). Fig. 4.1 shows this scaling collapse of simulation data at low temperatures when plotted as in (4.7), and confirms that the diffusion constant is independent of temperature. The prediction of the stochastic SCGA is also plotted in Fig. 4.1 with fitting parameter $\Gamma = 0.167$.

### 4.2.3 Relaxation of groundstate degrees of freedom

Away from $\mathbf{q} = 0$, most of the spectral weight which contributes at low temperatures to $S(\mathbf{q}, \omega)$ is in the flat bands. The flat bands of course represent the groundstate degrees of freedom. For the flat bands, the relaxation rate in Eq. (4.6) is $8\Gamma\lambda T$. Since $\lambda = \frac{3}{2} + \mathcal{O}(T/J)$, we conclude that the ground state degrees of freedom relax at a rate proportional to temperature and independent of wave vector.

There are several implications from this conclusion. First, the spin autocorrelation function is obtained as an integral over all wavevectors and so the leading behaviour of the spin auto-correlation function is $\langle s_i(0)s_i(t)\rangle \propto \exp(-8\Gamma\lambda T t)$. This agrees

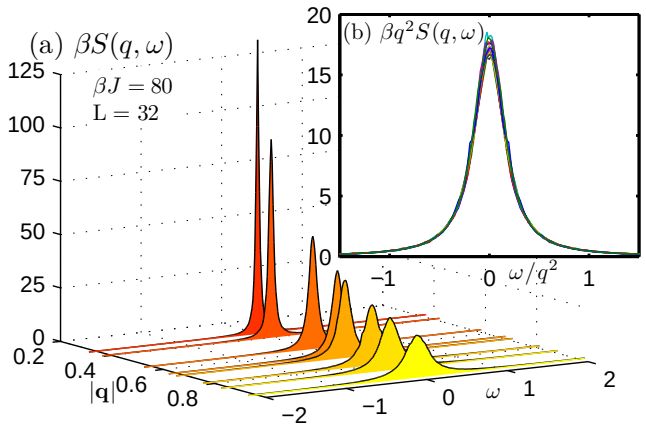


Figure 4.1: (a) The total spin is a conserved quantity and diffuses. (b) The diffusion constant is independent of temperature as shown by the scaling collapse onto one curve of multiple temperatures ($\beta J = 20, 40, 60, 80$) and wavevectors (4 of those plotted in (a) ). Also plotted on the same curve is the prediction of the stochastic SCGA with $\Gamma = 0.167$.

with the results of Moessner and Chalker [13] for the spin auto-correlation function. Secondly, in an approximation where only the flat bands contribute, the dynamic structure factor factorises as $S(\mathbf{q},\omega) = S(\mathbf{q})f(\omega)$, a possibility previously noted in [28].

The conclusions from the stochastic model are confirmed by simulations. In comparing simulations with the stochastic model, in principle there is a 5-dimensional space, $(q_x, q_y, q_z, t, T)$, to check. For reasons both of presentation and to avoid generating excessively large files in simulations, we restrict our analysis to the plane defined by $q_x = q_y$ in reciprocal space. And for detailed comparisons, it is helpful to choose a one-dimensional path in reciprocal space along which to present results. We define the path $\mathcal{P} = (0,0,0) \xrightarrow{\mathcal{P}_1} (2,2,2) \xrightarrow{\mathcal{P}_2} (0,0,2) \xrightarrow{\mathcal{P}_3} (0,0,0)$. The sections $\mathcal{P}_1$ and $\mathcal{P}_3$ are along the high symmetry nodal lines passing through pinch-points, whereas $\mathcal{P}_2$ represents more generic parts of reciprocal space, but includes the zone center. The path $\mathcal{P}$ is shown as a black line superimposed on the static structure factor in Fig. 4.2, in the $(hhl)$ plane.

To better understand the results of the stochastic SCGA, it is helpful to think in terms of the dynamical correlation function, Eq. (4.5), but normalized to unity at $t = 0$:

$$S_{\text{norm}}(\mathbf{q},t) = \frac{S(\mathbf{q},t)}{S(\mathbf{q},0)}. \tag{4.8}$$

If two wave vectors relax at the same rate then $S_{\text{norm}}(\mathbf{q}_1,t) = S_{\text{norm}}(\mathbf{q}_2,t)$ for all times.

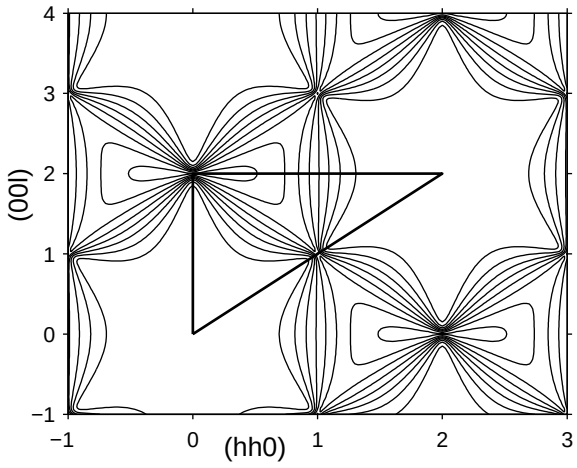


Figure 4.2: Navigating reciprocal space. The path $\mathcal{P}$ is indicated by a black line, superimposed on the contours of the static structure factor from the SCGA at $\beta J = 100$ in the $(hhl)$ plane.

The stochastic SCGA result is that as $T \to 0$ almost everywhere in reciprocal space relaxes at the same rate. This follows from the observation that as $T \to 0$, all weight is in the flat bands, and the structure factor for the flat bands only vanishes along the nodal lines. In Fig. 4.3 we have plotted $S_{\text{norm}}(\mathbf{q}, t)$ calculated from the stochastic SCGA across a reciprocal space plane. The only features are the vicinity of the nodal lines. The rest of reciprocal space, including the pinch points, takes the same value, indicative of an identical decay rate in the way just discussed. To demonstrate that the stochastic model interpolates correctly between the diffusive and flat regimes for the relaxational dynamics, we compare $S_{\text{norm}}(\mathbf{q}, t)$ measured from simulations with that calculated. Fig. 4.4 shows the comparison along the path $\mathcal{P}_1 + \mathcal{P}_2$. The agreement of the curves across multiple temperatures, wavevectors and times is evidence that the stochastic SCGA correctly describes the essential nature of the relaxational dynamics. We emphasise that there is a single fitting parameter, $\Gamma \simeq 0.167$.

Fourier transforming in time gives the dynamic structure factor, which at most wavevectors is well described by a Lorentzian line shape centered on $\omega = 0$ with a line-width linear in temperature. As further evidence for our conclusions, in Fig. 4.5 we have plotted the relaxation rate fitted from simulations over a wide range of temperatures, at three different wavevectors which include a pinch point. The linear dependence on temperature is clear, as is the independence of wavevector. The Lorentzian line shape obtained from simulations is shown in the inset.

While there is no doubt that simulations of precessional dynamics are in excellent

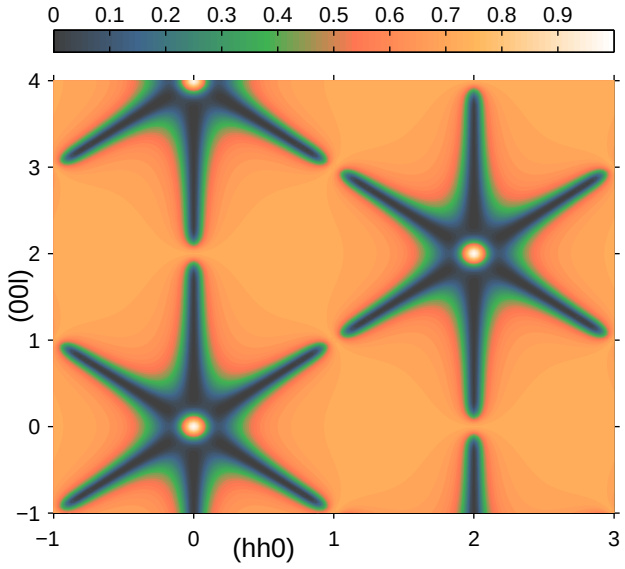


Figure 4.3: $S_{\text{norm}}(\mathbf{q}, t)$, Eq. (4.8), at $\beta J = 50$ and $t = 8$ across the $(q_x, q_x, q_z)$ plane in reciprocal space, calculated by the stochastic SCGA (with $\Gamma = 0.167$). At low temperatures, almost all points decay at the same rate, as shown by the large flat areas. The exception is along nodal lines, where dynamics is typically fast but controlled by spin diffusion in the vicinity of $\mathbf{q} = 0$.

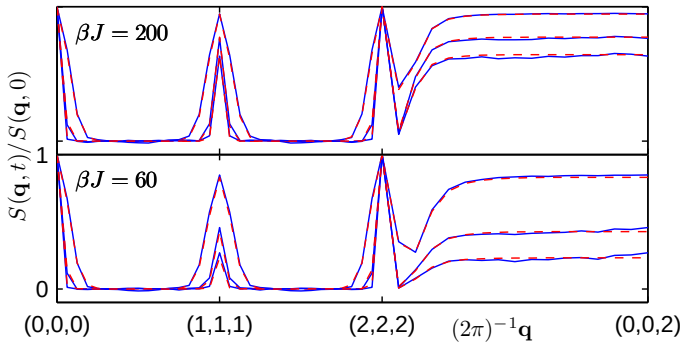


Figure 4.4: Solid blue lines: normalized correlation function $S_{\text{norm}}(\mathbf{q}, t)$ from simulations. Dashed red lines: prediction from 4.5 with $\Gamma = 0.167$, shown at 3 different times.

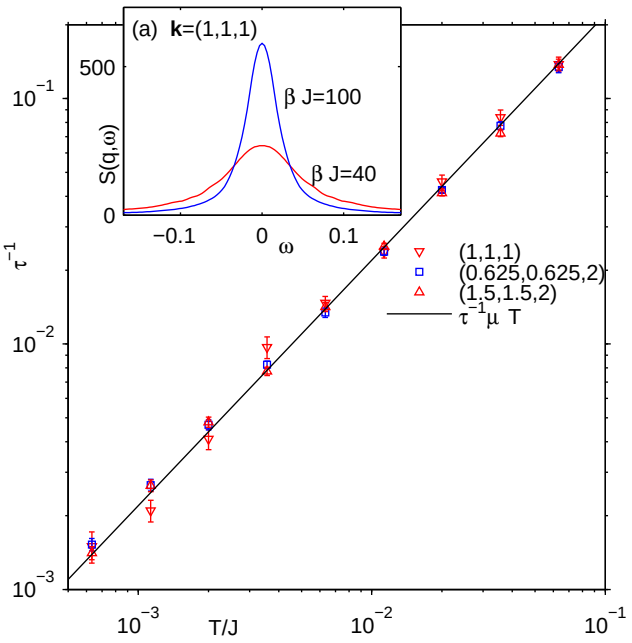


Figure 4.5: Relaxation rate across three orders of magnitude in temperature. Inset: Lorentzian lineshape for $S(\mathbf{q},\omega)$ narrows as $T \to 0$, shown for $\beta J = 40,\ 100$.

agreement with the stochastic SCGA for low temperature dynamics at large times, these results prompt certain questions. For example how is it that the pinch point correlations relax at the same rate as much shorter wavelength correlations? To answer this, we will discuss the dynamics of the long wavelength degrees of freedom in the next section.

## 4.3 Continuum dynamics

As explained in Chapter 2, the Coulomb-phase description of frustrated magnets is an appealing way to understand and have intuition for the physics of such models. The relevant degrees of freedom are the long-wavelength flux fields, $\mathbf{B}$, which we defined for the pyrochlore magnet in Chapter 3. By restricting the results of the stochastic SCGA to the field $\mathbf{B}$ at small $\mathbf{q}$, we obtain the correlation functions

$$\langle B_i(\mathbf{q},t)B_j(-\mathbf{q},0)\rangle = \frac{1}{\lambda}\left[\left(\delta_{ij}-\frac{q_iq_j}{q^2}\right)e^{-8\Gamma\lambda Tt}+\left(\frac{q_iq_j}{q^2}-\frac{q_iq_j}{q^2+\xi^{-2}}\right)e^{-(8-a^2q^2)\Gamma(Ja^2q^2+\lambda T)t}\right]. \tag{4.9}$$

The contribution from the flat bands is the first term on the rhs, while the second term represents contributions from the third band. The fourth band does not contribute.

The relaxation rate is never less than $8\Gamma\lambda T$, and there is no slowing down at the pinch points.

This result was arrived at from a microscopic model, but we now show how to derive it starting from the continuum. The dynamics of the classical model is driven by the presence of violations of the groundstate constraint, made clear in Eq. (4.2) since there is no time-dependence if all tetrahedra satisfy the ground state constraint. In deriving stochastic dynamics for the continuum model, we will interpret our equations to ensure that the dynamics is again driven solely by the violations, which in the continuum is the field $\rho = \nabla \cdot \mathbf{B}$, which we call the monopole density.

The flux fields are not microscopically conserved, and so our approach mimics *model A* [40]. Working from an effective Hamiltonian

$$\mathcal{H} = \frac{1}{2} \int d^3x \lambda T \mathbf{B}^2 + Ja^2 (\nabla \cdot \mathbf{B})^2$$

the stochastic dynamics for a non-conserved field $\mathbf{B}$ is

$$\begin{aligned} \frac{\partial \mathbf{B}}{\partial t} &= -8\Gamma \frac{\delta \mathcal{H}}{\delta \mathbf{B}} + \zeta(\mathbf{x}, t) \\ &= -8\Gamma\lambda T \mathbf{B} + 8\Gamma J a^2 \nabla(\nabla \cdot \mathbf{B}) + \zeta(\mathbf{x}, t) \end{aligned} \tag{4.10}$$

where the factor of 8 is included to unify notation with the lattice model. The noise correlation function takes the usual fluctuation-dissipation form $\langle \zeta(\mathbf{x}, t) \zeta(\mathbf{x}', t') \rangle = 16T\Gamma\delta(\mathbf{x}-\mathbf{x}')\delta(t-t')$. The problem is linear, and one can solve this equation to derive correlation functions which are exactly those in (4.9), except the factor $(8 - a^2q^2)$ is replaced by 8. However, insight comes from taking the divergence of Eq. (4.10) and recognising the resulting equation as a continuity equation for the monopole density $\rho$

$$\frac{\partial \rho}{\partial t} + \nabla \cdot \mathbf{j} = 0$$

with the monopole current density

$$\mathbf{j} = 8\Gamma\lambda T \mathbf{B} - 8\Gamma J a^2 \nabla \rho + \zeta(\mathbf{x}, t)\,.$$

There are two observations to make. The first is that all terms in (4.10) appear in the expression for the monopole current density, and so the dynamics is entirely driven by the defects. Stated in reverse, in integrating (4.3) to recover Eq. (4.10), the integration constant is any term which is purely solenoidal, which would correspond to dynamics entirely within the groundstate manifold. We set such contributions to zero to ensure a defect driven dynamics.

The second observation is to consider the contributions to the monopole current. There is a diffusive term, with diffusion constant $8\Gamma Ja^2$, which acts to eradicate gradients in the monopole density, which is entirely expected, as is the thermal noise which maintains the ensemble. The first term is more interesting though, since it describes monopole density advected along the flux lines of the $\mathbf{B}$ field at a rate given by the local strength of the flux. We interpret this term as the response of the monopole density to an entropic force, which mimics the usual response of charge to a field. Since our simulations reveal the flat temperature dependent relaxation rate dominates the dynamics, it is clear that the entropic contribution is essential to a complete understanding and a purely diffusive dynamics is insufficient.

Jaubert and Holdsworth [38] have argued similarly for the discrete equivalent spin ice, that a full description must take into account the network of Dirac strings, essentially an entropic term. For a nice discussion on the role of entropic forces in the dynamics of 2D models with height representations, see Ref. [42].

The proposed form of the monopole current density can be tested explicitly in simulations of precessional dynamics by starting in a non-equilibrium configuration and comparing the evolution to that predicted by (4.10). We start from a state with strong flux polarization and measure the return to equilibrium. The initial state was strongly polarised in the $B_z$ direction, corresponding to $\mathbf{S}_1 \simeq \mathbf{S}_4 \simeq -\mathbf{S}_2 \simeq -\mathbf{S}_3$, thermalised by Monte Carlo at a variety of temperatures. Fig. 4.6 shows the return to equilibrium of $B_z(t)$ with precessional dynamics, and the scaling collapse when time is rescaled by temperature, in strong support of entropically driven dynamics for the flux fields.

The other long-wavelength field is the total magnetisation, whose dynamics we established was purely diffusive with diffusion constant $8\Gamma Ja^2$. To emphasis the difference between the long wavelength dynamics of the staggered magnetisation $\rho$, and the magnetisation $m$, we write their stochastic equations together:

$$\frac{\partial \rho}{\partial t} = D\nabla^2 \rho - 8\Gamma\lambda T\rho + \xi(\mathbf{x}, t) \tag{4.11}$$

$$\frac{\partial m}{\partial t} = D\nabla^2 m + \xi(\mathbf{x}, t) \tag{4.12}$$

where the noise is conserved, with $\langle \xi(\mathbf{x}, t)\xi(\mathbf{x}, t)\rangle = 16\Gamma T\nabla^2\delta(\mathbf{x} - \mathbf{x}')\delta(t - t')$.

Our results are confirmed by simulations, but contradict previous expectations that the conserved nature of the flux field should make relaxation at long wavelengths slow [17]. The reason for the difference is that the dynamics is entirely driven by defects, so quite different from the conclusions one would draw for dynamics within

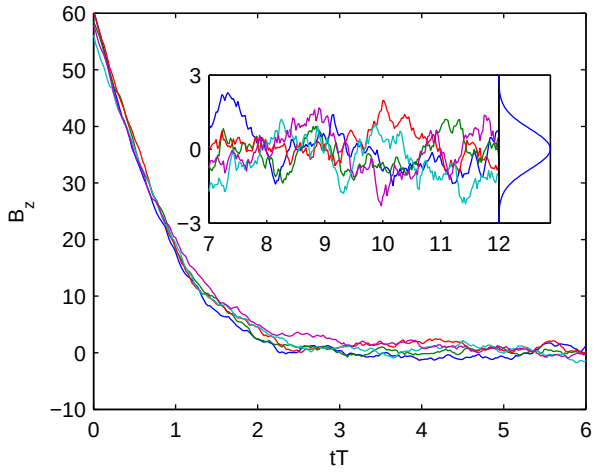


Figure 4.6: Flux relaxation from a non-equilibrium configuration shown for single simulation runs. The initial condition is strongly flux-polarised in $B_z$, but with different starting temperatures (energy densities) of $\beta J = 40, 60, 80, 100, 200$. The zero of time is chosen so that all systems have the same $B_z$. The time axis has been rescaled by temperature to demonstrate the scaling collapse of the relaxation to equilibrium, in strong support of Eq. (4.10). Inset: equilibrium fluctuations at later times, with $\exp(-\lambda B_z^2/2)$ plotted for comparison - the initial state corresponds to a $75\sigma$ fluctuation.

the ground state, such as in 2D close packed dimer models [20]. Furthermore, the defects essentially experience long range entropic forces, which facilitate the relaxation of the long wavelength degrees of freedom, much faster than would occur with a purely diffusive defect dynamics.

## 4.4 Short time behaviour

To conclude this chapter, we discuss the short time behaviour of the dynamics which is intrinsically precessional. The stochastic model is of no help here for two reasons: first, the short time behaviour is excluded in the formulation of the stochastic model and replaced by Gaussian noise; second, the precessional nature of the microscopic dynamics is excluded, so although successful for dissipative modes, the stochastic model would not predict the propagating spin wave modes which dominate the dynamics in in ordered ferromagnets or antiferromagnets.

