TOPICAL REVIEW

Monte Carlo studies of the dipolar spin ice model

Roger G Melko[1,2] and Michel J P Gingras[1,3]

[1] Department of Physics, University of Waterloo, Waterloo, ON, N2L 3G1, Canada
[2] Department of Physics, University of California, Santa Barbara, CA 93106, USA
[3] Canadian Institute for Advanced Research, 180 Dundas Street West, Toronto, ON, M5G 1Z8, Canada

Abstract
We present a detailed overview of numerical Monte Carlo studies of the dipolar
spin ice model, which has been shown to be an excellent quantitative descriptor
of the Ising pyrochlore materials $Dy_2Ti_2O_7$ and $Ho_2Ti_2O_7$. We show that
the dipolar spin ice model can reproduce an effective *quasi*-macroscopically
degenerate ground state and spin ice behaviour of these materials when the
long range nature of dipole–dipole interaction is handled carefully using Ewald
summation techniques. This degeneracy is, however, ultimately lifted at low
temperature. The long range ordered state is identified via Monte Carlo
simulation techniques. Finally, we investigate the behaviour of the dipolar
spin ice model in an applied magnetic field and compare our predictions to
experimental results. We find that a number of different long range ordered
ground states are favoured by the model, depending on field direction.

**Contents**

1. Introduction

1.1. Water ice and spin ice

Frustrated or competing interactions are a common feature of many condensed matter systems [1]. In magnetic materials, frustration arises when the system cannot minimize its total classical ground state energy by minimizing the energy of each spin–spin interaction individually [2–5]. When competing interactions cannot be simultaneously satisfied as a consequence of the arrangement of spins on a geometrical unit, such as a triangle or a tetrahedron, a system made of an assembly of such units is said to be geometrically frustrated. Geometric frustration has been studied extensively in recent years, with the discovery of classical systems that do not display any ordering or dynamical phase transitions down to the lowest temperatures (for recent reviews see [6–12]). Furthermore, much current research effort is being deployed to investigate the exotic behaviour of quantum frustrated systems [13–16]. In highly frustrated systems, weak quantum fluctuations may work to select a unique ground state that is not stabilized at the classical level, while strong quantum fluctuations (e.g. small spin number value, S) can give rise to novel quantum disordered states [17]. Real material [18–20] and model systems with strongly correlated electrons in the presence of strong magnetic frustration display interesting exotic properties.

While geometric frustration most commonly arises between spins interacting antiferromagnetically (AF), Harris and collaborators [21, 22] showed that the pyrochlore lattice of corner sharing tetrahedra with Ising spins pointing along a local cubic ⟨111⟩ axis constitutes a new class of geometrical frustration when nearest neighbour interactions are ferromagnetic (FM) (see figure 1) [23, 24][4,5]. As a consequence of the frustration on this lattice, the Ising pyrochlore ferromagnet has a lowest energy ground state configuration that is

[4] Anderson had already pointed out the analogy of the relationship between the Ising antiferromagnet model on the spinel lattice (which is the same magnetic lattice as the pyrochlore lattice of corner sharing tetrahedra) with that of ice water.