At low temperatures, the short time dynamics is governed by the full equations of motion, linearised about the instantaneous ground state spin configuration. Since the spin configuration has no long range order, conventional spin waves with well defined

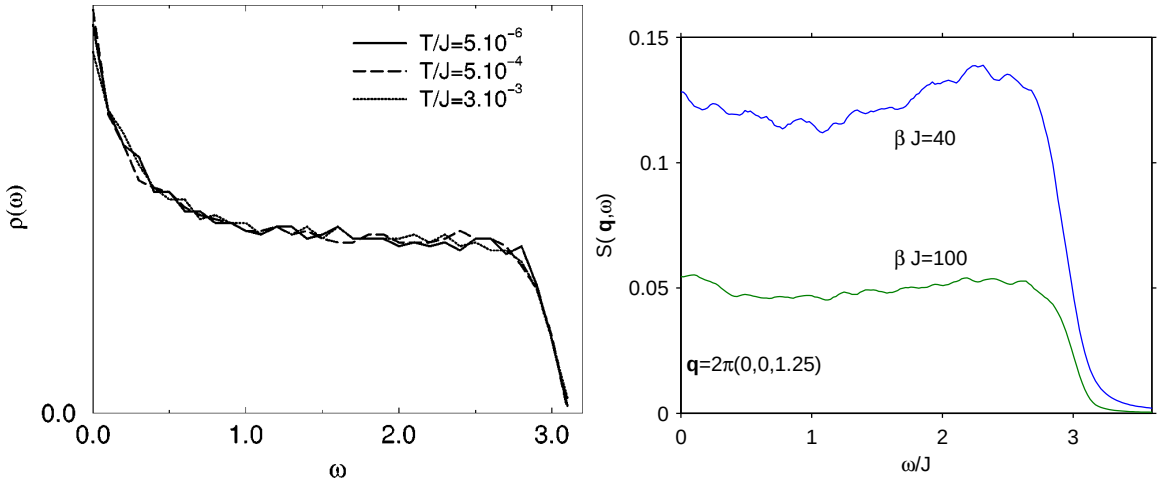


Figure 4.7: Left: density of states for harmonic spin waves, from Ref. 13. Right: $S(\mathbf{q},\omega)$ along a nodal line at (0,0,1.25) at two temperatures. Both figures share the cutoff at $\omega \simeq 3J$. At finite temperature, spin waves acquire a lifetime which explains the rounded edge in $S(\mathbf{q},\omega)$ near the cutoff.

wavevector do not exist, but we use 'spin wave' to refer to these harmonic modes. Whatever the nature of these spin waves, one does not expect them to be sharp in reciprocal space, and so asking for the spin wave dispersion does not make sense, but a valid question to ask is about the density of states in frequency.

The spin wave density of states was calculated numerically by Moessner and Chalker [13] by diagonalising the matrix which governs the linearised dynamics. Though they did not comment on it, the density of states goes abruptly to zero at $\omega \simeq 3J$. Since it is possible to construct periodic groundstates with a spinwave dispersion up to $4J$, the cutoff in frequency is not a trivial one, but due to the nature of typical groundstate configurations. Fig. 4.7 shows the density of states from Ref. 13 alongside the results of our simulations for $S(\mathbf{q},\omega)$, measured along a nodal line where the signal is not obscured by the groundstate degrees of freedom. It would be nice to understand the cutoff frequency better.

Despite the disordered background for the harmonic excitations, we do notice some structure in the finite frequency modes in reciprocal space. We think of these as remnants of the hydrodynamic spin wave modes which would exist, were the system in a fully flux polarised state. More explicitly, one can imagine the ensemble of systems with total flux $\overline{\mathbf{B}}$, obtained by coupling the macroscopic flux to a thermodynamically conjugate field. Then the dynamic structure factor generically depends on $\overline{\mathbf{B}}$, so we write $S(\mathbf{q},\omega,\overline{\mathbf{B}})$. For large $\overline{\mathbf{B}}$, propagating spin waves are guaranteed to exist. As $\overline{\mathbf{B}} \to 0$, the spin waves at different $\mathbf{q}$ mix, but some structure can remain.

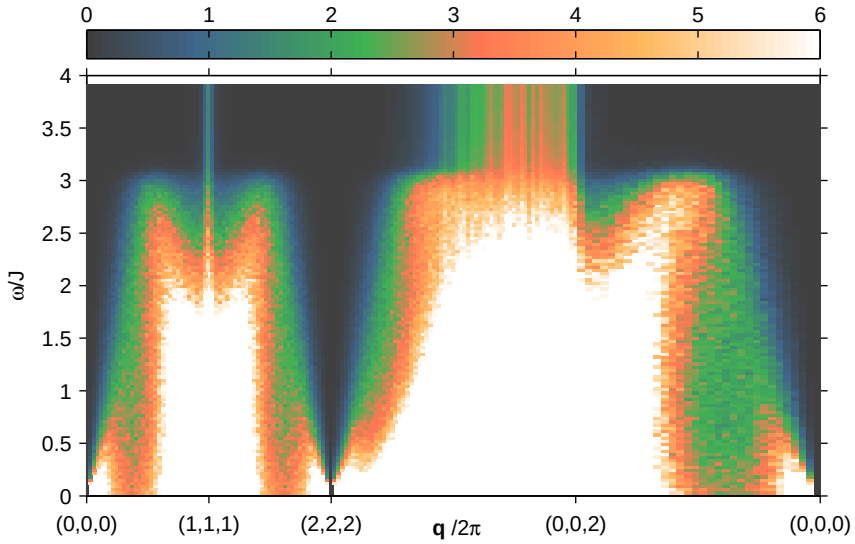


Figure 4.8: Intensity map of $S(\mathbf{q},\omega)$ along the path $\mathcal{P}$ (see Fig. 4.2) $\beta J = 320$. Ridges of intensity indicate remnant propagating spin wave modes, while the density of states has a cutoff at $\omega \simeq 3J$. The units of the colour scale are arbitrary, while the range is chosen to emphasise the high frequency behaviour. The low frequency behaviour saturates the scale, except in the vicinity of the diffusive peaks.

The results from simulations which support this idea are visual. In Fig. 4.8, we have plotted $S(\mathbf{q},\omega)$ along the path $\mathcal{P}$ (see Fig. 4.2) with an intensity scale chosen to highlight the high frequency behaviour. We claim that the ridges of intensity which emerge from the pinch points are the remnant propagating modes. The cutoff at $\omega \simeq 3J$ is visible, as are the wide Lorentzian tails from the relaxation of the groundstate degrees of freedom which extend to high frequency. In Fig. 4.9, we show the same features by plotting at fixed frequency $\omega = 2.95J$ across a plane in reciprocal space. The rings of intensity around the pinch points are reminiscent of the dispersion of propagating spin waves.

## 4.5 Summary

We have considered wavevector and frequency resolved dynamics of the classical pyrochlore antiferromagnet. The relaxational behavior is well captured by a stochastic model which conserves total spin. Spin diffuses with a diffusion constant independent of temperature, and entropic forces drive currents of monopole density which relax configurations at a rate independent of wave vector and proportional to temperature. With a single fitting parameter, the stochastic model correctly interpolates between

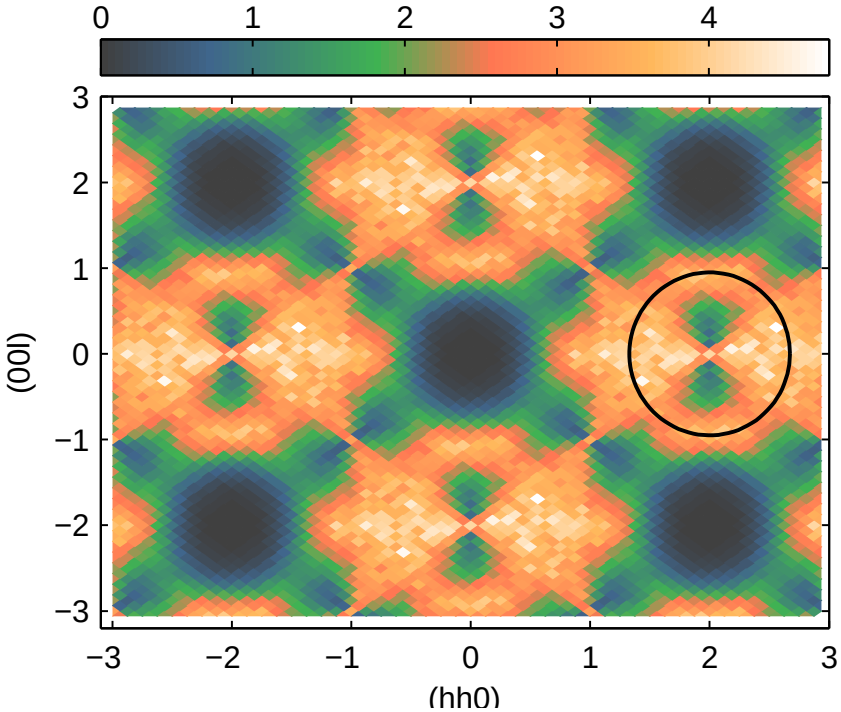


Figure 4.9: Intensity map of $S(\mathbf{q}, \omega)$ in the $(q_x, q_x, q_z)$ plane at $\beta J = 500$ and $\omega = 2.95J$. A black circle indicates the rings of intensity centered on pinch-points which correspond to remnant propagating modes. The units of the colour scale are arbitrary.

the diffusive regime and the flat relaxation of the ground state degrees of freedom. An analytic expression for the dynamic structure factor was obtained which we hope can be useful to experimentalists as a minimum theoretical model with which to compare results. The derivation of the stochastic SCGA, provided in detail for the nearest neighbour pyrochlore model, can be generalised to other frustrated lattices, and those with flat bands should show the same low temperature relaxation of the groundstate degrees of freedom, driven by entropic forces acting on the defects.

We have also considered the short-time behaviour, controlled by precessional dynamics. Simulations reveal that the density of states of finite-frequency modes has a cut-off at $\omega \simeq 3J$ in typical groundstates and there is evidence in the simulations for remnant propagating modes.

# Chapter 5

# Experimental review

In the previous chapters we have developed a comprehensive understanding, supported by numerical simulation, of the statics and dynamics of the nearest neighbour classical pyrochlore Heisenberg antiferromagnet, a paradigmatic model of frustrated magnetism. We first summarise the predictions which come from this understanding. Thermodynamically, there is no phase transition: the model remains paramagnetic at all temperatures. As temperature is decreased, the static structure factor develops pinch points with a width which vanishes as $\sim \sqrt{T/J}$. The pinch points have their origin in a mapping of spins to a lattice flux, which is divergenceless in the limit $T \to 0$. The low temperature dynamics is dominated by the relaxation of ground state degrees of freedom with a relaxation rate that is linear in temperature and independent of wave vector. Low temperature correlations in the vicinity of pinch points behave in the same way as at more generic places in reciprocal space. Entropic forces are responsible for the relaxation of long wavelength ground state correlations. Finite frequency excitations are characterized by a density of states which is approximately flat up to a cutoff $\omega \simeq 3J$.

In this chapter we assess to what extent experimental results are consistent with these predictions. We will find that when wave vector resolved information is available for candidate materials, it does not agree with the predictions in earlier chapters. We will use this to motivate the study of further neighbour interactions for understanding scattering in various spinel compounds. The consequences of further neighbour interactions are developed in Chapter 6 using the methods of Chapters 3 and 4.

Deviations from the classical Heisenberg model are of course inevitable. For a start, real systems have quantum spins with finite $S$, and phenomena on the scale $T \lesssim JS$ are sensitive to quantum effects. Even discarding the quantum nature, most of the compounds we discuss undergo ordering transitions at sufficiently low temperature, indicating the existence of a mechanism to lift the classical degeneracy. Being due to

secondary interactions beyond nearest neighbour exchange, this mechanism can vary considerably across materials.

To organise the presentation, we will discuss pyrochlore compounds, and subsequently spinels, both of which realise a pyrochlore lattice.

## 5.1 Pyrochlores

The pyrochlore lattice inherits its name from the pyrochlore compounds. In these compounds, $A_2B_2O_7$, both the $A$ and $B$ sites realize a pyrochlore lattice. A wide variety of compounds exist with magnetic ions on either or both sites. Typically, the $A$ site is a magnetic rare earth ion, and the significant magnetic interactions will include magnetic dipole interactions which can dominate exchange terms. Furthermore, crystal field terms readily spoil the Heisenberg symmetry of the degrees of freedom. For these reasons, the rare earth pyrochlores do not yield a nearest neighbour Heisenberg model, suitable for comparison with earlier chapters, but nevertheless they include some of the most interesting phenomena in frustrated magnetism. The spin ice compounds $Ho_2Ti_2O_7$ and $Dy_2Ti_2O_7$ are examples where strong crystal fields and ferromagnetic nearest neighbour (NN) exchange result in an effective Ising degree of freedom with antiferromagnetic NN interactions, a mapping which survives, somewhat surprisingly, even after taking into account strong long range dipolar interactions [43]. The magnetostatic dipolar interaction leads to a prediction that spin ice supports magnetic monopole excitations which interact by a true magnetic interaction, in addition to entropic interactions [44]. Magnetic monopoles of this kind have been found experimentally in $Dy_2Ti_2O_7$ [45, 46] and the pinch-point scattering has also been confirmed in $Ho_2Ti_2O_7$ [47]. The essential behaviour of spin ice is captured by the minimal theoretical model of Ising degrees of freedom with dipolar interactions, with Coulomb phase correlations. For an extensive review of the properties of other magnetic pyrochlore oxides, see Ref. 48.

$Y_2Ru_2O_7$ has a pyrochlore crystal structure, but the magnetism derives from the transition metal $Ru^{4+}$ ions with $S = 1$ on the $B$ sites, and in fact is a good approximation to a Heisenberg nearest neighbour antiferromagnet. The Curie-Weiss temperature is $\Theta_{CW} \simeq -1100$ K and the compound undergoes a magnetic phase transition at $T_N = 77K$. The large discrepancy in temperature scales is evidence for strong magnetic frustration. van Duijn *et al.* [49] report neutron scattering measurements on a powder sample, observing that spin correlations develop as high as 300 K and evolve little upon cooling towards $T_N$. They find excellent agreement between the

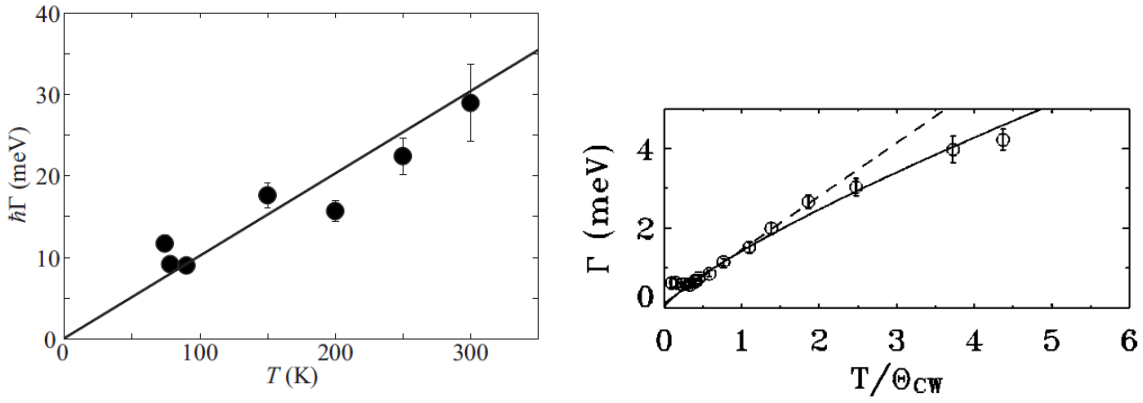


Figure 5.1: Left: spin relaxation rate in $Y_2Ru_2O_7$averaged over all wavevectors, from Ref. 49. Right: spin relaxation rate in $CdFe_2O_4$, from Ref. 50. The linear dependence on temperature is consistent with the conclusions of chapter 4 and Refs. 13, 12

neutron intensity and that predicted from zero temperature structure factor in Ref. 16 (i.e. the $T \to 0$ limit of Eq. (3.9)). Unfortunately the spherical averaging inherent in powder-sample measurements makes a detailed comparison with theoretical models impossible. The predicted pinch point singularities would not show up in a spherical average so the existence of pinch-points in $Y_2Ru_2O_7$ remains unclear. Without full wave-vector information, it is also impossible to test the dynamic predictions of the previous chapter in their entirety. Nevertheless, Fig. 5.1 shows the inelastic width derived from an integral over all wavevectors, thereby probing the decay of the spin autocorrelation function. The data is consistent with, though not in overwhelmingly agreement with, a linear dependence on temperature as predicted for the autocorrelation function by Moessner and Chalker [13]. The mechanism for the phase transition is not fully understood but local anisotropy is thought to be responsible [49].

## 5.2 Spinels

The spinel compounds, $AB_2O_4$, have a structure in which the $B$-sites form a pyrochlore lattice. In the compounds we will discuss, the $A$ ions are non-magnetic and the magnetism lives on the pyrochlore lattice. We will mainly refer to the series of chromite compounds, $ZnCr_2O_4$, $MgCr_2O_4$and $CdCr_2O_4$, for which single crystal inelastic neutron scattering measurements have been performed by a number of groups. $Cr^{3+}$ is magnetic with $S = 3/2$, and provides a good realization of a Heisenberg degree of freedom and so such measurements are a crucial test of the applicability of the ideas in previous chapters to real compounds. For a review of frustrated spinels,

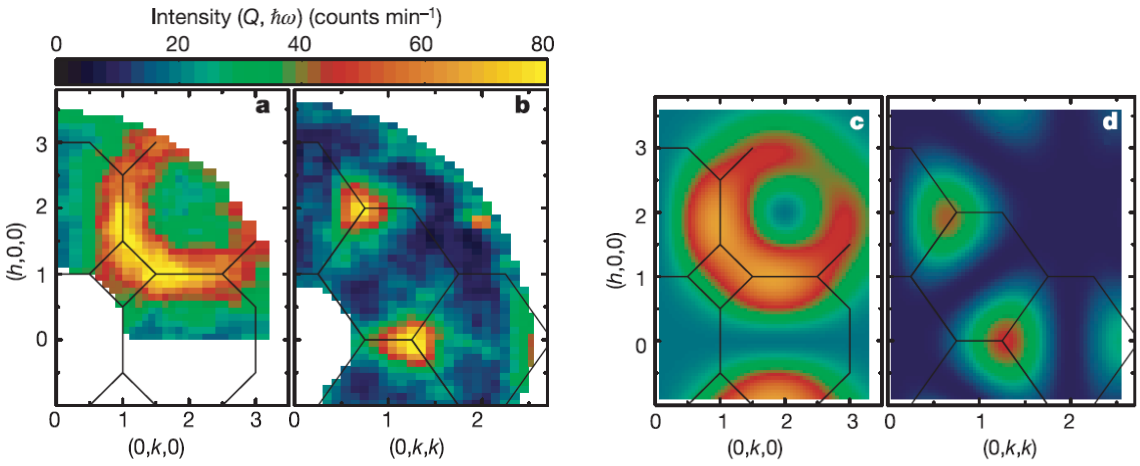


Figure 5.2: Left: inelastic neutron scattering intensity for $ZnCr_2O_4$ at T=15K, $\hbar\omega = 1meV$, from Ref. 52. Right: cubic-averaged form factor for hexagon of antiferromagnetically correlated spins, also from Ref. 52. The contrast with the nearest neighbour model, Fig. 3.4, is striking.

including more exotic compounds for which orbital degrees of freedom are important, see Ref. 51.

The chromites mentioned undergo a structural phase transition to a tetragonal phase with long range magnetic order due to spin-lattice coupling, but the leading interaction above the transition is NN AFM exchange. Lee *et al.* [52] performed inelastic neutron scattering on $ZnCr_2O_4$ and the low energy dynamic structure factor is shown on the left in Fig. 5.2. In Chapter 4 we showed that the classical nearest neighbour model exhibits an approximate factorisation of the dynamic structure factor as $S(\mathbf{q},\omega) = S(\mathbf{q})f(\omega)$, and so our expectation would be that inelastic magnetic scattering should strongly resemble the static structure factor, Fig.3.4. In fact they are quite different - the pinch point features seem entirely absent in Fig. 5.2.

Lee *et al.* [52] explained the scattering in Fig. 5.2 as due to hexagonal clusters of antiferromagnetically aligned spins. They observed that every lattice site can be considered a member of a hexagonal plaquette. By having loops of antiparallel spins on each hexagon, one obtains a ground state. The spin orientation of each hexagonal cluster can fluctuate independently without violating the ground state constraint, and these fluctuations are thus a zero mode. In the flux picture, the hexagonal clusters are equivalent to tiling the lattice with small loops of flux. With this prescription, there are no flux degrees of freedom available at long wavelengths and pinch points are thus completely suppressed. The right panel in Fig. 5.2 shows the calculated form factor for such independently fluctuating clusters of spins, averaged over the orientations

of hexagonal plaquettes found in the pyrochlore lattice. The agreement with the experimental data is taken as strong evidence for just such independent clusters. The ground states of antiferromagnetically correlated hexagonal loops obtained in this way are just the simplest of the class of ground states based on AFM loops proposed by Villain [4]. Similar scattering is observed in other chromites, $ACr_2O_4$, with $A =$ Mg [53] and Cd [54] as well as in the ferrite $CdFe_2O_4$[50], and explained as due to spin clusters. Incidentally, the spin relaxation rate in $CdFe_2O_4$ is approximately linear in temperature at low temperatures [50], and this is shown in the right panel in Fig. 5.1.

The hexagonal cluster phenomenology is successful in so far as Fig. 5.2 goes, but theoretically unsatisfactory: the lattice tiling implies a breaking of lattice symmetry for which no phase transition is seen; and the phenomenological picture gives little insight into the mechanism which stabilizes the hexagonal correlations, and suppresses the pinch point scattering. Unless it is a truly quantum effect within the nearest neighbour model, there must be other interactions which cause the dramatic change in scattering. There has been little attempt to reconcile the experimental measurements in the paramagnetic phase, which are described by the hexagonal cluster phenomenology, with the known facts about the correlations in the (classical) nearest neighbour model. In some cases, the similarity of the scattering to that found in $ZnCr_2O_4$ is taken in itself as evidence for a nearest neighbour Hamiltonian [54].

Further neighbour exchange was proposed as a mechanism for the cluster like scattering by Yavors'kii *et al.* [55], but in the context of correlations in spin ice. There, extensive Monte Carlo simulations on a refined microscopic model which included further neighbour exchange reproduced experimental results on $Dy_2Ti_2O_7$ better than the minimal theoretical model, and better than the spin cluster phenomenology. We propose that weak further neighbour interactions can account for both the suppression of pinch point scattering as well as the cluster-like scattering. In Chapter 6 we study the consequences of further neighbour interactions using the methods of Chapters 3 and 4 in detail to justify this proposition. Before turning to the details, we provide some motivation here.

## 5.3 Further neighbour exchange

Further neighbour exchange has been implicated in other pyrochlore magnets. $Gd_2Sn_2O_7$ is a Heisenberg magnet, but with strong dipolar interactions, due to the large moment on the rare earth Gd. Its ground-state ordering is the same as that predicted by Palmer and Chalker [56]. $Gd_2Ti_2O_7$ is nominally analogous but has a different

low temperature ordering pattern. Wills *et al.* [57] explain the difference as due to the influence of third neighbour exchange interactions across the hexagonal plaquettes of the pyrochlore lattice. Although not directly related, this example emphasises the importance further neighbour terms may play in selecting low temperature states from a manifold of otherwise degenerate groundstates.

For the frustrated spinels which show hexagonal cluster correlations, *ab initio* calculations reveal that third neighbour couplings are not negligible. In a study of $ACr_2X_4$ spinels including zinc and cadmium chromite, Yaresko [58] finds that AFM coupling between third neighbours is important in all the compounds considered. For both $ZnFe_2O_4$ and $CdFe_2O_4$, Cheng [59] also reports that third neighbour interactions are estimated to be much stronger than second-neighbour interactions and even on the same order in magnitude as first-neighbour interactions.

While $CdFe_2O_4$ has dominant nearest neighbour AFM interactions, $ZnFe_2O_4$ is curious in having a positive $\Theta_{\text{C.W.}} = +100K$ indicating ferromagnetic exchange yet despite this Kamazawa *et al.* [60] report that sufficiently pure samples[1] of $ZnFe_2O_4$ remain paramagnetic to the lowest measured temperatures of 1.5 K. This state of affairs is partly explained by a temperature dependent bond angle in the exchange path for first neighbour exchange, which results in a cross over from ferromagnetic to antiferromagnetic $J_1$ as temperature is lowered. In addition to this though, Yamada *et al.* [61] showed that the inclusion of third neighbour couplings is necessary for an adequate theoretical understanding of the observed scattering.

$CdFe_2O_4$ is a candidate for the effects of further neighbour exchange on the nearest neighbour Heisenberg AFM, and has not had very much attention. Single crystal measurements by Kamazawa *et al.* [50] find a peak in scattering at incommensurate wavevectors very close to those subsequently identified by Chern *et al.* [32] in a study of weak second neighbour interactions, although those authors do not refer to the work on $CdFe_2O_4$. There is no structural transition in $CdFe_2O_4$ so it seems plausible that the spin-lattice coupling mechanism responsible for the structural transitions in chromites is not enough to explain the nature of the paramagnetic correlations shared by many frustrated spinels.

Y(Sc)$Mn_2$ is an itinerant electron system, where the Mn ions form a pyrochlore lattice. Magnetic scattering shows broadly similar features to $ZnCr_2O_4$ with suppressed or absent pinch point scattering [62].

[1] from Ref. 61, the Bragg peaks which develop at $T_N = 13$ K grow progressively weaker in intensity and broader in width as sample purity increases, indicating the role of quenched atomic disorder in driving long range order.

### 5.3.1 Interactions in the Coulomb phase

We have established that some interactions beyond nearest neighbour are necessary to explain experimental scattering, and *ab initio* results suggest further neighbour interactions will have relevance to experimental systems. Through the mapping to the Coulomb phase description, we can motivate the study of further neighbour terms to compare to other Coulomb phase systems with other interactions.

The low temperature paramagnet is approximately described by a free field action for solenoidal fields where the stiffness which controls the amplitude of fluctuations is the only relevant parameter at long wavelengths, as described in Chapter 2. One expects additional microscopic interactions to change the value of the stiffness, and the theme of additional interactions causing a flow of the stiffness parameter is common to other Coulomb phase systems.

It is established that 3-D close packed dimer models on bi-partite lattices are in a Coulomb phase with dipolar correlations [18]. One can add interactions between the dimers which favour particular local dimer configurations, such as favouring flippable plaquettes. Alet *et al.* [63] considered the cubic close-packed dimer model with such an interaction and they find that the stiffness diverges as the system undergoes a transition into a dimer crystal. Later work on the same model discusses the correlations in the crystal and Coulomb phases [64]. Two dimensional models have their own complications, but by way of analogy, models with height representations have similarities to three-dimensional Coulomb phases. Alet *et al.* [65] studied the close-packed dimer model with dimer-aligning interactions and showed that the stiffness (Coulomb gas coupling constant) increases towards the transition to the dimer crystal.

# Chapter 6

# The role of further neighbour interactions

In this chapter we address the effect of further neighbour interactions on the classical pyrochlore Heisenberg antiferromagnet. The predictions of the nearest neighbour model discussed in Chapters 3 and 4 are at variance with experimental observations of correlations in the paramagnetic phases of frustrated spinels, and so an understanding must necessarily go beyond a nearest neighbour classical description. One approach is to derive effective classical models which include the effects of quantum fluctuations which lift the classical degeneracy [66]. Our approach is to modify the classical model by including further interactions, motivated by the reasons in the previous chapter. We will find that the effects of further neighbour terms can be quite dramatic, even at temperatures much greater than the further neighbour exchange. Thus even if further neighbour interactions are not *the* reason for the experimentally observed scattering, they should probably be included in the analysis of any microscopic model which is more refined than purely nearest neighbour.

Further neighbour terms, both ferromagnetic (FM) and antiferromagnetic (AFM), will in general have two effects: (i) to change scattering in the paramagnetic phase; (ii) to induce ordering at a low $T$. The nature of such ordering is interesting in its own right: at the mean field level, Reimers *et al.* [22] examined possible ordered states with further neighbour exchange; and following simulations by Tsuneishi *et al.* [67], a detailed study by Chern *et al.* [32] of the phase diagram with weak second neighbour interactions confirmed the mean field predictions at the lowest temperatures, as well as finding a phase with only partial order. Our aim is to develop a full understanding of the effects on paramagnetic correlations, as appropriate for comparing to experiments on frustrated magnets in their paramagnetic phase.

Most of the work reported in this chapter has appeared in Ref. 68.

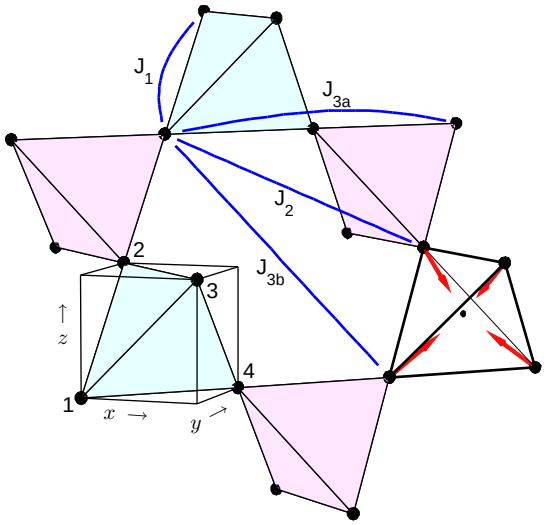


Figure 6.1: Pyrochlore lattice, illustrating the first, second and third neighbour couplings. There are two inequivalent third neighbour interactions, but for $J_{3a} = J_{3b}$, we use $J_3$ to denote both. Numbers indicate sublattice labeling. The centres of the tetrahedra form a diamond lattice and the red arrows show the choice made to orient the diamond lattice bonds.

Our main results are as follows: first, further neighbour exchange terms of the right sign (FM $J_2$ or AFM $J_3$) can be effective both at suppressing pinch point scattering intensity and at broadening pinch points; and second, within the manifold of NN ground states, they favour states with correlations that are similar to those obtained from the phenomenological spin cluster model outlined in Chapter 5. Importantly for the relevance of this work, these effects on paramagnetic scattering are noticeable even at temperatures two orders of magnitude larger than the temperature of the transition that the additional interactions induce.

A visual impression of the findings in relation to experimental reports of spin clusters is given by comparing figures that show the computed diffuse scattering which should be compared with Fig. 5.2. The first case is the nearest neighbour model, for which pinch points, broadened by finite temperature, are apparent in Figs. 6.3 (a) and 6.3 (b). The second case is a model with additional weak further neighbour interactions, for which behaviour is shown in Figs. 6.3 (c) and 6.3 (d). Scattering at the positions of pinch points is heavily suppressed. In its place, rings of scattering intensity are formed of the kind that have been interpreted as indications of spin cluster formation in Refs. 52, 54, 50, 53.

## 6.1 Self consistent Gaussian approximation

In Chapter 3 we detailed the techniques for studying the nearest neighbour magnet within the self-consistent gaussian approximation (SCGA). We extend the calculation to include further neighbour interactions. Ultimately the validity of the SCGA approach is confirmed by the excellent agreement with Monte Carlo (MC) simulations on the Heisenberg model. Fig. 6.1 shows the labels for the various couplings which we study. Unfortunately, there does not seem to be a universally agreed labeling for the two inequivalent third neighbour couplings. When we consider $J_{3a} = J_{3b}$, we refer simply to $J_3$.

As in Chapter 3, we distinguish between the adjacency matrix for the interaction, and an interaction matrix, related to the respective adjaceny matrix by the addition of a multiple of the identity matrix. The choice is made so that the parameter which controls the spin-length, $\lambda$, remains the stiffness of a Coulomb-phase like action. As notation for this chapter, the adjacency and interaction matrices for $m$-th neighbour couplings will be denoted $A^{(m)}$ and $V^{(m)}$ respectively.

Within the SCGA, spin configurations are weighted by $e^{-\beta\mathcal{H}}$ where

$$\beta\mathcal{H} = \frac{1}{2}\sum_{ij}(\lambda\delta_{ij} + \beta\sum_n J_n V^{(n)}_{ij})s_i s_j \,. \tag{6.1}$$

where, as earlier, $\lambda$ is determined by the condition that $\langle s_i^2\rangle = 1/3$, which leads to the self-consistency equation

$$\frac{1}{3} = \frac{1}{4N}\mathrm{Tr}\left[\lambda\mathsf{I} + \beta\sum_n J_n \mathsf{V}^{(n)}\right]^{-1} .$$

The translational symmetry of the lattice and the interactions ensures the matrices are block-diagonal in reciprocal space, and the self-consistency condition can be written in that basis as

$$\frac{1}{3} = \frac{1}{4N}\sum_{\mathbf{q}\in \mathrm{B.Z}}\mathrm{Tr}\left[\lambda\mathsf{I} + \beta\sum_n J_n \mathsf{V}^{(n)}(\mathbf{q})\right]^{-1} \tag{6.2}$$

where the matrices $\mathsf{V}^{(n)}(\mathbf{q})$ act within the space of the four sublattices.

The correlation functions in reciprocal space between different sublattices are

$$\langle s_\mu(\mathbf{q})s^*_\nu(\mathbf{q})\rangle = \left[\lambda\mathsf{I} + \beta\sum_n J_n \mathsf{V}^{(n)}(\mathbf{q})\right]^{-1}_{\mu\nu} \tag{6.3}$$

and the static structure factor is the sum of all 16 such correlators. In Chapter 3, we provided an analytic expression for the structure factor obtained in this way for

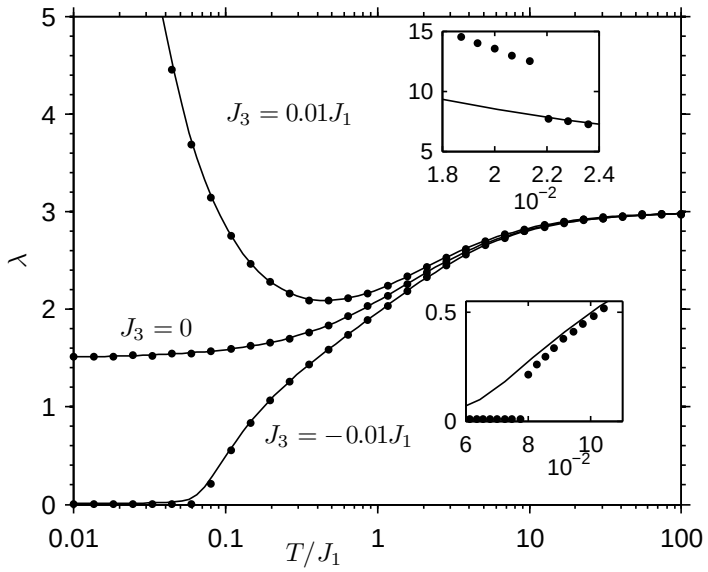


Figure 6.2: Dependence of flux stiffness $\lambda$ on temperature, showing influence of weak further neighbour interactions. Solid lines: solutions of Eq. (6.2) for the cases $J_3 = \pm 0.01 J_1$ and $J_3 = 0$, all with $J_2 = 0$. Points: results from simulations on 2048 spins interpreted via Eq. (6.6). The effects of third neighbour interactions are clearly noticeable even at temperatures two orders of magnitude larger than $J_3$. Insets: detailed behaviour at low temperature; axes as in main panel. First order phase transitions occur in simulations for both signs of $J_3$.

purely nearest neighbour interaction. In the general case where further neighbour terms are non-zero and analytic expressions are unilluminating, spin correlations at any temperature can be straightforwardly obtained numerically from Eq. (6.3). The numerical procedure is the inversion of a $4 \times 4$ matrix, which is vastly quicker than Monte Carlo simulations.

Fig. 6.2 shows the solution of the self-consistency condition for $J_3 = -0.01J_1, 0, 0.01J_1$ with $J_2 = 0$. The magnitude of further neighbour exchange is then the left most point on the logarithmic temperature axis. It is notable that the effect of $J_3$ is visible at temperatures $T \gg J_3$. The interpretation of $\lambda$ as the stiffness of a Coulomb-phase like action allows a direct measurement in simulations by measuring the long wavelength flux variables - the points in Fig. 6.2 come from Monte Carlo measurements interpreted in this way.

## 6.2 Gradient terms in the Coulomb phase

We can pursue the same mapping from microscopic degrees of freedom to long wavelength flux variables as described in Chapter 3, which affords a description of the effect of further neighbour terms which can be compared with other systems with a Coulomb phase, such as those discussed in Ref. 15 and references therein.

The mapping proceeds by the same change of basis in reciprocal space as Eq. (3.1.5), and restriction of the Fourier transformed interactions to $\mathcal{O}(q^2)$. Integrating out the long wavelength fluctuations of the total magnetization induces couplings at $\mathcal{O}(q^4)$ between the flux variables and to order $\mathcal{O}(q^2)$, we can legitimately restrict to the 3x3 subspace of the flux variables. The stiffness $\lambda$ must of course be calculated without any restrictions on wave-vector, using Eq. (6.2).

### 6.2.1 Explicit forms for interaction matrices

To make the procedure more concrete, we include here the $\mathcal{O}(q^2)$ expansions of the Fourier transformed interaction matrices expressed in the basis, given in Eq. (3.1.5), the 'physical basis'. We mentioned above that the interaction matrices are related to the adjacency matrices by the addition of a term proportional to the identity matrix - this diagonal term is chosen such that, when expressed in the physical basis, the $3 \times 3$ subblock which acts among the flux components vanishes at $\mathbf{q} = 0$. This ensures that $\lambda$ can be interpreted as the stiffness of the flux fluctuations. Explicitly, the interaction matrices are $V_{ij}^{(1)} = A_{ij}^{(1)} + 2\delta_{ij}$; $V_{ij}^{(2)} = A_{ij}^{(2)} + 4\delta_{ij}$; $V_{ij}^{(3a)} = A_{ij}^{(3a)} - 6\delta_{ij}$ and $V_{ij}^{(3b)} = A_{ij}^{(3b)} - 6\delta_{ij}$. With $\mathsf{P}$ denoting the basis change, $a$ the nearest neighbour pyrochlore distance, and $q^2 = \sum_i q_i q_i$, the $\mathcal{O}(q^2)$ expansions are

$$\mathsf{P}\mathsf{V}^{(1)}\mathsf{P}^{\mathsf{T}} \simeq \begin{pmatrix} 8 - a^2q^2 & -a^2q_yq_z & -a^2q_xq_z & -a^2q_xq_y \\ -a^2q_yq_z & a^2q_x^2 & a^2q_xq_y & a^2q_xq_z \\ -a^2q_xq_z & a^2q_xq_y & a^2q_y^2 & a^2q_yq_z \\ -a^2q_xq_y & a^2q_xq_z & a^2q_yq_z & a^2q_z^2 \end{pmatrix}$$

and

$$\mathsf{P}\mathsf{V}^{(2)}\mathsf{P}^{\mathsf{T}} \simeq \begin{pmatrix} 16 - 6a^2q^2 & 2a^2q_yq_z & 2a^2q_xq_z & 2a^2q_xq_y \\ 2a^2q_yq_z & -2a^2q_x^2 & -2a^2q_xq_y & -2a^2q_xq_z \\ 2a^2q_xq_z & -2a^2q_xq_y & -2a^2q_y^2 & -2a^2q_yq_z \\ 2a^2q_xq_y & -2a^2q_xq_z & -2a^2q_yq_z & -2a^2q_z^2 \end{pmatrix} + 4a^2 \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & q_y^2 + q_z^2 & 0 & 0 \\ 0 & 0 & q_x^2 + q_z^2 & 0 \\ 0 & 0 & 0 & q_x^2 + q_y^2 \end{pmatrix}$$

and

$$\mathsf{P}\mathsf{V}^{(3a)}\mathsf{P}^{\mathsf{T}} \simeq$$
$$4a^2\begin{pmatrix} -q^2 & -q_yq_z & -q_xq_z & -q_xq_y \\ -q_yq_z & -q_x^2 & -q_xq_y & -q_xq_z \\ -q_xq_z & -q_xq_y & -q_y^2 & -q_yq_z \\ -q_xq_y & -q_xq_z & -q_yq_z & -q_z^2 \end{pmatrix} - 4a^2\begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & q_y^2+q_z^2 & 0 & 0 \\ 0 & 0 & q_x^2+q_z^2 & 0 \\ 0 & 0 & 0 & q_x^2+q_y^2 \end{pmatrix}$$

$$\mathsf{P}\mathsf{V}^{(3b)}\mathsf{P}^{\mathsf{T}} \simeq$$
$$4a^2\begin{pmatrix} -q^2 & q_yq_z & q_xq_z & q_xq_y \\ q_yq_z & q_x^2 & q_xq_y & q_xq_z \\ q_xq_z & q_xq_y & q_y^2 & q_yq_z \\ q_xq_y & q_xq_z & q_yq_z & q_z^2 \end{pmatrix} - 4a^2\begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 2q_x^2+q_y^2+q_z^2 & 0 & 0 \\ 0 & 0 & q_x^2+2q_y^2+q_z^2 & 0 \\ 0 & 0 & 0 & q_x^2+q_y^2+2q_z^2 \end{pmatrix}$$

In each case we have split the interaction into a part which acts like $J_1$ within the flux field subspace, and another part. Chern *et al.* [32] argue that $J_2$ and $-J_{3a}$ have the same effect when $J_1 \to \infty$. This fact is evident in the small $q$ expansions above. If the two symmetry inequivalent third neighbour couplings are taken to have the same strength, then third neighbour couplings are just fcc nearest neighbour couplings for each sublattice. In this case,

$$\mathsf{P}[\mathsf{V}^{(3a)} + \mathsf{V}^{(3b)}]\mathsf{P}^{\mathsf{T}} \simeq -8a^2\mathrm{diag}(q^2)\,. \tag{6.4}$$

### 6.2.2 Long wavelength theory for further neighbour interactions

By expanding the interactions to second order in $\mathbf{q}$ we obtain as a long wavelength theory for the fourier transformed flux components

$$\begin{aligned}\beta\mathcal{H} = \frac{1}{2}\sum_{\mathbf{q}} \lambda|\mathbf{B}|^2 + \beta[J_1 - 2J_2 - 4J_{3a} + 4J_{3b}]a^2|\mathbf{q}\cdot\mathbf{B}|^2 \\ + 4\beta[J_2 - J_{3a} - J_{3b}]a^2((q_y^2+q_z^2)|B_x|^2 + (q_x^2+q_z^2)|B_y|^2 + (q_x^2+q_y^2)|B_z|^2) \\ - 8\beta J_{3b}a^2(q_x^2|B_x|^2 + q_y^2|B_y|^2 + q_z^2|B_z|^2)\,. \end{aligned} \tag{6.5}$$

It is clear from Eq. (6.5) that $J_1$ acts purely to enforce the constraint $\mathbf{q}\cdot\mathbf{B} = 0$, a longitudinal constraint on the flux fields. The various extra terms in Eq. (6.5) due to further neighbour interactions couple to transverse gradients in the flux. Depending on the sign of the further neighbour couplings, transverse gradients may either be suppressed or promoted. Suppression of transverse gradients leads to ordering at

zero wavevector and is naturally described in a long wavelength theory. When transverse gradients are promoted, one can object that the long wavelength description is unstable - this is not of much concern since Eq. (6.5) is really a conceptual short hand for understanding the behaviour of Eq. (6.1).