[5] We write the nearest neighbour ferromagnetic model as $H = -J \sum_{\langle ia,jb \rangle} S_i^a \cdot S_j^b$, where $J > 0$ is ferromagnetic, i, j denote FCC lattice points and a, b label the sublattices. For Ising spins that point along a local ⟨111⟩ direction we write $S_i^a = \sigma_i^a \hat{z}^a$, with $\sigma_i^a = \pm 1$ and $\hat{z}^a$ defining the local ⟨111⟩ axis on the a sublattice, so we have $H = (+J/3) \sum_{\langle ia,jb \rangle} \sigma_i^a \sigma_j^b$ since $\hat{z}^a \cdot \hat{z}^b = -1/3$ on the pyrochlore lattice. This is the Hamiltonian of a pyrochlore antiferromagnet with global (fictitious) $\hat{z}$ quantization axis direction and with nearest neighbour coupling constant $+J/3$. Hence, the local ⟨111⟩ Ising ferromagnet on the pyrochlore lattice maps onto the global Ising antiferromagnet on the same lattice, as studied by Anderson in [23] (see footnote 4). For antiferromagnetic exchange, $J < 0$, the same argument shows that there are two ground states related by a global spin reversal symmetry $S_i \rightarrow -S_j$ and which correspond to all spins in or all spins out on the tetrahedra unit cells, therefore constituting a simple and non-frustrated problem (i.e. a global Ising ferromagnet on an FCC lattice), first noticed in [24].
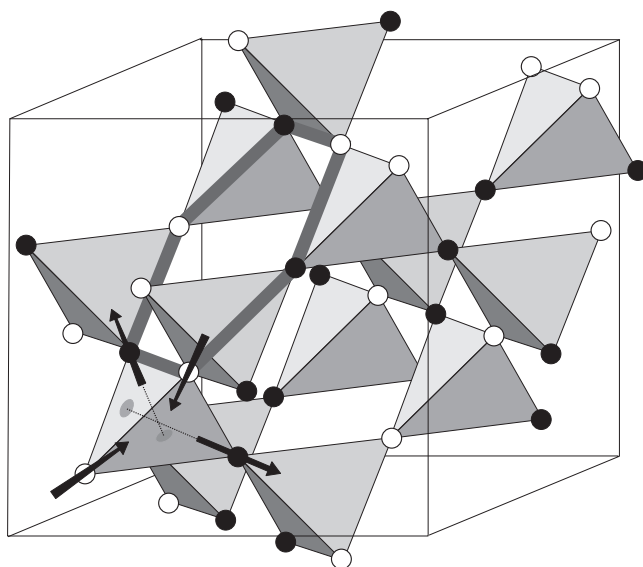
Figure 1. The $\langle 111 \rangle$ Ising pyrochlore lattice. The lower left 'downward' tetrahedron of the pyrochlore lattice shows Ising spins as arrows. Each spin axis is along the local $\langle 111 \rangle$ quantization axis, which goes from one site to the middle of the opposing triangular face (as shown by the discs) and meets with the three other $\langle 111 \rangle$ axes in the middle of the tetrahedron. For clarity, black and white circles on the lattice points denote other spins. White represents a spin pointing into a downward tetrahedron while black has the opposite meaning. The entire lattice is shown in an ice rule state (two black and two white sites for every tetrahedron). The hexagon (thick grey line) shows a minimal size loop move, which corresponds to reversing all colours (spins) on the loop to produce a new ice rule state.

very closely analogous to an entirely different, yet very common frustrated condensed matter system—namely water ice [10, 12]. In the low temperature–low pressure phase of water ice (the so-called 'hexagonal ice', phase I_h), the oxygen atoms are arranged on a hexagonal lattice, each oxygen having four nearest neighbours. Bernal and Fowler [25] and Pauling [26] were the first to propose that the hydrogen atoms (protons) within the H_2O lattice are not arranged periodically, but are disordered. These hydrogen atoms on the O–O bonds are not positioned at the mid-point between the two oxygen atoms, but rather each proton is (covalently) bonded 'near' one oxygen and (hydrogen bonded) 'far' from the other such that the water solid consists of hydrogen bonded H_2O molecules (see figure 2). In the Pauling model, ice I_h is established when the whole system is arranged according to the two *ice rules*:

 (i) Precisely one hydrogen atom is on each proton bond that links two nearest neighbour oxygen atoms.
 (ii) Precisely two hydrogen atoms are near each oxygen atom (spin in) and two are far from them (spin out; see figure 2).

A consequence of this structure, and the subsequent ice rules, is that there is no single unique lowest energy state. Indeed, there exist an infinitely large number of degenerate low energy states that fulfil the ice rules and, if the degeneracy was truly exact, would manifest themselves as a residual entropy at zero temperature (called the zero-point entropy). This set of all configurations that obey the ice rules and contribute to the degeneracy is called the 'ice rules manifold'. Pauling [26, 27] estimated theoretically the residual entropy per hydrogen atom as $S \approx k_B/2 \ln(3/2)$, where k_B is Boltzmann's constant. To make connection with experimental
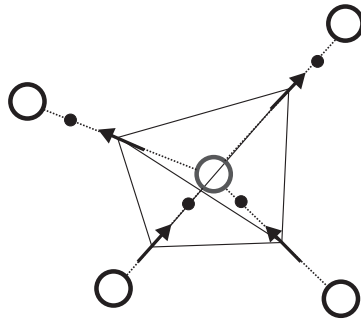
Figure 2. The local proton arrangement in ice, showing oxygen atoms (large white circles) and hydrogen atoms (small black circles) arranged to obey the ice rules. The displacements of the hydrogen atoms from the mid-points of the oxygen–oxygen bonds are represented as arrows, which translate into spins on the pyrochlore lattice in figure 1.