### 6.2.3 Width and intensity of pinch points

We now consider the flux correlations that follow from Eq. (6.5). At $\mathbf{q} = 0$ the flux fluctuations are straightforwardly obtained as

$$\langle B_i(\mathbf{q}) B_j^*(\mathbf{q}) \rangle_{\mathbf{q}=\mathbf{0}} = \frac{1}{\lambda} \delta_{ij} \tag{6.6}$$

which confirms our identification of the parameter $\lambda$ as the stiffness of the flux fields in the long wavelength description, when the further neighbour interactions are defined in the appropriate way. The correlation function in Eq. (6.6) is directly accessible from simulations and the data points in Fig. 6.2 follow from this relation; the same relation is used to measure the stiffness in simulations of dimer models (see e.g. Eq. (2) in Ref. 63). It is impressive that the SCGA captures behaviour in the Heisenberg model so accurately, even with further neighbour interactions. The pinch point intensity in scattering is the inverse of the stiffness $\lambda$: the flow of the stiffness with temperature, seen in Fig. 6.2, will cause pinch points to be suppressed or intensified relative to the nearest neighbour model.

The profile of the pinch points is given by the correlator for longitudinal flux fluctuations:

$$\langle B_x(q_x) B_x^*(q_x) \rangle_{q_y=q_z=0} = \frac{1}{\lambda} \left[ \frac{\xi_\parallel^{-2}}{\xi_\parallel^{-2} + q_x^2} \right] \text{ where } \xi_\parallel = a \sqrt{\beta (J_1 - 2J_2 - 4J_{3a} - 4J_{3b})/\lambda} \tag{6.7}$$

The length $\xi_\parallel$ is the longitudinal correlation length and we emphasise that $\xi_\parallel$ depends on $\beta$ both explicitly and through the $\beta$-dependence of $\lambda$, via Eq. (6.2). A diverging stiffness thus leads to suppression (via the prefactor) and broadening (via $\xi_\parallel$) of the pinch points relative to the nearest neighbour model.

The transverse flux correlations are

$$\langle B_z(q_x) B_z^*(q_x) \rangle_{q_y=q_z=0} = \frac{1}{\lambda} \left[ \frac{1}{1 + \frac{4\beta(J_2 - J_{3a} - J_{3b})a^2}{\lambda} q_x^2} \right] .$$

A transverse length scale is $\xi_\perp = a\sqrt{4\beta |J_2 - J_{3a} - J_{3b}|/\lambda}$. In the case that $(J_2 - J3a - J_{3b}) > 0$, $\xi_\perp$ is the correlation length for transverse flux fluctuations but otherwise $\xi_\perp$

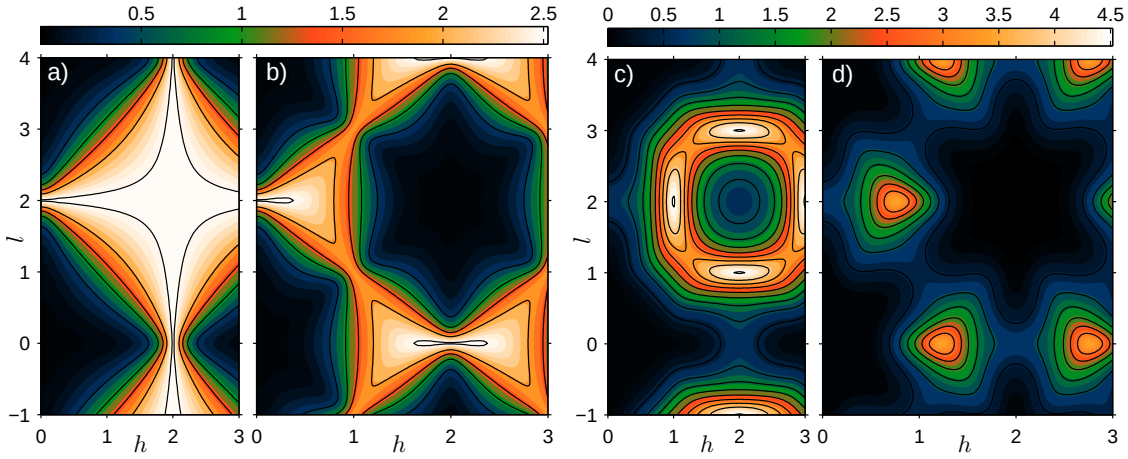


Figure 6.3: Structure factor $S(\mathbf{q})$ at $T = 0.1J_1$. (a) and (b): NN interactions only; (c) and (d): $J_3 = 0.025J_1$. (a) and (c): $(h0l)$ plane; (b) and (d): $(hhl)$ plane. Contour lines are at the levels marked on the color scale. Compared to a nearest neighbour model, the effect of AFM $J_3$ is, first, to suppress and broaden the pinch points, and second, to shift spectral weight so that behaviour resembles more the hexagonal cluster scattering. MC simulations yield plots that are indistinguishable by eye.

is not a correlation length, instead it sets a length scale at which higher order terms in $q$ must be included in the long wavelength theory. Thus qualitatively different behaviour is expected depending on the sign of $J_2 - J_{3a} - J_{3b}$.

In their consideration of polarization fluctuations in ferroelectrics, Youngblood and Axe [69] considered a phenomenological Landau free energy essentially identical to Eq. (6.5). Our discussion differs in that Eq. (6.5) is derived directly from a microscopic model, the parameter $\lambda$ is allowed to flow, and moreover both positive and negative coupling to transverse gradient terms are considered.

Having established that $\lambda$ plays an important role in understanding the nature of the pinch point correlations, we now turn to putting bounds on $\lambda$ by analysing the self-consistency equation, Eq. (6.2). Let $\varepsilon_{\min}$ be the minimum eigenvalue of $\sum_n J_n \mathsf{V}^{(n)}(\mathbf{q})$ over all $\mathbf{q}$. The requirement that the quadratic form in Eq. (6.1) is positive-definite demands that $\lambda + \beta\varepsilon_{\min} > 0$ with equality achieved only in the thermodynamic limit at a phase transition. It follows that

$$\lambda > -\beta\varepsilon_{\min} \ .$$

To illustrate the result, we consider two special cases with AFM $J_1$. First for an AFM $J_3$ and all other $J_{n>1} = 0$, we have $\varepsilon_{\min} = -16J_3$. Second, for a FM $J_2$ and all other $J_{n>1} = 0$, the position of the minimum in reciprocal space is given by Chern *et al.* [32]

in their Eq. (4), with $\varepsilon_{\min} \simeq -8.57|J_2|$ for $|J_2| \ll J_1$. The relatively large numerical prefactors are an indication that the effects of further neighbour terms are stronger than one might first imagine.

A second order phase transition is predicted by the SCGA at the temperature such that $\lambda = -\beta\varepsilon_{\min}$ solves the self-consistency equation, in the thermodynamic limit. Rearranging for the temperature, this occurs at

$$T_c = \frac{S^2}{3}\left[\int \frac{d\mathbf{q}}{v_{\mathrm{B.Z}}}\frac{1}{4}\mathrm{Tr}\left[\sum_n J_n \mathsf{V}^{(n)}(\mathbf{q}) - \varepsilon_{\min}\mathsf{I}\right]^{-1}\right]^{-1} . \tag{6.8}$$

where $S$ is the length of the classical vector spins and $v_{\mathrm{B.Z}}$ is the volume of the Brillouin Zone. While in practice, we find that these transitions are preempted by first order transitions in simulations, it is nevertheless interesting to determine the predicted transition temperature.

We now apply the theory developed above in more detail to particular choices of further neighbour exchange.

## 6.3 Consequences of AFM $J_3$ or FM $J_2$

### 6.3.1 Paramagnetic phase

Within the paramagnetic phase, AFM $J_3 > 0$ or FM $J_2 < 0$ act to suppress pinch points, and to promote cluster-like correlations. These effects are shown in Fig. 6.3, which should be compared with Fig. 5.2. The behaviour illustrated here is one of our main results. Its origin is discussed in this section.

#### Consequences for pinch point visibility

The pinch points are the signature of emerging algebraic correlations in a Coulomb phase. There are three ways in which they may become less visible. (i) Further neighbour interactions or finite temperature endow the pinch points with a width set by the inverse of the longitudinal correlation length. (ii) The absolute amplitude of the pinch points is set by the inverse of the flux field stiffness. (iii) The visibility of the pinch points depends also on the relative weight at the pinch points compared to that at the maximum in scattering. We address these mechanisms quantitatively below.

(i) For AFM $J_3$, the longitudinal correlation length is $a\sqrt{\beta(J_1 - 8J_3)/\lambda}$. The lower bound we derived on $\lambda$ puts an upper bound on the correlation length of

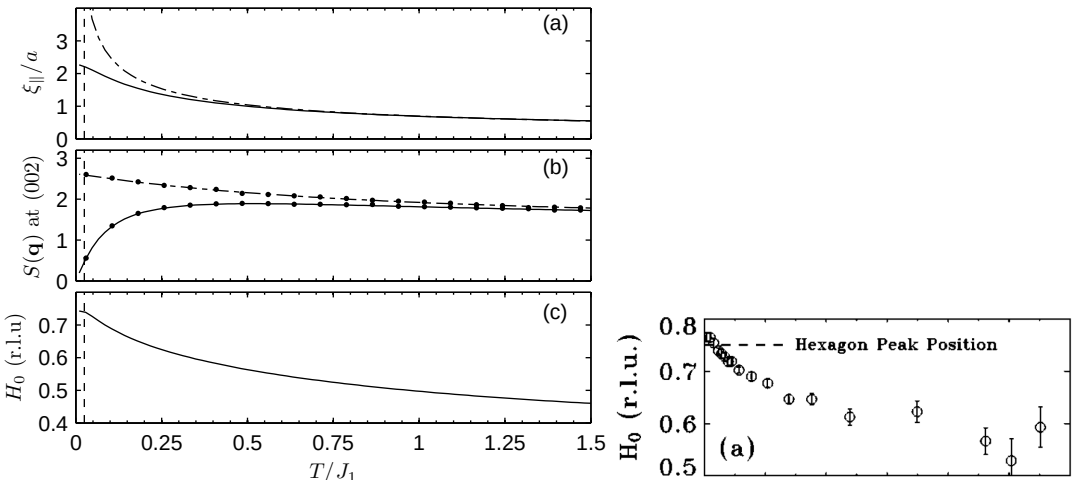


Figure 6.4: Left: Temperature dependence of quantities characterising the scattering, for $J_2 = -0.025J_1$. (a) Coulomb phase correlation length $\xi$ in units of site spacing for NN model (broken line) and for model with non-zero $J_2$ (solid line). (b) Amplitude of correlation function at pinch-point, for both models. (c) Position $H_0$ of maximum in scattering, as described in text. Dashed line indicates temperature $T = J_2$: the effect of non-zero $J_2$ extends to $T \gg J_2$. Curves are from SCGA. Points are from simulations on 2048 spins. Right: measurements on $CdFe_2O_4$ of the position of the peak in scattering, from Ref. 50.

$a\sqrt{(J_1/16J_3 - 1/2)}$. Even with $J_3 \simeq J_1/150$, the correlation length never exceeds 3 site spacings and the pinch points are thus never sharp. For FM $J_2$, a similar argument gives an upper bound on the correlation length of approximately $a\sqrt{(2 + J_1/|J_2|)/8.6}$. The effect on the correlation length is shown in Fig. 6.4(a) for $J_2 = -0.025J_1$.

(ii) Since the absolute amplitude of the pinch points is set by $1/\lambda$, the lower bound on $\lambda$ also puts an upper bound on the absolute intensity of pinch point correlations. For AFM $J_3$, the bound was that $\lambda > 16\beta J_3$, while for FM $J_2$, we have $\lambda > 8.6\beta J_2$, for $|J_2| \ll J_1$. The pinch point intensity is shown in Fig. 6.4(b) for $J_2 = -0.025J_1$, and compared with a nearest neighbour model.

(iii) Moreover, as spectral weight shifts away from the pinch points to other wavevectors, the relative amplitude given by the ratio of pinch point correlations to the maximum in the correlation function decreases even more rapidly.

These three mechanisms indicate how pinch points may both be suppressed and broadened by further neighbour interactions. The dramatic broadening of the pinch points would be missed in any approximation which projects onto classical ground-states.

**Hexagonal cluster-like correlations**

We have seen how further neighbour terms suppress spectral weight at the pinch points. We turn now to why antiferromagnetically correlated loops of spins such as hexagonal clusters may appear in a model with further neighbour terms. The mechanism is straightforward to explain in the long wavelength description. The further neighbour terms couple to transverse gradients of the flux fields and the effect of AFM $J_3$ or FM $J_2$ is to promote transverse gradients in the flux field. In the continuum picture, maximising transverse gradients requires having as many small loops of flux as possible; when understood in terms of the lattice geometry, the smallest loops of flux available are the hexagonal loops on the diamond lattice, which yield hexagonal clusters of antiferromagnetically correlated spins. Considered microscopically, it is also clear from Fig. 6.1 that an AFM $J_{3b}$ is likely to induce antiferromagnetically correlated hexagons of spins.

Following Ref. 50, we define $H_0$ as the temperature-dependent position of the maximum in scattering along the (hh2) direction (see vertical dashed lines in Fig. 6.5). For FM $J_2$, at $T = 0$, $H_0$ coincides with $h^*$ given by Chern *et al.* [32] in their Eq. (4), the position in reciprocal space of the minimum eigenvalue of the exchange matrix. The evolution of $H_0$ with temperature extracted from the SCGA is shown in Fig. 6.4(c). Qualitatively similar behaviour for $H_0$ is found experimentally in $CdFe_2O_4$ which we show alongside [50]. In contrast, in the NN model (at least classically) the scattering maxima are always found at the pinch points, i.e $H_0 = 0$, so the NN model does not provide an explanation for why the maximum in scattering is elsewhere.

In Fig. 6.6, we show the structure factor on the same planes of reciprocal space as in Fig. 6.3, but for the case of a FM $J_2$ rather than an AFM $J_3$. The qualitative effect is similar, but the characteristic scattering shapes look less like the hexagonal form factor used to explain the scattering in many chromites. We summarise the characteristic shapes of such scattering in Sec. 6.5.

### 6.3.2 Ordering transitions

For $J_3 \ll J_1$, we observe a first order phase transition at $T_c \simeq 2.2J_3$ in our MC simulations (eg. see upper inset in Fig. 6.2). We have not attempted to characterize the ordered state since our focus is on the strongly correlated paramagnet in the window $T_c < T < J_1$. A phase transition is not predicted by mean field calculations[22], nor within the theory of Sec. 6.1 since precisely for $J_{3a} = J_{3b}$ the interaction minimum

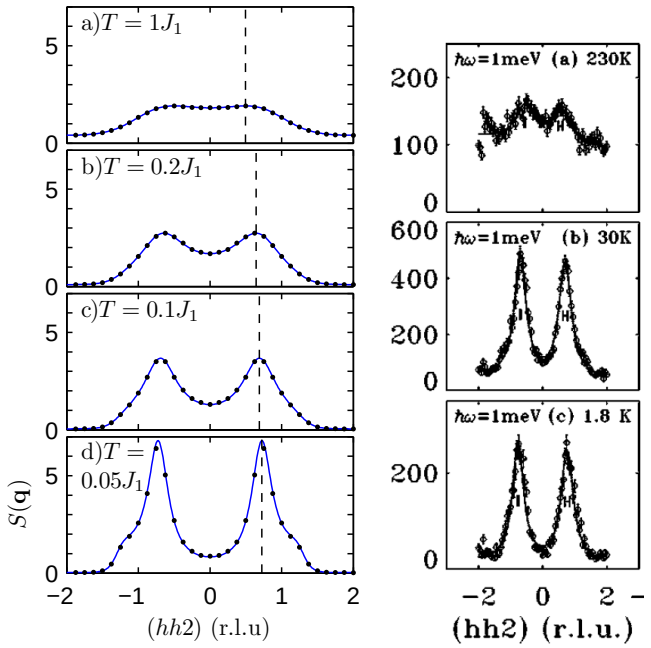


Figure 6.5: Left:Structure factor along $(hh2)$ for $J_2 = -0.025J_1$. Lines are from SCGA, points from simulations on 16384 spins. The position $H_0$ at each temperature is indicated by dashed line. Right: measurements on $CdFe_2O_4$ from Ref. 50

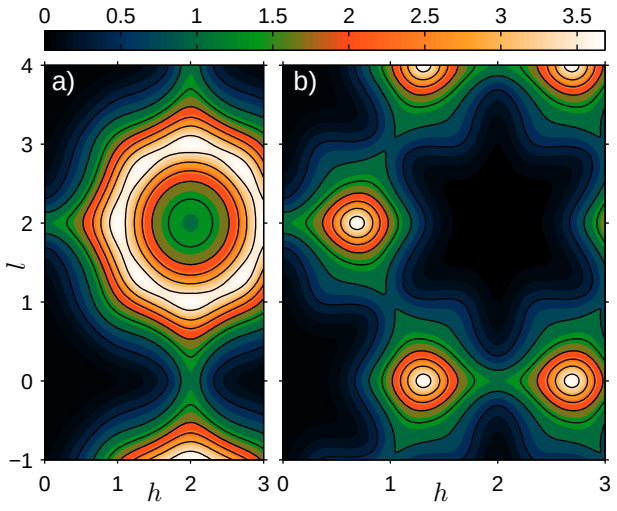


Figure 6.6: Structure factor at $T = 0.1J_1$ with $J_2 = -0.025J_1$. Panels as in Fig. 6.3.

occurs along a 1-dimensional minimum in reciprocal space. Understanding the phase transition with AFM $J_3$ is a topic for further work.

The phase diagram for weak FM $J_2$ has been understood by Chern *et al.* [32]. Ferromagnetic $J_2$ leads to incommensurate order at low temperatures consistent with mean field calculations [22], but also to collinear regime at intermediate temperatures, stabilised by thermal fluctuations, with first order transitions out of the paramagnetic phase. The transition temperature is approximately at $T_c \simeq J_2$, but see Ref. 32 for details.

## 6.4 Consequences of AFM $J_2$ or FM $J_3$

In understanding how further neighbour terms lead to an absence of pinch points, they needed to be of the right sign. We now also consider further neighbour terms with the opposite sign which enhance pinch point scattering.

### 6.4.1 Paramagnetic phase

It is clear from Eq. (6.5) that in the long wavelength flux description an antiferromagnetic $J_2$ or FM $J_3$ penalizes transverse gradients of the flux fields. Qualitatively then, one expects spectral weight to accumulate at small $q$ as short wavelength fluctuations are suppressed, developing into Bragg peaks as the transverse flux correlation length diverges at a phase transition.

As seen from the numerical solution of Eq. (6.2) shown in Fig. 6.2, FM $J_3$ causes the stiffness $\lambda$ to decrease with cooling, and the same behaviour is found with AFM $J_2$. The strength of flux fluctuations is given by the inverse stiffness as in Eq. (6.6) and consequently this increases. In Fig. 6.7 we show spin correlations in the paramagnetic phase at the same value of $\beta J_1$ as in Fig. 6.6: with antiferromagnetic $J_2$, pinch points are intensified and sharper than for the nearest neighbour model, as expected from our discussion. Youngblood and Axe [69] considered the Maxwell action with terms penalising transverse flux gradients and obtained a similar effect.

### 6.4.2 Ordering transitions

With an AFM $J_2 \ll J_1$, the classical Heisenberg model has a first order transition to a $\mathbf{q} = 0$ Néel-ordered phase at $T \simeq 3.2J_2$ [32]. The transition is second order within the SCGA and a numerical calculation of 6.8 gives $T_c \simeq 2.44J_2$ for $J_1/J_2 = 100$. For FM $J_3 \ll J_1$, our simulations also reveal a first order transition at $T \simeq 8|J_3|$

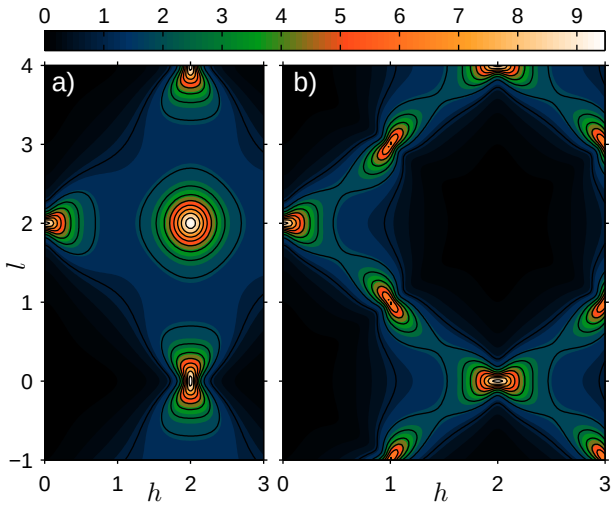


Figure 6.7: Structure factor $S(\mathbf{q})$ at $T = 0.1J_1$ and $J_2 = 0.025J_1$ in (a) the $(h0l)$ plane and (b) the $(hhl)$ plane. Contour lines are at the levels given in the color scale. Positive $J_2$ increases pinch point scattering, and eventually leads to Néel order. The shapes around the pinch points are almost identical to those discussed by Youngblood and Axe [69].

(see lower inset of Fig. 6.2). In both cases, the transition is to a fully flux-polarised state. The ferromagnetic third neighbour interactions acts simply to induce ferromagnetic ordering on all sublattices, so that energetic groundstates are given by the groundstates of a single tetrahedron. The order-by-disorder mechanism then selects a collinear configuration [32]. Were it continuous as in the SCGA, the transition would be associated with a divergence of the transverse flux correlation length, $\xi_\perp$.

Similar transitions occur in other Coulomb-phase systems: for interacting dimer models, this is a transition to a uniform tilted phase, while in the context of polarization fluctuations in paraelectrics, a ferroelectric transition out of the paraelectric phase [69].

## 6.5 General combinations of further neighbour interaction

In Secs. 6.3 and 6.4 we took both third neighbour couplings as equal, but in general they need not be and here we consider $J_{3a}$ and $J_{3b}$ as distinct. Within the long wavelength description, nothing qualitatively changes since further neighbour terms can only act to add gradient terms while renormalising the stiffness. However, the detailed nature of the short range paramagnetic correlations is not generic and depends

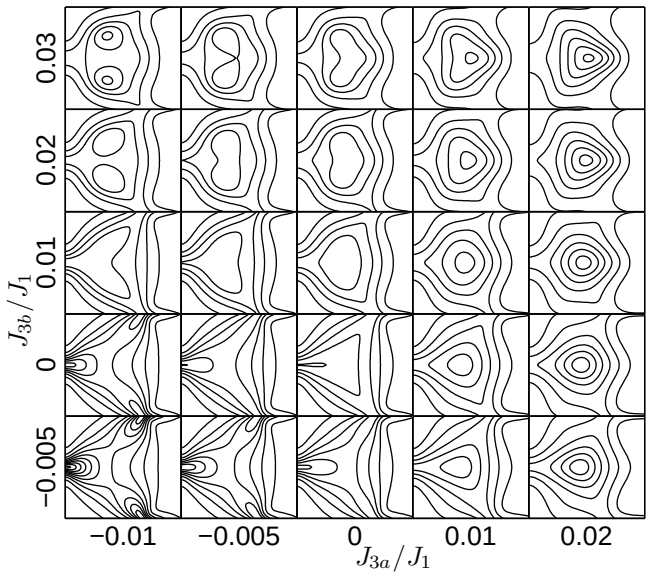


Figure 6.8: $S(\mathbf{q})$ in the $(hhl)$ plane for $h \in [0, 3/2]$, $l \in [1, 3]$ at $T = 0.1J_1$, $J_2 = 0$ and with $J_{3a}$, $J_{3b}$ interactions treated independently, showing the variety of short range correlations, and the effects on the pinch point at (002).

sensitively on the ratios of the different further neighbour couplings. Since this balancing act is separate for each candidate material which our ideas hope to describe, we provide a survey for several different combinations. In Fig. 6.8, we have plotted paramagnetic correlations in the reciprocal space patch described in the caption for a variety of different $J_{3a}$ and $J_{3b}$ at $T = 0.1J_1$. By the arguments of Ref. 32, $J_2$ acts as $-J_{3a}$ for large enough $J_1$, so the horizontal axis can be considered approximately to probe $J_{3a} - J_2$.

## 6.6 Dynamics and inelastic scattering

We have so far considered only equal time correlations but further neighbour terms will also affect dynamics. In Chapter 4, we showed how within the paramagnetic phase of the nearest neighbour model, the intrinsic precessional dynamics associated with the Heisenberg spins is dominated by relaxational modes which are very well captured by a purely relaxational stochastic model based on the SCGA. We find this statement about the low temperature paramagnet remains true even with further neighbour terms. As evidence for this, we present in Fig. 6.9 the results of molecular dynamics simulations using the same techniques as the simulations for Chapter 4, (described in Appendix. A). We plot the dynamic structure factor at 3 different

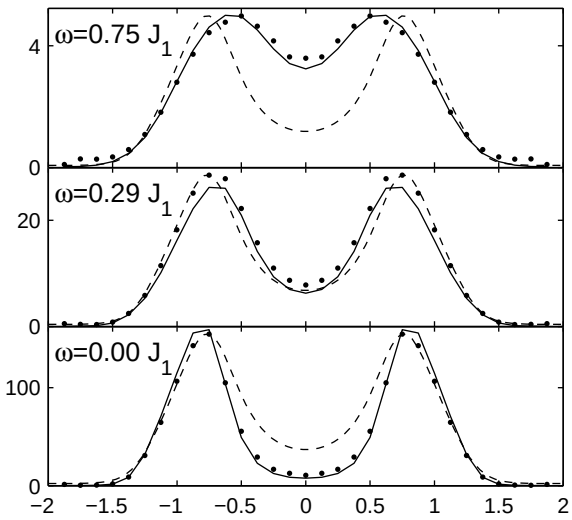


Figure 6.9: Points: $S(\mathbf{q},\omega)$ along $(hh2)$ at $T = 0.1J_1$, $J_3 = 0.025J_1$ from simulations. Solid line: $S(\mathbf{q},\omega)$ from stochastic model of Chapter 4 with further neighbour terms included. Broken line: $S(\mathbf{q})$ rescaled for comparison. Pinch points are more suppressed at small-$\omega$, and conversely less suppressed at high-$\omega$.

frequencies $\omega$. On the same plot, we show the prediction obtained from including further neighbour terms in the stochastic model, simplified by not conserving total magnetisation.

Compared to the static result, the pinch point scattering at (002) is more suppressed at small $\omega$, and conversely, less suppressed at higher $\omega$. The implications for quasi-elastic scattering experiments at $\omega \simeq 0$ are that we expect the effects of further neighbour terms to be more conspicuous (and conversely at larger $\omega$, to become less conspicuous) when compared to the static ($\omega$-integrated) scattering which we have described so far.

The enhanced suppression can be understood by reconsidering the stochastic model. Within that model the dynamic structure factor is understood as the sum of contributions from the four bands of the interaction matrix, given by

$$S(\mathbf{q},\omega) = \sum_{\mu=1}^{4} g_\mu(\mathbf{q}) f_\mu(\mathbf{q},\omega)\,, \tag{6.9}$$

where the $g_\mu(\mathbf{q})$ are form factors, independent of temperature, and where the thermal weight and dynamics for each band are contained in the functions $f_\mu(\mathbf{q},\omega)$. The quadratic form in Eq. (6.1) has eigenvalues $\alpha_\mu(\mathbf{q})$, labeled by band index $\mu$, which determine relaxation rates $\Gamma_\mu(\mathbf{q}) \propto \alpha_\mu(\mathbf{q})$ and these enter the dynamic structure

factor through the relation

$$f_\mu(\mathbf{q},\omega) = \alpha_\mu^{-1}(\mathbf{q})\frac{2\Gamma_\mu(\mathbf{q})}{\omega^2+\Gamma_\mu(\mathbf{q})^2}\,. \tag{6.10}$$

By integrating over $\omega$ we obtain the static structure factor, which has the form of Eq. (6.9), but with $\alpha_\mu^{-1}(\mathbf{q})$ in place of $f_\mu(\mathbf{q},\omega)$.

With only nearest neighbour interactions, the two lowest bands (which are dominant at low temperature) are dispersionless and degenerate, which has two consequences: the $\mathbf{q}$-dependence arises entirely from the form factors in Eq. (6.9) since the relevant $f_\mu$ are independent of $\mathbf{q}$, and the $\omega$-dependence approximately factorises out. With further neighbour interactions included, the lowest bands acquire dispersion and the relevant $f_\mu$ are no longer independent of $\mathbf{q}$. If their ratio in the static case for band $\mu$ at $\mathbf{q}_1$ and $\mathbf{q}_2$ is $r$, then in the dynamic case with $\omega = 0$ it is $r^2$. This is why the suppression effects we have discussed, which correspond to $r \neq 1$, are enhanced in quasi-elastic scattering.

## 6.7 Discussion

The appearance of antiferromagnetically correlated hexagonal clusters in the spinels is certainly puzzling in the context of a pure nearest neighbour model, since properties of the nearest neighbour model have been reliably calculated and give different behaviour. In Chapter 5, we discussed evidence from *ab initio* calculations that further neighbour exchange is important in many of the compounds which show antiferromagnetically correlated clusters without pinch points. Although further neighbour exchange, understood in a classical model, is one route to finding such correlations, we should also address other possibilities.

It has been suggested that spin clusters are induced by magnetoelastic coupling [53]. However, we note that among the compounds where spin clusters have been reported, $CdCr_2O_4$ undergoes a $c$-axis elongation [54], whereas both $ZnCr_2O_4$ and $MgCr_2O_4$ undergo a $c$-axis contraction, and $CdFe_2O_4$ has no structural transition at all; it is not clear how magnetoelastic coupling which leads to different structural transitions, or none, should explain the cluster-like scattering common to all of these materials. Moreover, scattering in Y(Sc)$Mn_2$ is broadly similar to the spinels [62], but due to the different lattice structure we expect the nature of spin lattice coupling to be different, and this throws further doubt over the idea that magnetoelastic coupling may be universally responsible for cluster-like scattering.

We should also discuss the temperature dependence of our results. The images in Figs. 6.3 and 6.6 are for a single temperature, where the hexagonal cluster-like scattering is most visible; at lower temperatures however, the further neighbour terms induce magnetic phase transitions as discussed earlier in this chapter. If indeed further neighbour terms are required to explain the hexagonal cluster scattering, one also expects to observe the associated magnetic phase transitions. In the chromites, we can avoid the issue by assuming the magnetic phase transition is preempted by the structural transitions driven by the independent mechanism of spin-lattice coupling; this is not entirely satisfactory as an explanation, since $CdFe_2O_4$ shows neither a structural transition, nor development of any Bragg peaks at any temperature. It is possible that quantum fluctuations or disorder can stabilise the paramagnetic phase in the way discussed by Saunders and Chalker [70], but it is hard to be conclusive about a range of materials at once, and the role of further neighbour exchange should be assessed individually for different materials.

We have presented further neighbour exchange as the simplest extra interaction which gives hexagonal cluster scattering, but further neighbour exchange is not the only perturbation one could imagine. Others include Dzyaloshinsky-Moriya interactions [71, 72] and biquadratic exchange, which can appear microscopically or as an effective term taking into account quantum fluctuations [66]. In future work it would be interesting to consider the effects of these other perturbations on the paramagnetic phase. The present results are sufficient to show that further neighbour exchange of a realistic strength can dramatically alter the diffuse scattering and should not be ignored in comparing experimental results to simple theoretical models, even if a full description requires a more complete understanding of the interplay between different perturbations away from the classical nearest neighbour limit.

Clusters should be understood, then, as short range order in a strongly correlated state, a similar view to that expressed by Yavors'kii *et al.* [55] when considering cluster-like scattering in spin ices; their analysis including further neighbour terms suggests that cluster-like scattering is the property of a strongly correlated liquid state which is sensitive to weak perturbations rather than due to the emergence of 'real' clusters.

Our results potentially explain why a clear signature of algebraic correlations is missing in candidates for pyrochlore Heisenberg magnets. Small further neighbour interactions are sufficient to wash out the pinch points. Our results may have implications for other pyrochlore models, but in Ising models where the ice-rules are exponentially enforced, we do not expect the dramatic broadening of pinch points

that we observe in the Heisenberg model. Nevertheless, the suppression of pinch point scattering amplitude is expected, and we note that in Ref. 47 such a suppression is reported in the spin ice $Ho_2Ti_2O_7$.

In summary, we have studied the effect of further neighbour interactions on the low temperature paramagnetic phase of the frustrated Heisenberg antiferromagnet on the pyrochlore lattice. Further neighbour terms induce transitions to ordered phases, but they also have a striking effect on paramagnetic correlations. In the description of the low temperature paramagnet as a Coulomb phase, further neighbour terms cause the Coulomb phase coupling constant, or stiffness, to flow. Pinch points in diffuse scattering have their amplitude and width controlled by the stiffness, and even very weak further neighbour terms can cause pinch points to be suppressed in amplitude and broadened. With FM $J_2$ or AFM $J_3$, paramagnetic scattering is altered to resemble more the hexagonal cluster scattering that is often observed in experimental systems, notably spinels. Further neighbour terms then provide a mechanism for the previously unexplained cluster-like scattering in frustrated spinels.

# Chapter Appendix: adjacency matrices

For completeness, we provide here the Fourier transformed adjacency matrices for the further neighbour interactions we considered. The matrix elements appear for example in Ref. 22, although in a different basis, so we provide them here using our conventions. We first define the fcc displacement vectors $\mathbf{t}_{\mu\nu} = 2(\mathbf{c}_\mu - \mathbf{c}_\nu)$ where the $\mathbf{c}_\mu$ are as in Sec. 3.1.1. Then, the matrix elements of the adjacency matrices are given for nearest neighbours:

$$A^{(1)}(\mathbf{q})_{12} = 2\cos\left(\frac{\mathbf{q}\cdot\mathbf{t}_{12}}{2}\right)$$

and for second neighbours,

$$A^{(2)}(\mathbf{q})_{12} = 2\cos\left(\frac{\mathbf{q}\cdot(\mathbf{t}_{13}+\mathbf{t}_{23})}{2}\right) + 2\cos\left(\frac{\mathbf{q}\cdot(\mathbf{t}_{14}+\mathbf{t}_{24})}{2}\right) .$$

In both cases, the diagonal elements are zero and the other matrix elements are obtained by permutation of the sublattice labels in the appropriate way. Since third neighbour couplings link sites on the same sublattice, they are diagonal:

$$A^{(3a)}(\mathbf{q})_{11} = 2\cos(\mathbf{q}\cdot\mathbf{t}_{12}) + 2\cos(\mathbf{q}\cdot\mathbf{t}_{13}) + 2\cos(\mathbf{q}\cdot\mathbf{t}_{14})$$

and

$$A^{(3b)}(\mathbf{q})_{11} = 2\cos(\mathbf{q}\cdot\mathbf{t}_{23}) + 2\cos(\mathbf{q}\cdot\mathbf{t}_{24}) + 2\cos(\mathbf{q}\cdot\mathbf{t}_{34}) .$$

# Chapter 7

# Dynamics of ordered phases

The classical nearest neighbour pyrochlore antiferromagnet remains disordered at low temperature due to macroscopic groundstate degeneracy. The degeneracy can be lifted by weak perturbations. In this situation, the Hamiltonian has two different energy scales, that of the nearest neighbour coupling, and the much smaller scale of the perturbations which induce ordering. In this chapter, we examine the dynamics of fluctuations around an ordered groundstate induced by such interactions.

There are numerous motivations for wanting to understand what happens. In an ordered state which breaks spin rotation symmetry, we are guaranteed to find Goldstone modes. We wish to understand the interplay between the different energy scales in determining the dispersion of the Goldstone modes. Further, it is interesting to examine exactly how the zero modes show up in the linearised dynamics of the nearest neighbour model, as well as their fate in the presence of perturbations. Additionally by understanding the variables which behave in a canonically conjugate way in the classical dynamics, hopefully some insight can be obtained into quantum versions of the same problem through the mapping of Poisson-brackets to commutators.

The ordered state we consider in this chapter is a collinear state with all tetrahedra the same, shown in Fig. 7.1. For a given ordered state, there are typically many microscopic interactions which can stabilise such an ordered state in the classical model. In the semiclassical limit, the selection of collinear states by quantum fluctuations can be included in an effective Hamiltonian by biquadratic spin interactions [73]. Biquadratic terms can also arise from spin-lattice coupling [74]. Spin-lattice coupling can lead to a structural distortion which strengthens some bonds while weakening others. We explain below how in the harmonic approximation about an ordered state, the effect of biquadratic terms is to alter exchange constants in the way one would expect from a structural distortion. In this chapter we will take the difference in the exchange constant between strong and weak bonds to be $2\Delta$.

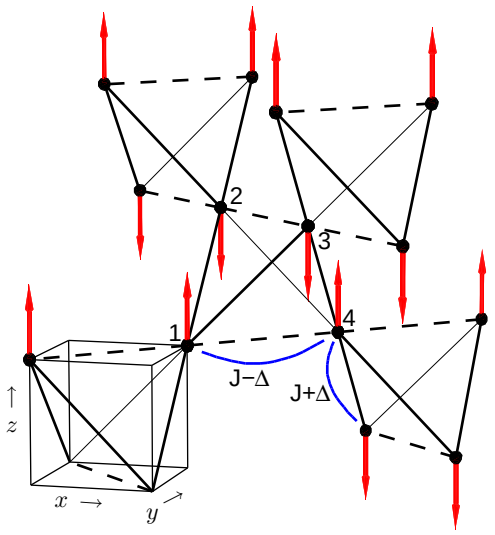


Figure 7.1: Collinear ordered state stabilised when $\Delta > 0$. In the flux language, the ordered state is fully polarized in the $z$ direction, corresponding to the maximum value possible for $|\mathbf{B}_z|$.