measurements, we will often quote this value scaled by Avogadro's number N_A, so that the residual Pauling entropy per mole is

$$S(T \to 0) \approx \frac{R}{2} \ln \frac{3}{2} = 1.686 \, \text{J mol}^{-1} \, \text{K}^{-1}, \tag{1}$$

where $R = N_A k_B$ is the molar gas constant. For ice water, Pauling's result is not exact, but is accurate to within a few per cent compared to Giauque's experiments [28].

Returning to the magnetic Ising pyrochlores, the analogy to water ice arises if the spins are chosen to represent hydrogen displacements from the mid-points of the O–O bonds (figure 2). The Bernal–Fowler ice rule of *two protons close, two protons further away* corresponds to the configuration with *two spins in and two spins out* of each tetrahedron on the pyrochlore lattice. Because of this direct analogy between water ice and the Ising pyrochlores, Harris *et al* [21, 22] called the latter *spin ice* [10, 12, 23, 24] (see footnotes 4, 5). We note, however, that common water ice at atmospheric pressure, ice I_h, has a hexagonal structure, while here the magnetic lattice has cubic symmetry. Strictly speaking, the Ising pyrochlore problem is equivalent to cubic ice, and not the hexagonal phase. Yet, this does not qualitatively modify the 'ice rule' analogy (or mapping) or the connection between the statistical mechanics of the local proton coordination in water ice and the low temperature spin structure of the spin ice materials.

An important point must be emphasized here. In both ice water and spin ice, the microscopic origin of the residual zero-point entropy arises from the 'simplicity' and 'underconstraints' in the problem. Indeed, the constraints (rules) for constructing a minimum energy ground state, which arise from the underlying microscopic Hamiltonian, are so 'simple' that an infinite number of configurations of the dynamical variables at stake (proton position in ice and spin direction in spin ice) can be used to make a minimum energy state from which the extensive residual ground state entropy $S(T \to 0)$ results. Henceforth, we define the term 'spin ice behaviour' as the observation of this Pauling-type entropy, retained by the material or model spin system at low temperatures, arising from an underlying degeneracy or near degeneracy in the spin configurations that make up an ice rules manifold. Spin ice behaviour also precludes the development of long range order in the system, down to the temperature where dynamical freezing sets in due to large energy barriers separating the ice rule obeying states.
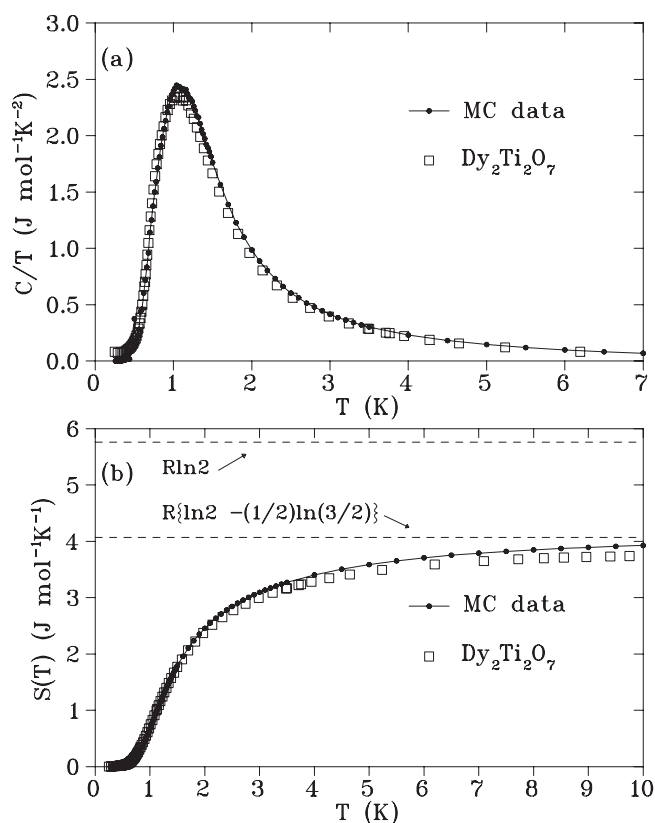
Figure 3. (a) Specific heat and (b) entropy data for $Dy_2Ti_2O_7$ from [30], compared with Monte Carlo simulation results for the dipolar spin ice model, with $J_{nn} = -1.24$ K and $D_{nn} = 2.35$ K.