There has been a limited amount of earlier work in this area which can be summarised as follows. Tchernyshyov *et al.* [74] investigated spin-lattice coupling in pyrochlore antiferromagnets and showed that for the state in Fig. 7.1, zero modes are lifted to finite frequency and cause a divergence in the density of states at $\omega = 8\Delta$. We show below that some zero modes are lifted at a different scale, set by $\sqrt{J\Delta}$, but remain degenerate. In any collinear state with only nearest neighbour interactions, Hizi and Henley [66] established the existence of zero modes in the harmonic approximation which appear along the cubic axes of reciprocal space. We investigate how such modes are lifted by perturbations. Canals [75] studied the checkerboard lattice with a parameter which behaves like our $\Delta$ to investigate the stability of Néel order. As we mentioned in Chapter 2, the pyrochlore lattice with periodic boundary conditions imposed in the $z$-direction is the same as the planar pyrochlore or checkerboard lattice, thus for $q_z = 0$, we expect an overlap in our results.

We first derive the Hamiltonian which governs fluctuations about the ordered state, examining the static correlations. We then calculate the dynamic mode spectrum both for the nearest neighbour case, for which the ordered state is not stable, and also the case with perturbations. We will work in the basis of $\mathbf{m}$ and $\{\mathbf{B}_i\}$ fields since they elucidate the physics, and act as canonically conjugate variables in the linearised dynamics. We will also discuss further neighbour interactions. Here, spatial

components will always be subscripted, spin components superscripted, and boldface indicate the vector of spin components.

## 7.1 Static correlations

A ferromagnetic third neighbour interaction, $J_3$, will induce ferromagnetic order on each sublattice and the extensive degeneracy is reduced to the degeneracy of a single tetrahedron of classical spins. The remaining degeneracy retains a non-trivial character because the degeneracy of the classical tetrahedron is accidental. We include $J_3$ here to investigate how the non-trivial degeneracy is distinguished from the symmetry related degeneracy in the dynamics. The other perturbation away from the nearest neighbour model we consider is a biquadratic interaction selecting collinear states. Thus we consider a model such that a nearest neighbour bond contributes an energy $J_1\mathbf{S}_i \cdot \mathbf{S}_j - \frac{\Delta}{2}(\mathbf{S}_i \cdot \mathbf{S}_j)^2$ while a third neighbour bond contributes an energy $J_3\mathbf{S}_i \cdot \mathbf{S}_j$. This leads to the Hamiltonian

$$H = \frac{1}{2}\sum_{ij}(J_1 A_{ij}^{(1)} + J_3 A_{ij}^{(3)})\mathbf{S}_i \cdot \mathbf{S}_j - \frac{1}{2}\sum_{ij}\frac{\Delta}{2}A_{ij}^{(1)}(\mathbf{S}_i \cdot \mathbf{S}_j)^2 \tag{7.1}$$

where $A_{ij}^{(m)}$ is an $m$-th neighbour adjacency matrix. Up to spin-rotation symmetry, Fig. 7.1 shows the ordered state selected.

At low temperatures, configurations will be close to that in Fig. 7.1 and we can expand spins in deviations from collinearity as

$$\begin{pmatrix} \sigma^x \\ \sigma^y \\ 1 - \frac{1}{2}[(\sigma^x)^2 + (\sigma^y)^2] \end{pmatrix} \text{ and } \begin{pmatrix} \sigma^x \\ \sigma^y \\ -1 + \frac{1}{2}[(\sigma^x)^2 + (\sigma^y)^2] \end{pmatrix}. \tag{7.2}$$

By substituting these into Eq. (7.1), one obtains the Hamiltonian that governs the low energy physics. In substituting, Eq. (7.2) into Eq. (7.1), it is only necessary to keep quadratic terms - constant terms can be removed by redefining the zero of energy to be the classical ground state, and at low enough temperatures, quartic terms are irrelevant. For this reason, we use the symbol $\sim$ to represent the equivalence relation defined by keeping only quadratic terms in an expression. In a collinear state, all bonds are either parallel (+) or antiparallel (−). For a parallel bond,

$$(\mathbf{S}_i \cdot \mathbf{S}_j)_+ \sim \frac{1}{2}\sum_{\alpha=x,y} 2\sigma_i^\alpha \sigma_j^\alpha - (\sigma_i^\alpha)^2 - (\sigma_j^\alpha)^2$$

and

$$(\mathbf{S}_i \cdot \mathbf{S}_j)_+^2 \sim 2(\mathbf{S}_i \cdot \mathbf{S}_j)_+$$

whereas for an antiparallel bond

$$(\mathbf{S}_i \cdot \mathbf{S}_j)_- \sim \frac{1}{2} \sum_{\alpha=x,y} 2\sigma_i^\alpha \sigma_j^\alpha + (\sigma_i^\alpha)^2 + (\sigma_j^\alpha)^2$$

and

$$(\mathbf{S}_i \cdot \mathbf{S}_j)_-^2 \sim -2(\mathbf{S}_i \cdot \mathbf{S}_j)_- \,.$$

It follows that in the Hamiltonian expanded to harmonic order, biquadratic exchange can be absorbed in a redefinition of the nearest neighbour exchange constant which depends on whether the bond is parallel or antiparallel, since

$$J_1(\mathbf{S}_1 \cdot \mathbf{S}_2)_\pm - \frac{\Delta}{2}(\mathbf{S}_1 \cdot \mathbf{S}_2)_\pm^2 \sim (J_1 \mp \Delta)(\mathbf{S}_1 \cdot \mathbf{S}_2)_\pm \,.$$

The net effect of the substitutions is a quadratic Hamiltonian

$$H = \frac{1}{2} \sum_{\alpha=x,y} \sum_{ij} (J_1 V_{ij}^{(1)} + \Delta V_{ij}^{(\Delta)} + J_3 V_{ij}^{(3)}) \sigma_i^\alpha \sigma_j^\alpha \tag{7.3}$$

where $V_{ij}^{(1)}$ and $V_{ij}^{(3)}$ are the interaction matrices of previous chapters, and $V_{ij}^{(\Delta)} = A_{ij}^{(1-)} - A_{ij}^{(1+)} + 6\delta_{ij}$. Eq. (7.3) is the Hamiltonian which governs low energy fluctuations.

The analysis of static properties of Eq. (7.3) follows previous chapters, since $\exp[-\beta H]$ is a Gaussian weight for two independent spin components but there is no entropic contribution. The temperature dependence in the Gaussian approximation is trivial, since temperature can be absorbed in a rescaling $\sigma \to \sigma/\sqrt{T}$. In reciprocal space, the change of basis, 3.1.5, to $\mathbf{m}$ and $\mathbf{B}_i$ fields diagonalises the quadratic form in 7.3 along certain symmetry directions. For $\Delta = 0$, the flux variables have correlations

$$\langle B_i^\alpha(\mathbf{q}) B_j^\beta(-\mathbf{q}) \rangle \propto \delta_{\alpha\beta} \frac{T}{8J_3 a^2 q^2} \left[ \delta_{ij} - \frac{q_i q_j}{(1 + 8J_3/J_1)q^2} \right] \quad \alpha, \beta = x, y; \quad i, j = x, y, z \tag{7.4}$$

where Greek letters label spin components and Roman letters spatial. Since the term in square brackets is finite (although singular), the correlation function in Eq. (7.4) indicates a susceptibility that diverges like $\frac{1}{q^2}$ as $q \to 0$. Spontaneously breaking spin rotation symmetry gives rise to two soft directions in phase space, but Eq. (7.4) indicates six equivalent soft directions, which is four too many to be accounted for by broken spin rotation symmetry alone. The missing 'symmetry' is the classical degeneracy of the single tetrahedron. The degeneracy is accidental, in not being a consequence of a symmetry group, yet at low enough energy the statistical mechanical

consequences can be indistinguishable from a true symmetry - the symmetry restoring direction is $\mathbf{B}_z$, while $\mathbf{B}_x$ and $\mathbf{B}_y$ move among classically degenerate yet inequivalent points in phase space. Eq. (7.4) is invariant under $O(3)$ transformations of the spatial flux components. In the dynamics, one expects to find Goldstone modes associated with the broken symmetry, as well as Goldstone-like modes associated with the breaking of the accidental degeneracy. We will find below that these modes have different dispersions.

The accidental degeneracy is lifted at finite $\Delta$ and one expects to find a divergent susceptibility only for $\mathbf{B}_z$, the true Goldstone mode. To understand the role of $\Delta$, we consider the interaction $V^{(\Delta)}$ which appears in Eq. (7.3), expressed in the flux field basis, $(\mathbf{m}, \{\mathbf{B}_i\})$:

$$PV^{(\Delta)}(\mathbf{q})P^T = \begin{pmatrix} 8-a^2q_z^2 & -a^2q_yq_z & -a^2q_xq_z & a^2q_xq_y \\ -a^2q_yq_z & 8-a^2q_y^2 & -a^2q_xq_y & a^2q_xq_z \\ -a^2q_xq_z & -a^2q_xq_y & 8-a^2q_x^2 & a^2q_yq_z \\ a^2q_xq_y & a^2q_xq_z & a^2q_yq_z & a^2q^2 \end{pmatrix} .$$

The first comment is the structural similarity with the nearest neighbour interaction matrix, Eq. (3.11). This follows from the exact symmetry of any classical model of inverting a set of spins while simultaneous changing the sign of the affected exchange interactions. The second is that $\mathbf{B}_x$ and $\mathbf{B}_y$ acquire a flux stiffness $8\Delta$, which ensures a non-divergent susceptibility. The stiffness here is energetic in origin, rather than entropic as in previous chapters. By lifting the accidental classical degeneracy, only the symmetry-related degeneracy gives rise to a divergent susceptibility. The explicit form is complicated for general $J_3$, and $\Delta$, even at long wavelengths. With $J_3 = 0$ and $q_z = 0$, the transverse flux correlations are approximately

$$\langle B_i^\alpha(\mathbf{q})B_j^\beta(-\mathbf{q})\rangle_{q_z=0} \propto \frac{T}{8\Delta}\left[\delta_{ij} - \frac{q_iq_j}{q^2+\xi^{-2}}\right] \quad \alpha,\beta = x,y; \quad i,j = x,y$$

with $\xi = \sqrt{J_1a^2/8\Delta}$. We have ignored corrections proportional to $\Delta/J_1$. While statically, $\Delta$ gives a correlation length to transverse flux fluctuations, we will see how it also lifts modes to finite frequency in the dynamics.

## 7.2 Spin wave dynamics

We turn now to the dynamic behaviour of the harmonic fluctuations about the ordered state in Fig. 7.1. By solving the linearised classical equations of motion, one obtains the mode frequencies. Equivalently, a linear spin wave calculation using the Holstein-Primakoff boson representation, Eq. (1.8), for the quantum spin operators gives the

same frequencies, interpreted as a spectrum for quantized oscillator modes. The approach we take is to study the equations of motion. All the results derived below have been verified with a Holstein-Primakoff boson calculation.

For the frustrated pyrochlore antiferromagnet, we have argued throughout that it is useful to think in terms of the variables $\mathbf{m}$ and $\mathbf{B}_i$, rather than the separate sublattice spins, when in reciprocal space. If considered as continuum fields, the Hamiltonian is then a functional, $\mathcal{H}[\mathbf{m}, \{\mathbf{B}_i\}]$. By their definition in terms of the sublattice spins, the fields have Poisson-bracket relations

$$\begin{aligned} \{m^\alpha(\mathbf{r}), m^\beta(\mathbf{r}')\} &= \frac{1}{2}\epsilon^{\alpha\beta\gamma} m^\gamma(\mathbf{r})\delta(\mathbf{r}-\mathbf{r}') \\ \{m^\alpha(\mathbf{r}), B_i^\beta(\mathbf{r}')\} = \{B_i^\alpha(\mathbf{r}), m^\beta(\mathbf{r}')\} &= \frac{1}{2}\epsilon^{\alpha\beta\gamma} B_i^\gamma(\mathbf{r})\delta(\mathbf{r}-\mathbf{r}') \\ \{B_i^\alpha(\mathbf{r}), B_j^\beta(\mathbf{r}')\} &= \frac{1}{2}\epsilon^{\alpha\beta\gamma}\left(|\epsilon_{ijk}| B_k^\gamma(\mathbf{r}) + \delta_{ij} m^\gamma(\mathbf{r})\right)\delta(\mathbf{r}-\mathbf{r}') \end{aligned}$$

where $|\epsilon_{ijk}|$ is the symbol which is 1 whenever all indices are different. For notational simplicity below, we will summarise this information in a reduced bracket, $\{\cdots\}_R$, where the spatial delta function and the dependence on spin components are implied, for example

$$\{\mathbf{B}_i, \mathbf{B}_j\}_R = \frac{1}{2}\left(|\epsilon_{ijk}|\mathbf{B}_k + \delta_{ij}\mathbf{m}\right) .$$

The presence of $|\epsilon_{ijk}|$ indicates a cubic anisotropy. This is expected, since the only symmetry-equivalent groundstates of a single frustrated tetrahedron of classical spins are those linked by rotation of all spins, and the discrete symmetries of the tetrahedron.

The equations of motion of the classical dynamics are obtained from Hamilton's equations:

$$\begin{aligned} \frac{\partial \mathbf{m}}{\partial t} &= -\{\mathbf{m}, \mathbf{m}\}_R \times \frac{\delta \mathcal{H}}{\delta \mathbf{m}} - \sum_j \{\mathbf{m}, \mathbf{B}_j\}_R \times \frac{\delta \mathcal{H}}{\delta \mathbf{B}_j} \\ \frac{\partial \mathbf{B}_i}{\partial t} &= -\{\mathbf{B}_i, \mathbf{m}\}_R \times \frac{\delta \mathcal{H}}{\delta \mathbf{m}} - \sum_j \{\mathbf{B}_i, \mathbf{B}_j\}_R \times \frac{\delta \mathcal{H}}{\delta \mathbf{B}_j} . \end{aligned} \tag{7.5}$$

Although the total magnetization drops out of the statics at low temperatures and long wavelengths, its role as a conserved density and the generator of the global spin rotations makes it an essential ingredient of the dynamics. With an appropriately defined Hamiltonian, the functional derivatives in 7.5 will vanish in the groundstate, and be linear in small deviations away from it, thus for linearised equations of motion we need only consider the groundstate contribution to the reduced bracket expressions in Eq. (7.5), i.e., the only non-vanishing Poisson-brackets are those which equal

$\mathbf{B}_z^{(0)}$, with magnitude $2S$. Thus in the linearised dynamics of the transverse spin components, (alternate spin components of) $\mathbf{m}$ and $\mathbf{B}_z$ form a canonically conjugate pair, as do $\mathbf{B}_x$ and $\mathbf{B}_y$. The resulting equation of motion are

$$\frac{\partial \mathbf{m}}{\partial t} = -\frac{1}{2}\mathbf{B}_z^{(0)} \times \frac{\delta \mathcal{H}}{\delta \mathbf{B}_z} \tag{7.6}$$

$$\frac{\partial \mathbf{B}_x}{\partial t} = -\frac{1}{2}\mathbf{B}_z^{(0)} \times \frac{\delta \mathcal{H}}{\delta \mathbf{B}_y} \tag{7.7}$$

$$\frac{\partial \mathbf{B}_y}{\partial t} = -\frac{1}{2}\mathbf{B}_z^{(0)} \times \frac{\delta \mathcal{H}}{\delta \mathbf{B}_x} \tag{7.8}$$

$$\frac{\partial \mathbf{B}_z}{\partial t} = -\frac{1}{2}\mathbf{B}_z^{(0)} \times \frac{\delta \mathcal{H}}{\delta \mathbf{m}}\,. \tag{7.9}$$

Since the Hamiltonian functional is a symmetric quadratic form acting among $\mathbf{m}$ and $\{\mathbf{B}_i\}$, one can recognise Eq. (7.6) as a matrix equation where the matrix elements are those of the quadratic form in the Hamiltonian, but with the row order reversed. We call this matrix the dynamic matrix, and the mode frequencies are the eigenvalues of the dynamic matrix.

### 7.2.1 Nearest neighbour interactions only

For the nearest neighbour model, it is helpful to look explicitly at the dynamic matrix:

$$\begin{pmatrix} -2s_xs_y & 2s_xs_z & 2s_ys_z & 4-G_z \\ -2s_xs_z & 2s_xs_y & 4-G_y & 2s_ys_z \\ -2s_ys_z & 4-G_x & 2s_xs_y & 2s_xs_z \\ 4-G_M & -2s_ys_z & -2s_xs_z & -2s_xs_y \end{pmatrix}$$

where $s_a \equiv \sin\left(\frac{q_a}{4}\right)$, $c_a \equiv \cos\left(\frac{q_a}{4}\right)$ and the $G$ functions were defined in Chapter 3 preceding 3.11. For a general $\mathbf{q}$, the matrix is diagonalisable by similarity transformation. It is clear though, that along the axes of reciprocal space, the majority of the entries vanish. In fact along $q_x$ for example,

$$\begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 8\sin^2(q_x/8) & 0 & 0 \\ 8\cos^2(q_x/8) & 0 & 0 & 0 \end{pmatrix}$$

This matrix is singular, with all eigenvalues zero. Hizi and Henley [66] studied general collinear states and identified modes whose frequencies vanish due to Jordan blocks in the normal form of the matrix governing the linearised dynamics as above, and called them divergent modes. We are considering a specific collinear state, and the divergent modes appear along the $x$ and $y$ reciprocal space axes.

There are other places in reciprocal space also, where the frequencies go to zero. This occurs along the lines $(1, 0, q_z)$ and $(0, 1, q_z)$ in reciprocal space. These lines of degeneracy are interesting since they remain exactly degenerate even in the presence of the perturbations. Here, the dynamic matrix takes the form

$$\begin{pmatrix} 0 & 0 & 2s_z & 2-2c_z \\ 0 & 0 & 2+2c_z & 2s_z \\ -2s_z & 2-2c_z & 0 & 0 \\ 2+2c_z & -2s_z & 0 & 0 \end{pmatrix}$$

with all eigenvalues zero.

At a general point in reciprocal space for the nearest neighbour model, the eigenvalues are

$$\omega = 0, 0, \pm\sqrt{2}\sqrt{4 - G_z(2\mathbf{q})}$$

where $G_z$ is as defined in Eq. (3.10) (and the factor of 2 is not a typo). In the plane defined by $q_y = 0$, it reduces to $\omega = \pm 4J_1 S|\cos(q_x/4)||\sin(q_z/4)|$. At long wavelengths, this describes antiferromagnetic spin waves with a group-velocity $J_1 S\sqrt{8a^2}$, propagating in the $z$ direction. The group velocity vector in this plane varies very little with the direction of the wavevector. In the plane defined by $q_z = 0$, the dispersion is quadratic in wavevector $\omega = \pm 4J_1 S|\sin(q_x/4)||\sin(q_y/4)|$. The cubic anisotropy is clear. Since the ordered state picks out the $z$-direction, one might expect behaviour in the plane defined by $q_z = 0$ to be isotropic whereas in fact it is highly anisotropic. Fig. 7.2 shows the dispersion in the plane $q_z = 0$.

In summary, the nearest neighbour harmonic modes has two flat bands at $\omega = 0$ of generic zero modes, and two degenerate bands of finite frequency. The finite frequency modes touch $\omega = 0$ along special lines in reciprocal space, where the dynamic matrix is non-diagonalisable, giving non-generic zero modes.

The pure nearest neighbour case is slightly pathological, since the ordered state is a neutral, rather than stable, equilibrium, and the macroscopic number of zero modes is evidence for this. As we discussed earlier in this chapter, the ordered state in Fig. 7.1 can be stabilised by additional interactions. Such interactions will change the mode frequencies. Typically, the frequency change will be proportional to the strength of the perturbation. For a perturbation of strength $\varepsilon$, one expects to be able to express the change in eigenvalue as a series $\delta\omega = \varepsilon\delta\omega^{(1)} + O(\varepsilon^2)$. However, when the unperturbed matrix is not diagonalisable, this generically fails. To demonstrate the point, the eigenvalues of

$$\begin{pmatrix} 0 & \varepsilon \\ J & 0 \end{pmatrix}$$

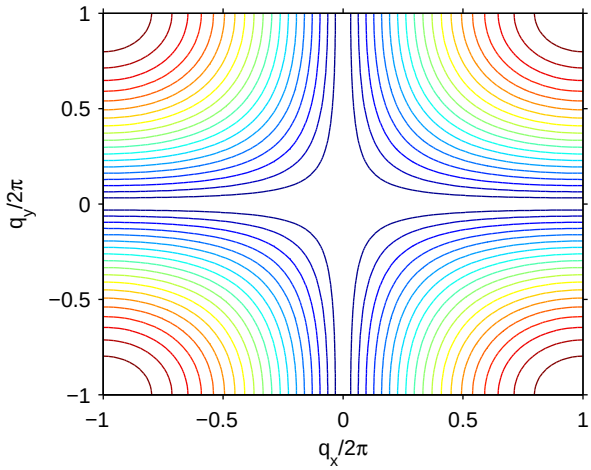


Figure 7.2: Contour plot of $\omega = 4J_1S|\sin(q_x/4)||\sin(q_y/4)|$, the spin wave dispersion for nearest neighbour interactions in the plane $q_z = 0$. The frequency vanishes along the reciprocal space axes.

are $\omega = \pm\sqrt{J\varepsilon}$, which is non-analytic at $\varepsilon = 0$. In analysing the effect of additional interactions which stabilise the ordered state, we expect to find this non-analytic behaviour for the non-generic zero modes.

The procedure to obtain the mode frequencies for the ordered state shown is as with the nearest neighbour case: the mode frequencies are the eigenvalues of the dynamic matrix, which is obtained from reversing the row order of the interaction matrix in Eq. (7.3) when expressed in the flux basis. We now solve the eigenvalue problem 'by hand' along the symmetry axes in the presence of the perturbations, and complement this with plots obtained from numerical solutions.

### 7.2.2 $J_3$

The purpose of studying the effect of $J_3$ while $\Delta = 0$ is to show that the dynamics distinguishes between the three flux components, even though their static correlations are identical. By working with the long wavelength forms for the interaction matrix, the functional derivatives in Eq. (7.6) are

$$\frac{\delta H}{\delta \mathbf{m}} = 8J_1\mathbf{m} + \frac{1}{2}|\varepsilon_{ijk}|J_1a^2\partial_i\partial_j\mathbf{B}_k \tag{7.10}$$

$$\frac{\delta H}{\delta \mathbf{B}_i} = -J_1a^2\partial_i(\partial_k\mathbf{B}_k) + \frac{1}{2}|\varepsilon_{ijk}|J_1a^2\partial_j\partial_k\mathbf{m} + 8J_3a^2\partial_k\partial_k\mathbf{B}_i \tag{7.11}$$

The equations are simplified along the cubic axes of reciprocal space, since the terms proportional to $|\varepsilon_{ijk}|$ vanish along those high symmetry directions. For wavevectors

along the $z$ direction, the equations of motion are

$$\frac{\partial \mathbf{m}}{\partial t} = \frac{1}{2}\mathbf{B}_z^{(0)} \times (J_1 - 8J_3)a^2\partial_z^2 \mathbf{B}_z \qquad \frac{\partial \mathbf{B}_x}{\partial t} = -\frac{1}{2}\mathbf{B}_z^{(0)} \times 8J_3 a^2 \partial_z^2 \mathbf{B}_y \tag{7.12}$$

$$\frac{\partial \mathbf{B}_z}{\partial t} = -\frac{1}{2}\mathbf{B}_z^{(0)} \times 8J_1 \mathbf{m} \qquad \frac{\partial \mathbf{B}_y}{\partial t} = -\frac{1}{2}\mathbf{B}_z^{(0)} \times 8J_3 a^2 \partial_z^2 \mathbf{B}_x \,. \tag{7.13}$$

Using that for the reference groundstate, $|\mathbf{B}_z^{(0)}| = 2S$, the mode frequencies are

$$\omega \simeq \pm(J_1 - 4J_3)S\sqrt{8a^2}|q_z| \qquad \omega \simeq \pm|J_3|S8a^2 q_z^2 \,. \tag{7.14}$$

Compared to the nearest neighbour result, the zero modes have acquired a quadratic dispersion. The true Goldstone modes are the linearly dispersing antiferromagnetic spin waves, which come from dynamic couplings between $\mathbf{m}$ and the antiferromagnetic order parameter $\mathbf{B}_z$ and this behaviour is just as in an ordinary antiferromagnet. In contrast, the modes with quadratic dispersion are associated with the accidental degeneracy of the frustrated tetrahedron.

Along the $x$ and $y$ reciprocal space axes, we explained that frequencies vanish due to a non-diagonalisable dynamic matrix. To see how mode frequencies are lifted with additional interactions, we consider modes with wavevector in the $x$ direction. In this case the equations of motion are

$$\frac{\partial \mathbf{m}}{\partial t} = -\frac{1}{2}\mathbf{B}_z \times 8J_3 a^2 \partial_x^2 \mathbf{B}_z \qquad \frac{\partial \mathbf{B}_x}{\partial t} = -\frac{1}{2}\mathbf{B}_z \times 8J_3 a^2 \partial_x^2 \mathbf{B}_y \tag{7.15}$$

$$\frac{\partial \mathbf{B}_z}{\partial t} = -\frac{1}{2}\mathbf{B}_z \times 8J_1 \mathbf{m} \qquad \frac{\partial \mathbf{B}_y}{\partial t} = \frac{1}{2}\mathbf{B}_z \times (J_1 - 8J_3)a^2 \partial_x^2 \mathbf{B}_x \,. \tag{7.16}$$

with mode frequencies

$$\omega \simeq \pm\sqrt{8|J_3|J_1}\sqrt{8a^2}S|q_x| \qquad \omega \simeq \pm(1 - 4J_3/J_1)\sqrt{\frac{J_1|J_3|}{8}}S(8a^2)q_z^2 \,. \tag{7.17}$$

Again, the linearly dispersing modes are the Goldstone modes associated with the breaking of spin rotation symmetry and they connect smoothly to the linearly dispersing modes in the $z$-direction. The dispersion is anisotropic and the group velocity is mainly in the direction of flux polarisation, as indicated in Fig. 7.3. The dependence on $\sqrt{J_3}$ is just as in the example of a two-dimensional matrix close to being non-diagonalisable given above. In addition to the linearly dispersing true Goldstone mode, we find the quadratic dispersing Goldstone-like mode associated with the accidental degeneracy, arising from dynamic coupling between $\mathbf{B}_x$ and $\mathbf{B}_y$.

The other non-generic zero modes along $(1, 0, q_z)$ are lifted to finite frequency, but remain four-fold degenerate, and will be discussed at the end of the next section.

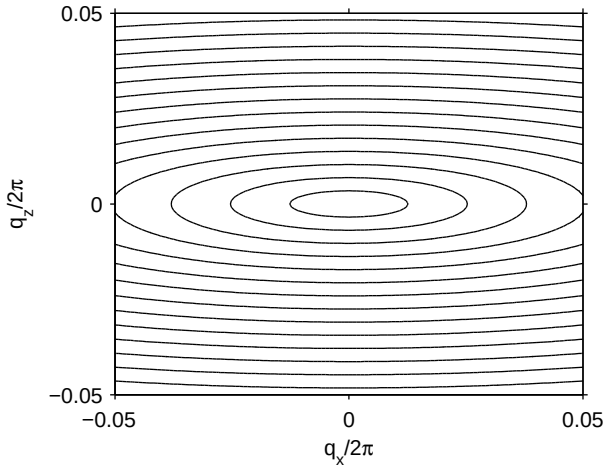


Figure 7.3: Contours of dispersion for linearly dispersing spin waves for $q_y = 0$ at $J_3 = -0.01$. Strong anisotropy ensures the group velocity is along the $z$-direction, the direction of flux polarization.

### 7.2.3 $\Delta$

Having established that the Goldstone-like modes associated with the accidental degeneracy behave very differently to the true Goldstone modes which come from breaking spin symmetry, we turn to situations where the accidental degeneracy is lifted by weak perturbations and the Goldstone-like modes are gapped. We set $J_3 = 0$.

Proceeding with a similar analysis to those presented above, one finds for modes propagating in the $z$-direction

$$\frac{\partial \mathbf{m}}{\partial t} = -\frac{1}{2}\mathbf{B}_z \times \left[(J+\Delta)(-a^2\partial_z^2 \mathbf{B}_z)\right] \qquad \frac{\partial \mathbf{B}_x}{\partial t} = -\frac{1}{2}\mathbf{B}_z \times 8\Delta \mathbf{B}_y \tag{7.18}$$

$$\frac{\partial \mathbf{B}_z}{\partial t} = -\frac{1}{2}\mathbf{B}_z \times [(J+\Delta)8\mathbf{m}] \qquad \frac{\partial \mathbf{B}_y}{\partial t} = -\frac{1}{2}\mathbf{B}_z \times 8\Delta \mathbf{B}_x \tag{7.19}$$

with frequencies

$$\omega = \pm(J+\Delta)S\sqrt{8a^2}|q_z| \qquad \omega = \pm 8\Delta\,. \tag{7.20}$$

Along with conventional linearly dispersing spin waves, the zero modes of the nearest neighbour model are lifted to non-dispersing finite frequency modes. Such non-dispersing modes at finite frequency are a remnant of the extensive ground state degeneracy - Tchernyshyov *et al.* [74] called them string modes, and they give rise to a divergence in the density of states at $8\Delta$. The non-dispersing modes at $\omega = 8\Delta$ can be seen in the final segment of Fig. 7.4, where the dispersion of the four bands is plotted along a path in reciprocal space. This has important experimental consequences since it implies a finite frequency resonance in scattering experiments.

The lifting of the zero modes to finite frequency can come from any interaction which breaks the classical degeneracy of the single tetrahedron - $Y_2Ru_2O_7$ has a magnetic transition without a structural distortion, and in the ordered phase shows a large band of spectral weight at finite frequency with no dispersion. However, no Goldstone modes are observed at all, suggesting spin rotation symmetry is explicitly broken and spin space anisotropy may be responsible for lifting the zero modes [49].

As before, more interesting behaviour is found in spatially transverse fluctuations. For modes along $q_x$, there are linearly dispersing Goldstone modes with $\omega = cq_x$ where $c = \sqrt{J_1\Delta}S\sqrt{8a^2}$, in the limit $\Delta \ll J_1$. The dispersion of the gapped modes is $\omega = \sqrt{(8\Delta)^2 + q_x^2c^2}$ which is the dispersion of a relativistic particle with mass $8\Delta$ with respect to the velocity $c$. Here, the zero modes are not just lifted to finite frequency, they also acquire a relativistic dispersion. The first segment of Fig. 7.4 shows this.

Finally, we turn to the fate of the non-generic zero-modes along $(1, 0, q_z)$. Along these lines in reciprocal space, all four bands are degenerate and remain so when lifted to finite frequency. Reinstating $J_3$, the frequency is

$$\omega = 4\sqrt{(J_1 + 2\Delta - 4J_3)(\Delta - 4J_3)}\,. \tag{7.21}$$

This can be seen in the second segment in Fig. 7.4. As described above, the behaviour is non-analytic at $\Delta = J_3 = 0$, due to the non-diagonalisability of the dynamic matrix.

## 7.3 Discussion

To conclude we discuss the relevance of these results. In the ordered magnet, the spin wave dispersion for fluctuations out of the nearest neighbour groundstate manifold is highly anisotropic, with a group velocity in the direction of flux polarisation. We can speculate that this will carry over to locally ordered sections in the paramagnetic phase. A group-velocity for spin waves in the direction of flux polarisation may play a role in the microscopic mechanism for the entropic forces which drive the relaxation in the low temperature paramagnet discussed in Chapters 3 and 4.

The precessional spin dynamics gives rise to a mode structure which is more subtle than one might expect from thinking about the statics alone. Cubic anisotropy is built in to the Poisson-bracket relations of the collective variables which describe the low energy states of the pyrochlore magnet. It would be interesting to understand in more detail how the various features uncovered in the linearised semiclassical description appear in the strongly quantum limit.

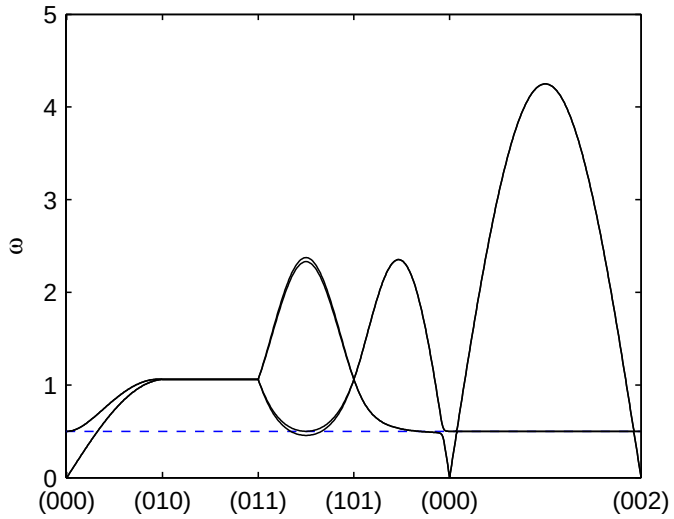


Figure 7.4: Mode frequencies at $J_1 = 1, \Delta = 0.0625, J_3 = 0$ along lines in reciprocal space. The dashed line shows the scale $8\Delta$. Generically, the four bands have different frequencies. Along certain lines in reciprocal space, all four bands are exactly degenerate with frequency in Eq. (7.21).

# Appendix A

# Simulations and numerical work

This appendix provides some details about the simulations which have been performed, and numerical calculations.

## A.1 Simulations

The simulations of the classical dynamics of the pyrochlore antiferromagnet were written in C++, compiled with the GNU compiler collection, and run on Apple iMac and MacPro desktop machines. The simulation results reported for the dynamic structure factor are obtained from an average over a large number of runs. Introduce a notation whereby $\mathcal{C}$ refers to a configuration of the spins, i.e, a single point in the classical phase space of the system. A run involves three processes:

1. Monte Carlo sampling is used to draw an initial configuration from the thermal distribution:

$$P[\mathcal{C}] \propto \exp(-\beta H[\mathcal{C}])$$

   Call this initial configuration $\mathcal{C}^r$ where $r$ is a run label.

2. Numerical integration of the classical equations of motion evolves the configuration forward in time by $\tau = (n-1)\Delta t$, while generating a configuration every $\Delta t$. The set of $n$ time-evolved configurations is

$$\mathcal{E}^r = \{\mathcal{C}^r(0), \mathcal{C}^r(\Delta t), \mathcal{C}^r(2\Delta t), \ldots \mathcal{C}^r(\tau - \Delta t), \mathcal{C}^r(\tau)\}\,.$$

   The time-evolution is Hamiltonian, and so conserves the energy. Thus the time-evolution is microcanonical, but we associate the configurations with the temperature of the thermal distribution from which the initial configuration was sampled, $1/\beta$.

3. Finally, the set $\mathcal{E}_\beta^r$ is Fourier transformed in time and space to calculate $S_\beta^r(\mathbf{q},\omega)$. Defining $\Omega = \frac{2\pi}{\Delta t}$, and $\Delta\omega = \frac{2\pi}{n\Delta t}$, the dynamic structure factor is evaluated at $\omega = -\Omega, -\Omega + \Delta\omega, \ldots \Omega - 2\Delta\omega, \Omega - \Delta\omega$. It is important that $\Delta t$ is sufficiently small to avoid aliasing effects from fast modes, and that $\tau$ is sufficiently large to uncover the slowest relaxation times.

The procedure above is repeated for many runs, $R$, and the results averaged:

$$S(\mathbf{q},\omega) = \frac{1}{R}\sum_{r=1}^{R} S^r(\mathbf{q},\omega)\,.$$

For a sufficiently large number of runs, this will be a good estimate of the true dynamic structure factor as long as the simulation time $\tau$ is much longer than the autocorrelation time. Separate programs were written to implement the three steps above, and UNIX scripts managed the whole process.

### A.1.1 System sizes and Monte Carlo

The Monte Carlo sampling was done by the Metropolis Monte Carlo algorithm, such as described in Ref. 76 with single spin updates. The random number generator algorithm was the Tausworthe generator, as implemented in the GNU Scientific Library [77]. The step size was scaled with $\sqrt{T}$ (as in the simulations described in Ref. 78) so that the acceptance ratio was 0.5 which ensures efficient sampling. At low temperatures, energy equilibration is rapid, relative to the equilibration time for the collective flux degrees of freedom. This timescale is given by the spin autocorrelation time, which was found to be approximately $20\beta$, measured in Monte Carlo steps per spin (MCS).

The system was rhombohedral, generated by the fcc lattice vectors in Sec. 3.1.1 and a tetrahedron at each site. The total number of spins was $4L^3$, where $L$ took values 8, 16 and 32. With hindsight, a simple cubic lattice with a 16 site basis would have been preferable for the presentation of results in reciprocal space.

### A.1.2 Molecular Dynamics

Numerical integration of equations of motion, often called molecular dynamics (MD), evolves configurations in time. The equations of motion for the spins are a set of coupled first-order ordinary differential equations in $12L^3$ dynamical variables. For the largest systems, $L = 32$, this is 393216 variables. The equations of motion conserve total energy, and total spin, by virtue of symmetries of the Hamiltonian.

Each spin also has conserved length. While numerical integration schemes such as symplectic integrators exist which conserve energy by construction, our integration routine did not build in the conservation laws - instead they are regarded as tests of the accuracy of the implementation.

The integration scheme was a 4th order Runge-Kutta algorithm. The implementation used `gsl_odeiv_step_rk4` with adapative step size control from the GNU Scientific Library [77]. Even with the largest systems, the routine was unproblematic in so far as the conservation laws were not violated outside the specified error tolerance. As a further check of the correctness of the results, the numerical spin wave spectrum in ordered states at very low temperatures matched calculations from linear spin wave theory.

The program output a configuration every $\Delta t$. We took $\Delta t = 0.8J^{-1}$, probing a maximum frequency of $\frac{\pi}{0.8}J \simeq 3.92J$ which comfortably includes the bandwidth of the frequency spectrum, as for example in Fig. 4.8.

### A.1.3 Fourier transformation

The Fourier transform routines were implemented using the *Fastest Fourier Transform in the West* library [79], with appropriately adjusted normalisation for the transforms. Each of the four sublattices is transformed separately, and multiplied by the appropriate phase factor from the position in the primitive cell. The calculation for the dynamic structure factor of a run is then $S(\mathbf{q},\omega) = |\sum_{\mu=1}^{4} \mathbf{S}_\mu(\mathbf{q},\omega)|^2$. Symmetries of reciprocal space were used to give better statistics from a single run in the paramagnetic phase.

We also calculated a dynamical correlation function $S(\mathbf{q},t) \equiv \langle \mathbf{S}(\mathbf{q},t) \cdot \mathbf{S}(-\mathbf{q},0)\rangle$ such as appears in Fig. 4.4. For this, there was no artificial imposition of periodic boundary conditions in time, as generically exists for the discrete Fourier transform. Rather, $S(\mathbf{q}, m\Delta t)$ is calculated as an average of $n - m$ products of configurations. This time-averaging is in addition to the Monte Carlo run averaging.

## A.2 Numerical calculations

The calculations required for deriving correlation functions in the self-consistent Gaussian approximation were written as self-contained MATLAB programs. MATLAB routines for root finding and matrix inversion were used to numerically solve the self-consistency condition. Calculation of structure factors such as in Fig. 6.3 can be done very rapidly on a desktop machine in this way.