1.2. Dipolar spin ice

Experimentally, it is known that the single ion ground states of the rare earth ions Dy^{3+} and Ho^{3+} in the pyrochlore structure are described by an effective classical Ising doublet [21, 29]. Specific heat measurements by Ramirez [30] on the compound $Dy_2Ti_2O_7$ have shown that the 'missing' magnetic entropy not recovered upon warming the system from $T \approx 0.4$ to 10 K agrees reasonably well with Pauling's entropy calculation above, $S \approx S(T \to 0)$, thereby providing compelling thermodynamic evidence that $Dy_2Ti_2O_7$ is a spin ice material [31] (see figure 3). While early neutron scattering and magnetization measurements first suggested that $Ho_2Ti_2O_7$ was a spin ice material [21], some subsequent specific heat measurements and numerical simulations by Siddharthan and co-workers were interpreted as evidence for a freezing transition to a partially ordered state as opposed to spin ice behaviour in that material [12, 32, 33]. However, more recent specific heat [34, 35], magnetization [35, 36] and neutron scattering experiments [34], supported by Monte Carlo (MC) simulations [34], confirm the initial proposal [21] that $Ho_2Ti_2O_7$ is indeed a spin ice material akin to $Dy_2Ti_2O_7$. Other magnetization measurements have recently been reported that also argue for spin ice behaviour in the closely related $Ho_2Sn_2O_7$ [37, 38] and $Dy_2Sn_2O_7$ [39] materials. The dynamical properties of these materials at the spin ice freezing point appear somewhat puzzling and are the subject of an increasing number of studies [41–44].

Following the initial spin ice proposal in 1997 by Harris and co-workers [21, 22], it appeared that the spin ice materials obeyed the simple ferromagnetic nearest neighbour model mentioned above. This model intuitively gives rise to a degenerate spin ice ground state because of the equivalent energies of the six different tetrahedron configurations that make up the ground state of this geometrically frustrated unit. However, the nearest neighbour spin ice model is too simple to accurately describe the physical properties of real materials composed of the rare earth ions Ho^{3+} and Dy^{3+} (see [32]). Firstly, the magnetic cations Ho^{3+} and Dy^{3+} in $Ho_2Ti_2O_7$ and $Dy_2Ti_2O_7$ carry a large magnetic moment [21, 29], μ, of approximately $10\ \mu_B$. This entails strong magnetic dipole–dipole interactions in these materials. Indeed, the strength of the dipolar interaction at nearest neighbour distances, D_{nn}, is of order 2 K, which is of the same order of magnitude as the overall magnetic interaction energy scale in these materials as estimated by the Curie–Weiss temperature, $\theta_{CW} \sim 1$ K, extracted from DC magnetization measurements. Secondly, rare earth ions possess very small exchange energies, which are roughly of the same order of magnitude as θ_{CW} and D_{nn}. Consequently, dipole–dipole interactions in $Ho_2M_2O_7$ and $Dy_2M_2O_7$ (M = Ti, Sn) are very significant and constitute an order one energy scale in the problem. This is the reverse of what is observed in transition metal compounds, where the exchange interaction predominates and the dipolar interaction can be treated as a very weak perturbation. Finally, the nearest neighbour exchange interaction in $Ho_2Ti_2O_7$ and $Dy_2Ti_2O_7$ is actually antiferromagnetic, which would by itself cause a phase transition to a Néel long range ordered $\mathbf{q} = 0$ state [22, 24] (footnote 5; see figure 5). Consequently, we consider the simplest model of $\langle 111 \rangle$ Ising pyrochlore magnets with both nearest neighbour exchange and long range magnetic dipole–dipole interactions with the Hamiltonian

$$H = -J \sum_{\langle (i,a),(j,b) \rangle} \mathbf{S}_i^a \cdot \mathbf{S}_j^b + D r_{nn}^3 \sum_{\substack{i>j \\ a,b}} \frac{\mathbf{S}_i^a \cdot \mathbf{S}_j^b}{|\mathbf{R}_{ij}^{ab}|^3} - \frac{3(\mathbf{S}_i^a \cdot \mathbf{R}_{ij}^{ab})(\mathbf{S}_j^b \cdot \mathbf{R}_{ij}^{ab})}{|\mathbf{R}_{ij}^{ab}|^5}. \quad (2)$$

Here the spin vector $\mathbf{S}_i^a = \sigma_i^a \hat{z}^a$ labels the Ising moment of magnitude $|\mathbf{S}_i^a| = 1$ at FCC lattice site $\mathbf{R}_i$ and tetrahedral sublattice site coordinate $\mathbf{r}^a$, where the local Ising axis is denoted by $\hat{z}^a$ and the Ising variable is $\sigma_i^a = \pm 1$. The vector $\mathbf{R}_{ij}^{ab} = \mathbf{R}_{ij} + \mathbf{r}^{ab}$ connects spins $\mathbf{S}_i^a$ and $\mathbf{S}_j^b$. J represents the exchange energy and D the dipolar energy scale ($J > 0$ and $D = 0$ in the spin ice model originally proposed by Harris *et al* [22] which we refer to as the 'near neighbour spin ice model'). Because of the relative local $\langle 111 \rangle$ Ising orientations, $\hat{z}^a \cdot \hat{z}^b = -1/3$, the nearest neighbour exchange energy between two spins is $J_{nn} \equiv J/3$. The dipole interaction is calculated from

$$D = \frac{\mu_0}{4\pi} \frac{\mu^2}{r_{nn}^3}. \quad (3)$$

Experimentally, from magnetization measurements [22] and analysis of the crystal field levels via inelastic neutron scattering [29], it is known that the moments of the Dy^{3+} and Ho^{3+} rare earth ions in the pyrochlore lattice are $\mu \approx 10\ \mu_B$ and the nearest neighbour distance r_{nn} is approximately 3.54 Å. We get the dipole–dipole interaction at nearest neighbour distances to be $D_{nn} \equiv 5D/3$, since $\hat{z}^a \cdot \hat{z}^b = -1/3$ and $(\hat{z}^a \cdot \mathbf{R}_{ij}^{ab})(\mathbf{R}_{ij}^{ab} \cdot \hat{z}^b) = -2/3$ in equation (2). For both $Ho_2Ti_2O_7$ and $Dy_2Ti_2O_7$, $D_{nn} \approx 2.35$ K.

In order to consider the combined role of exchange and dipole–dipole interactions, it is useful to define an effective nearest neighbour energy scale, J_{eff}, for $\langle 111 \rangle$ Ising spins:

$$J_{eff} \equiv J_{nn} + D_{nn}, \quad (4)$$

where $J_{nn} \equiv J/3$ is the nearest neighbour exchange energy between $\langle 111 \rangle$ Ising moments. This simple near neighbour description of the system suggests that a $\langle 111 \rangle$ Ising system could

display spin ice properties, even for antiferromagnetic nearest neighbour exchange, $J_{nn} < 0$, as long as $J_{eff} = J_{nn} + D_{nn} > 0$. Fits to experimental data give $J_{nn} \sim -0.52$ K for $Ho_2Ti_2O_7$ [34] and $J_{nn} \sim -1.24$ K for $Dy_2Ti_2O_7$ [45]. Thus, J_{eff} is positive (using $D_{nn} = 2.35$ K), and hence ferromagnetic and frustrated, for both $Ho_2Ti_2O_7$ ($J_{eff} \sim 1.8$ K) and $Dy_2Ti_2O_7$ ($J_{eff} \sim 1.1$ K). It would therefore appear natural to ascribe the spin ice behaviour in both $Ho_2Ti_2O_7$ and $Dy_2Ti_2O_7$ to the positive J_{eff} value as in the simple model of Bramwell and Harris [22]. This naive interpretation goes a long way in explaining the origin of spin ice in $Ho_2Ti_2O_7$ and $Dy_2Ti_2O_7$. However, the situation is more complex than it appears.