# Appendix B

# $O(n)$ invariant models

The classical Heisenberg model is the $n = 3$ member of a family of $O(n)$ invariant models, where the degrees of freedom (spins) live on an sphere in $n$-dimensions, and the Hamiltonian is invariant under global rotations of all spins by an element of the $O(n)$ group. The partition function of an $O(n)$ model is

$$Z = \int \left( \prod_{i=1}^{N} \rho(\mathbf{S}_i) d\mathbf{S}_i \right) \exp(-\beta H[\{\mathbf{S}_1, \ldots, \mathbf{S}_N\}])$$

where $\rho(\mathbf{S}) = \delta(\mathbf{S}^2 - n)$ is the single site distribution, expressing the spin-length constraint for the $n$-component vector $\mathbf{S}$ of spin components. When the Hamiltonian is $O(n)$ invariant and the symmetry is not spontaneously broken, integrating out all but one spin will recover the single site distribution:

$$\rho(\mathbf{S}_i) \propto \int \left( \prod_{j \neq i} \rho(\mathbf{S}_j) d\mathbf{S}_j \right) \exp(-\beta H) \,.$$

Thus single-site moments are thus independent of the Hamiltonian, depending only on the uniform distribution on the $n$-dimensional sphere.

The single-site distribution is clearly a joint distribution for $n$ components. Integrating out $n - q$ spin components leaves a joint distribution for the remaining $q$ components which can be calculated in hyperspherical coordinates. The result is that the joint distribution of $q < n$ Cartesian components is

$$P(x_1, x_2, \ldots, x_q) = \frac{1}{\mathcal{N}} \left( 1 - \frac{\sum_a x_a^2}{n} \right)^{(n-2-q)/2} \overset{\lim n \to \infty}{\simeq} \frac{1}{(2\pi)^{q/2}} \exp(-\sum_{a=1}^{q} x_a^2/2) \,. \tag{B.1}$$

where it is understood that $P = 0$ for $\sum_a x_a^2 > n$.[1] The derivation is not valid for $q = n$, but the simple pole which results from naively putting $q = n$ into Eq. (B.1)

[1] The family of marginal distributions on hyperspheres in Eq. (B.1) are also known as q-Gaussian distributions.

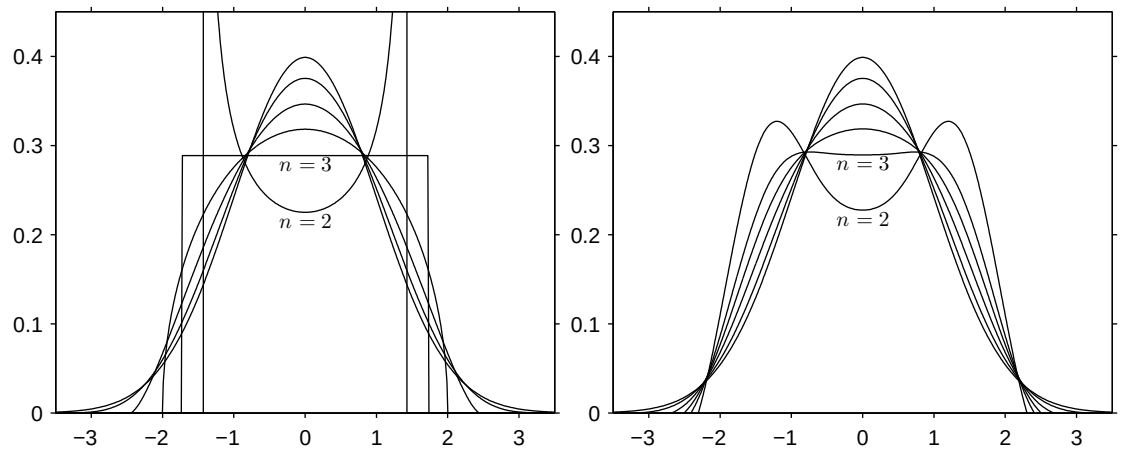


Figure B.1: Left: normalised probability density, Eq. (B.1), for a single Cartesian component induced by a uniform measure on the sphere of radius $\sqrt{n}$ in $n$-dimensions, shown for $n = 2, 3, 4, 6, 13, \infty$. Right: asymptotic expansion to second order in $1/n$ of the same quantity. The asymptotic expansion is negative for large $|x|$.

may loosely be understood as the Dirac delta function of the single-site distribution. For $q = 1$, the normalisation is $\mathcal{N} = \sqrt{n\pi}\Gamma(\frac{1}{2}(n-1))/\Gamma(n/2)$ and the convergence to the Gaussian limit as $n \to \infty$ is shown in Fig. B.1. An asymptotic expansion about $n = \infty$ can be obtained from the Taylor expansion of $\ln(1-x)$, and is shown for comparison also in Fig. B.1.

In a self-consistent Gaussian approximation, the single site distribution is

$$P(x_1, x_2, \ldots, x_q) = \frac{1}{(2\pi)^{q/2}} \exp(-\sum_{a=1}^{q} x_a^2/2)\,,$$

exactly the same as obtained in the $n \to \infty$ limit above. That the single site distribution is Gaussian follows from the fact that marginal distributions of a multivariate Gaussian are themselves Gaussian; that the variance is fixed follows from the self-consistency condition.

One also may consider the question in reverse, and ask how the spin length constraint is recovered from a Gaussian distribution when the number of components is large. Assume spin components are independent and normally distributed with variance $\sigma^2 = \langle x_a^2 \rangle = 1$. The variance of $y = x_a^2$ is 2. It follows that the random variable $r^2 = \mathbf{x}^2$ has mean $n$ and variance $2n$. Fluctuations of $r^2$ relative to the mean are $\mathcal{O}(n^{-1/2})$ and the spin length constraint, $r^2 = n$, is recovered in probability as $n \to \infty$.

# Bibliography


[1] N. W. Ashcroft and N. D. Mermin, *Solid state physics* (Thomson Learning, 1976).

[2] H. Callen, *Thermodynamics and an introduction to thermostatistics*, 2nd ed. (Wiley, 1985).

[3] B. I. Halperin and P. C. Hohenberg, "Hydrodynamic theory of spin waves," Phys. Rev., **188**, 898–918 (1969).

[4] J. Villain, "Insulating spin glasses," Z. Phys. B, **33**, 31–42 (1979).

[5] K. H. Fischer and J. A. Hertz, *Spin glasses* (Cambridge University Press, 1991).

[6] A. P. Ramirez, "Strongly geometrically frustrated magnets," Annual Review of Materials Science, **24**, 453–480 (1994).

[7] M. J. P. Gingras, "Spin ice," in *Introduction to Frustrated Magnetism*, edited by C. Lacroix, P. Mendels, and F. Mila (Springer, 2010).

[8] G. H. Wannier, "Antiferromagnetism. The Triangular Ising Net," Phys. Rev., **79**, 357–364 (1950).

[9] J. Stephenson, "Ising-model spin correlations on the triangular lattice," Journal of Mathematical Physics, **5**, 1009–1024 (1964).

[10] D. A. Huse and A. D. Rutenberg, "Classical antiferromagnets on the kagomé lattice," Phys. Rev. B, **45**, 7536–7539 (1992).

[11] J. T. Chalker, P. C. W. Holdsworth, and E. F. Shender, "Hidden order in a frustrated system: Properties of the Heisenberg Kagomé antiferromagnet," Phys. Rev. Lett., **68**, 855–858 (1992).

[12] R. Moessner and J. T. Chalker, "Properties of a Classical Spin Liquid: The Heisenberg Pyrochlore Antiferromagnet," Phys. Rev. Lett., **80**, 2929–2932 (1998).

[13] R. Moessner and J. T. Chalker, "Low-temperature properties of classical geometrically frustrated antiferromagnets," Phys. Rev. B, **58**, 12049–12062 (1998).

[14] P. W. Anderson, "Ordering and antiferromagnetism in ferrites," Phys. Rev., **102**, 1008–1013 (1956).

[15] C. L. Henley, "The 'Coulomb phase' in frustrated systems," Annu. Rev. Condens. Matter Phys., **1**, 179 (2010).

[16] S. V. Isakov, K. Gregor, R. Moessner, and S. L. Sondhi, "Dipolar spin correlations in classical pyrochlore magnets," Phys. Rev. Lett., **93**, 167204 (2004).

[17] C. L. Henley, "Power-law spin correlations in pyrochlore antiferromagnets," Phys. Rev. B, **71**, 014424 (2005).

[18] 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).

[19] R. Youngblood, J. D. Axe, and B. M. McCoy, "Correlations in ice-rule ferroelectrics," Phys. Rev. B, **21**, 5212–5220 (1980).

[20] C. L. Henley, "Relaxation time for a dimer covering with height representation," Journal of Statistical Physics, **89**, 483–507 (1997).

[21] A. E. Brouwer and W. H. Haemers, "Spectra of graphs," Chapter 1 of lecture notes at `http://www.win.tue.nl/~aeb/`.

[22] J. N. Reimers, A. J. Berlinsky, and A.-C. Shi, "Mean-field approach to magnetic ordering in highly frustrated pyrochlores," Phys. Rev. B, **43**, 865–878 (1991).

[23] J. N. Reimers, "Absence of long-range order in a three-dimensional geometrically frustrated antiferromagnet," Phys. Rev. B, **45**, 7287–7294 (1992).

[24] C. L. Henley, "Ordering by disorder: Ground-state selection in fcc vector antiferromagnets," Journal of Applied Physics, **61**, 3962–3964 (1987).

[25] T. H. Berlin and M. Kac, "The spherical model of a ferromagnet," Phys. Rev., **86**, 821–835 (1952).

[26] H. E. Stanley, "Spherical model as the limit of infinite spin dimensionality," Phys. Rev., **176**, 718–722 (1968).

[27] D. A. Garanin, "Self-consistent Gaussian approximation for classical spin systems: Thermodynamics," Phys. Rev. B, **53**, 11593–11605 (1996).

[28] B. Canals and D. A. Garanin, "Spin-liquid phase in the pyrochlore antiferromagnet," Can. J. Phys, **79**, 1323 (2001).

[29] K. Gregor, Ph.D. thesis, Princeton University (2006).

[30] M. P. Zinkin, M. J. Harris, and T. Zeiske, "Short-range magnetic order in the frustrated pyrochlore antiferromagnet $CsNiCrF_6$," Phys. Rev. B, **56**, 11786–11790 (1997).

[31] G.-W. Chern, C. J. Fennie, and O. Tchernyshyov, "Broken parity and a chiral ground state in the frustrated magnet $CdCr_2O_4$," Phys. Rev. B, **74**, 060405 (2006).

[32] G.-W. Chern, R. Moessner, and O. Tchernyshyov, "Partial order from disorder in a classical pyrochlore antiferromagnet," Phys. Rev. B, **78**, 144418 (2008).

[33] B. Canals and D. Garanin, "Classical spin liquid properties of the infinite-component spin vector model on a fully frustrated two dimensional lattice," The European Physical Journal B - Condensed Matter and Complex Systems, **26**, 439–447 (2002).

[34] R. Moessner and A. J. Berlinsky, "Magnetic susceptibility of diluted pyrochlore and $SrCr_{9-9x}Ga_{3+9x}O_{19}$ antiferromagnets," Phys. Rev. Lett., **83**, 3293–3296 (1999).

[35] P. G. de Gennes, "Inelastic magnetic scattering of neutrons at high temperatures," J. Phys. Chem. Solids, **4**, 223–226 (1958).

[36] M. De Leener and P. Résibois, "Irreversibility in Heisenberg Spin Systems. II. Approximate Solution of the High-Temperature Kinetic Equations," Phys. Rev., **152**, 318–325 (1966).

[37] A. Bunker, K. Chen, and D. P. Landau, "Critical dynamics of the body-centered-cubic classical Heisenberg antiferromagnet," Phys. Rev. B, **54**, 9259–9266 (1996).

[38] L. D. C. Jaubert and P. C. W. Holdsworth, "Signature of magnetic monopole and Dirac string dynamics in spin ice," Nature Physics, **5**, 258–261 (2009).

[39] P. H. Conlon and J. T. Chalker, "Spin dynamics in pyrochlore Heisenberg antiferromagnets," Phys. Rev. Lett., **102**, 237206 (2009).

[40] P. C. Hohenberg and B. I. Halperin, "Theory of dynamic critical phenomena," Rev. Mod. Phys., **49**, 435–479 (1977).

[41] P. M. Chaikin and T. C. Lubensky, *Principles of condensed matter physics* (Cambridge University Press, 1995).

[42] C. Moore, M. G. Nordahl, N. Minar, and C. R. Shalizi, "Vortex dynamics and entropic forces in antiferromagnets and antiferromagnetic Potts models," Phys. Rev. E, **60**, 5344–5351 (1999).

[43] S. V. Isakov, R. Moessner, and S. L. Sondhi, "Why spin ice obeys the ice rules," Phys. Rev. Lett., **95**, 217201 (2005).

[44] C. Castelnovo, R. Moessner, and S. L. Sondhi, "Magnetic monopoles in spin ice," Nature, **451**, 42–45 (2006).

[45] D. J. P. Morris, D. A. Tennant, S. A. Grigera, B. Klemke, C. Castelnovo, R. Moessner, C. Czternasty, M. Meissner, K. C. Rule, J.-U. Hoffmann, K. Kiefer, S. Gerischer, D. Slobinsky, and R. S. Perry, "Dirac Strings and Magnetic Monopoles in the Spin Ice $Dy_2Ti_2O_7$," Science, **326**, 411–414 (2009).

[46] S. T. Bramwell, S. R. Giblin, S. Calder, R. Aldus, D. Prabhakaran, and T. Fennell, "Measurement of the charge and current of magnetic monopoles in spin ice," Nature, 956–959 (2009).

[47] 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 $Ho_2Ti_2O_7$," Science, **326**, 415–417 (2009).

[48] J. S. Gardner, M. J. P. Gingras, and J. E. Greedan, "Magnetic pyrochlore oxides," Rev. Mod. Phys., **82**, 53 (2010).

[49] J. van Duijn, N. Hur, J. W. Taylor, Y. Qiu, Q. Z. Huang, S.-W. Cheong, C. Broholm, and T. G. Perring, "From cooperative paramagnetism to Néel order in $Y_2Ru_2O_7$: Neutron scattering measurements," Phys. Rev. B, **77**, 020405 (2008).

[50] K. Kamazawa, S. Park, S.-H. Lee, T. J. Sato, and Y. Tsunoda, "Dissociation of spin objects in geometrically frustrated $CdFe_2O_4$," Phys. Rev. B, **70**, 024418 (2004).

[51] S.-H. Lee, H. Takagi, D. Louca, M. Matsuda, S. Ji, H. Ueda, Y. Ueda, T. Katsufuji, J.-H. Chung, S. Park, S.-W. Cheong, and C. Broholm, "Frustrated magnetism and cooperative phase transitions in spinels," Journal of the Physical Society of Japan, **79**, 011004 (2010).

[52] S. Lee, C. Broholm, W. Ratcliff, G. Gasparovic, Q. Huang, T. Kim, and S. Cheong, "Emergent excitations in a geometrically frustrated magnet," Nature, **418**, 856–858 (2002).

[53] K. Tomiyasu, H. Suzuki, M. Toki, S. Itoh, M. Matsuura, N. Aso, and K. Yamada, "Molecular spin resonance in the geometrically frustrated magnet $MgCr_2O_4$ by inelastic neutron scattering," Phys. Rev. Lett., **101**, 177401 (2008).

[54] J.-H. Chung, M. Matsuda, S.-H. Lee, K. Kakurai, H. Ueda, T. J. Sato, H. Takagi, K.-P. Hong, and S. Park, "Statics and dynamics of incommensurate spin order in a geometrically frustrated antiferromagnet $CdCr_2O_4$," Phys. Rev. Lett., **95**, 247204 (2005).

[55] T. Yavors'kii, T. Fennell, M. J. P. Gingras, and S. T. Bramwell, "$Dy_2Ti_2O_7$ Spin Ice: A Test Case for Emergent Clusters in a Frustrated Magnet," Phys. Rev. Lett., **101**, 037204 (2008).

[56] S. E. Palmer and J. T. Chalker, "Order induced by dipolar interactions in a geometrically frustrated antiferromagnet," Phys. Rev. B, **62**, 488–492 (2000).

[57] A. S. Wills, M. E. Zhitomirsky, B. Canals, J. Sanchez, P. Bonville, P. D. de Reotier, and A. Yaouanc, "Magnetic ordering in $Gd_2Sn_2O_7$ : the archetypal Heisenberg pyrochlore antiferromagnet," Journal of Physics: Condensed Matter, **18**, L37 (2006).

[58] A. N. Yaresko, "Electronic band structure and exchange coupling constants in $ACr_2X_4$ spinels (A = Zn, Cd, Hg; X = O, S, Se)," Phys. Rev. B, **77**, 115106 (2008).

[59] C. Cheng, "Long-range antiferromagnetic interactions in $ZnFe_2O_4$ and $CdFe_2O_4$: Density functional theory calculations," Phys. Rev. B, **78**, 132403 (2008).

[60] K. Kamazawa, Y. Tsunoda, H. Kadowaki, and K. Kohn, "Magnetic neutron scattering measurements on a single crystal of frustrated $ZnFe_2O_4$," Phys. Rev. B, **68**, 024412 (2003).

[61] Y. Yamada, K. Kamazawa, and Y. Tsunoda, "Interspin interactions in $ZnFe_2O_4$: Theoretical analysis of neutron scattering study," Phys. Rev. B, **66**, 064401 (2002).

[62] R. Ballou, E. Lelièvre-Berna, and B. Fåk, "Spin Fluctuations in $(Y_{0.97}Sc_{0.03})Mn_2$: A Geometrically Frustrated, Nearly Antiferromagnetic, Itinerant Electron System," Phys. Rev. Lett., **76**, 2125–2128 (1996).

[63] F. Alet, G. Misguich, V. Pasquier, R. Moessner, and J. L. Jacobsen, "Unconventional continuous phase transition in a three-dimensional dimer model," Phys. Rev. Lett., **97**, 030403 (2006).

[64] G. Misguich, V. Pasquier, and F. Alet, "Correlations and order parameter at a Coulomb-crystal phase transition in a three-dimensional dimer model," Phys. Rev. B, **78**, 100402 (2008).

[65] F. Alet, Y. Ikhlef, J. L. Jacobsen, G. Misguich, and V. Pasquier, "Classical dimers with aligning interactions on the square lattice," Phys. Rev. E, **74**, 041124 (2006).

[66] U. Hizi and C. L. Henley, "Effective Hamiltonian for the pyrochlore antiferromagnet: Semiclassical derivation and degeneracy," Phys. Rev. B, **73**, 054403 (2006).

[67] D. Tsuneishi, M. Ioki, and H. Kawamura, "Novel ordering of the pyrochlore Heisenberg antiferromagnet with the ferromagnetic next-nearest-neighbour interaction," J. Phys.: Condens. Matter, **19**, 145273 (2007).

[68] 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).

[69] R. W. Youngblood and J. D. Axe, "Polarization fluctuations in ferroelectric models," Phys. Rev. B, **23**, 232–238 (1981).

[70] T. E. Saunders and J. T. Chalker, "Structural phase transitions in geometrically frustrated antiferromagnets," Phys. Rev. B, **77**, 214438 (2008).

[71] V. N. Kotov, M. Elhajal, M. E. Zhitomirsky, and F. Mila, "Dzyaloshinsky-Moriya-induced order in the spin-liquid phase of the $S = 1/2$ pyrochlore antiferromagnet," Phys. Rev. B, **72**, 014421 (2005).

[72] 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).

[73] C. L. Henley and N.-g. Zhang, "Semiclassical eigenstates of four-sublattice antiferromagnets," Phys. Rev. Lett., **81**, 5221–5224 (1998).

[74] O. Tchernyshyov, R. Moessner, and S. L. Sondhi, "Order by distortion and string modes in pyrochlore antiferromagnets," Phys. Rev. Lett., **88**, 067203 (2002).

[75] B. Canals, "From the square lattice to the checkerboard lattice: Spin-wave and large-$n$ limit analysis," Phys. Rev. B, **65**, 184408 (2002).

[76] M. E. J. Newman and G. T. Barkema, *Monte Carlo Methods in Statistical Physics* (Oxford University Press, 1999).

[77] M. Galassi, J. Davies, J. Theiler, B. Gough, G. Jungman, P. Alken, M. Booth, and F. Rossi, *GNU Scientific Library Reference Manual*, 3rd ed. (2009).

[78] R. Moessner, Ph.D. thesis, University of Oxford (1997).

[79] M. Frigo and S. Johnson, "The design and implementation of fftw3," Proceedings of the IEEE, **93**, 216 –231 (2005).