Dipole–dipole interactions (equation (2), second term) are 'complicated': (i) they are strongly anisotropic since they couple the spin, $\mathbf{S}_i^a$, and space, $\mathbf{R}_{ij}^{ab}$, directions; and (ii) they are also very long ranged ($\propto |\mathbf{R}_{ij}^{ab}|^{-3}$). For example, the second nearest neighbour distance is $\sqrt{3}$ times larger than the nearest neighbour distance, which means that the second nearest neighbour dipolar energy is $D_{nnn} \sim 0.2 D_{nn}$. This implies an important perturbation compared to $J_{eff} = J_{nn} + D_{nn} < D_{nn}$, especially for antiferromagnetic (negative) J_{nn}. Specifically, for $Dy_2Ti_2O_7$, the second nearest neighbour energy scale is about 40% of the effective nearest neighbour energy scale, J_{eff}, a large proportion! Therefore, one might have expected that the dipolar interactions beyond nearest neighbour would cause the different ice rule states to have different energies, hence possibly breaking the degeneracy of the spin ice manifold, similar to what happens in the kagomé [46] and pyrochlore Heisenberg antiferromagnets [47] when exchange interactions beyond nearest neighbour are considered. In equation (2), if the dipolar term is summed beyond nearest neighbours, one might naively expect a long ranged Néel ordered state at a critical temperature $T_N \sim O(D_{nn})$. Thus, here arises one of the main puzzling and interesting problems posed by the dipolar spin ice materials that can be summarized by two questions:

(i) Are the experimental observations of spin ice behaviour in real materials consistent with dominant long range dipolar interactions?
(ii) If so, why do long range dipolar interactions fail to destroy spin ice behaviour and give rise to long range order at a temperature $T_N \sim O(D_{nn})$?

Results from Monte Carlo simulations on the dipolar spin ice model attempting to answer the first question above were first reported in [32] and [33]. In that work, the dipole–dipole interactions were cut off at a distance of five [32] or ten and twelve nearest neighbours [33]. In those studies the thermodynamic behaviour was found to be consistent with spin ice behaviour for a model of $Dy_2Ti_2O_7$, provided that the size of the magnetic moment of Dy^{3+} was rescaled and the exchange interaction was made to extend far beyond nearest neighbours [48][6], but not for a model of $Ho_2Ti_2O_7$. Another work [45] considered the Hamiltonian of equation (2) with only nearest neighbour exchange and the value of J as an adjustable parameter. In that work, the long range dipole–dipole interaction was handled using the Ewald method, which derives an effective dipole–dipole interaction between spins within the cubic simulation cell. The Monte Carlo simulations were carried out by slowly cooling the simulated lattice, subject to the usual Metropolis algorithm. Numerical integration of the specific heat divided by temperature was performed to determine the entropy of the system [45]. For a parameter J appropriate for the $Dy_2Ti_2O_7$ spin ice material (see section 2 below), the dipolar spin ice model retained Pauling's entropy (equation (1)), in good agreement with experiments on $Dy_2Ti_2O_7$ (figure 3). Following the same approach as in [45], Monte Carlo simulations have found good agreement between the dipolar spin ice model, specific heat measurements and elastic neutron

---

[6] However, we note that for the open pyrochlore lattice structure, one would naively expect further neighbour exchange coupling to be very small. See the discussion in [48].

scattering, as well as experiments on $Ho_2Ti_2O_7$ [34]. Finally, mean field theory calculations of the neutron scattering intensity, valid in the (paramagnetic) temperature regime $T \gg \theta_{CW}$, that consider a large distance cut-off of the dipole–dipole interactions, have been found to be in good agreement with experiments on $Ho_2Sn_2O_7$ [38] and $Ho_2Ti_2O_7$ [49]. Consequently, there is now strong compelling evidence that the long range dipolar interaction is responsible for the ice behaviour and the subsequent retention of spin entropy in rare earth based insulating pyrochlore magnets down to very low temperature [10].

1.3. True long range order at low temperature in the dipolar spin ice model

Having answered question No 1 above in the affirmative, one is then faced with addressing question No 2. The Monte Carlo results mentioned above [45] show that spin ice behaviour arises from the combination of nearest neighbour exchange, J_{nn}, and dipole energies, D_{nn}, which create an *effective* ferromagnetic Ising model J_{eff} at nearest neighbours as long as $J_{eff} = J_{nn} + D_{nn} > 0$, akin to Harris and Bramwell's simple nearest neighbour Ising model [21, 22, 50, 51]. However, the long range dipolar interaction does not appear to destroy the spin ice degeneracy (and subsequent retention of zero-point entropy) created by this effective ferromagnetic nearest neighbour interaction. In support of this picture, a mean field theory (MFT) calculation finds that the remaining (beyond nearest neighbour) dipole–dipole interaction terms, which couple every spin in the system with varying strength depending on their separation distance, are 'self-screened' to a large degree [52, 53]. This means that the degeneracy between different ice rule obeying states is almost exactly fulfilled by carefully including the long distance dependence of the dipolar term in the Hamiltonian. However, and perhaps most interestingly for these Ising pyrochlore systems, the same mean field calculation suggests that the screening of the long range terms is not perfect and that the associated spin ice manifold is only *quasi*-degenerate, due to some small remaining effective energy scale, so a unique ordering wavevector is selected [49, 52, 53]. This suggests that at some temperature below the onset temperature of spin ice correlations, the dipolar spin ice model should in principle favour the unique long range ordered state selected by this remaining ('unscreened') perturbative dipole–dipole energy.

One might naively expect that such an ordered state should be found in the MC simulations [34, 45]. However, this does not happen, as measurements of the temperature dependent acceptance rate of the simulations make it apparent that the standard single (Ising) spin flip Metropolis algorithm experiences a dynamical 'freezing' at a temperature of ≈ 0.4 K for J_{nn} and D_{nn} parameters appropriate for $Dy_2Ti_2O_7$ [45] and $T \approx 0.6$ K for $Ho_2Ti_2O_7$ [34]. If the dipolar interactions are cut off at some arbitrary distance, R_c, one can generate scenarios where, depending on specific numerical values for J_{nn}, D_{nn} and R_c, a selected state is dynamically accessible before the spin ice manifold freezes out, as was found in simulations where dipole interactions are cut off [32, 33] (see also appendix B). Consequently, akin to the case for the approaches used in ice lattice (vertex) models [54, 55], one must introduce non-local dynamics in the simulation to combat this freezing out and maintain simulation equilibrium down to lower temperatures. The inclusion of non-local 'loop moves' in the dipolar spin ice model promotes the development of a long range ordered phase via a sharp first order phase transition at $T \approx 0.18$ K [56, 57], a much lower temperature than the onset temperature for spin ice correlations at $T \sim 1.2$ K in $Dy_2Ti_2O_7$ [30, 45] and $T \sim 1.9$ K in $Ho_2Ti_2O_7$ [34]. The ground state found in the loop MC simulations has zero total (bulk) magnetization (recall that each tetrahedron individually carries a net magnetic moment in each of the ice rule obeying states). See figure 4 for the spin configurations in this ground state. The pre-transitional specific heat and the latent heat associated with the first order transition
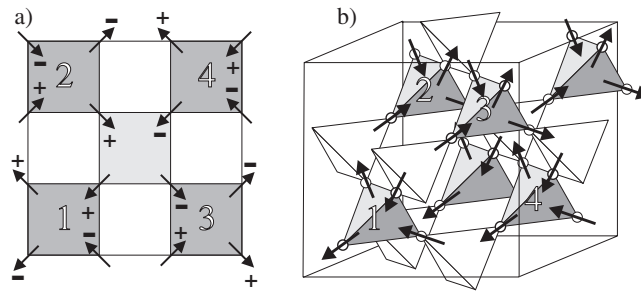
Figure 4. The long range ordered $\mathbf{q} = (0, 0, 2\pi/a)$ dipolar spin ice ground state. Projected down the z axis (a), the four tetrahedra making up the cubic unit cell appear as dark grey squares. The light grey square in the middle does not represent a tetrahedron, but its diagonally opposing spins are in the same lattice plane. The component of each spin parallel to the z axis is indicated by + and − signs. In perspective (b), the four tetrahedra of the unit cell are numbered to enable comparison with (a).
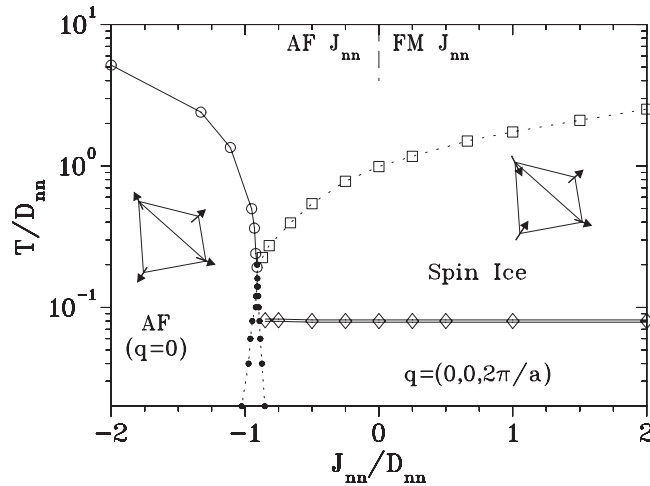
Figure 5. The phase diagram for the dipolar spin ice model in zero applied magnetic field. The antiferromagnetic ground state is an all-spins-in or all-spins-out configuration for each tetrahedron. The spin ice configuration, which includes the $\mathbf{q} = (0, 0, 2\pi/a)$ ground state, is a two-spins-in–two-spins-out configuration for each tetrahedron. The region encompassed between the quasi-vertical dotted curves displays hysteresis in the long range ordered state selected ($\mathbf{q} = 0$ versus $\mathbf{q} = (0, 0, 2\pi/a)$) as J_{nn}/D_{nn} is varied at fixed temperature T.

recover all of Pauling's missing entropy in the model. The ordered state that is found in the loop MC simulations [56] corresponds to the ordered state predicted by mean field theory [52, 53]. In other words, the dipolar spin ice model possesses on its own, without invoking energetic perturbations and/or thermal and quantum fluctuations, a unique (up to trivial global symmetry relations) classical ground state with zero entropy in the thermodynamic limit.

Pauling's entropy can also be recovered in spin ices exposed to an external magnetic field. In particular, for fields of sufficient magnitude aligned along different crystal axes, distinct ground state ordering patterns can be realized which destroy the macroscopic spin ice degeneracy through various mechanisms. Using MC simulations and direct Ewald energy calculations, we are able to investigate the behaviour of the dipolar spin ice model in an external magnetic field. With application of a large field along three different crystal symmetry

directions, three different long range ordered ground states appear. With large fields parallel to the [100] crystal direction, the ground state is the ice rules $\mathbf{q} = 0$ structure identified by Harris [21]. For large fields parallel to [110], the ground state is the ice rule $\mathbf{q} = X$ state [21] and for large fields along [111], the ice rules are broken and a spin configuration wit three spins in, one spin out becomes the lowest energy state. The experimentally determined field dependence of the magnetization and specific heat for fields along the [100], [110] and [111] directions in the $Dy_2Ti_2O_7$ spin ice material agree quantitatively with the Monte Carlo results for the long range dipolar spin ice model [58–63].

1.4. Phases of dipolar spin ice

Using Monte Carlo simulations, the zero-field phase diagram for the dipolar spin ice model can be mapped out (figure 5). To summarize our results, spin ice correlations develop for all cases where the effective nearest neighbour energy scale $J_{eff}/D_{nn} > 0.095$ (ferromagnetic) and the temperature is below the broad peak in the specific heat, T_{peak}. For $T/D_{nn} \leqslant 0.08$, independently of the value of J_{nn} (as long as $J_{eff}/D_{nn} > 0.095$), the system orders into the long range ordered state of figure 4, with the help of the loop moves in the simulation. For J_{nn}/D_{nn} less than -0.905 ($J_{eff}/D_{nn} < 0.095$), the system orders into an antiferromagnetic $\mathbf{q} = 0$ Néel ground state, where every tetrahedron in the system has an all-in or all-out spin configuration at low temperatures [22, 24, 64, 65] (see footnote 5). The region around $J_{nn}/D_{nn} = -1$ shows hysteresis at low temperatures. Because of the close cancellation of energy scales, we imagine that real materials which fall into this region, e.g., $Tb_2Ti_2O_7$ [66–69], will be particularly susceptible to the influence of small perturbations (such as exchange beyond nearest neighbour or finite, as opposed to infinite, Ising anisotropy [53, 70] and, in particular, to quantum fluctuations) with the result of possible ordering into long range ordered states [71] distinct from the two shown in figure 5.

1.5. Outline

The rest of the paper is organized as follows. In the next section we present results from conventional Monte Carlo simulations with single spin flips, that show how spin ice behaviour develops at finite temperatures in the dipolar spin ice system whenever the effective nearest neighbour coupling is ferromagnetic ($J_{eff} = J_{nn} + D_{nn} > 0$). Results from mean field theory, presented elsewhere [52, 53], show that there exists a weak selection of a unique ordering mode at $\mathbf{q} = (0, 0, 2\pi/a)$. Motivated by these mean field results, we undertake a numerical search for a long range ordered state in the model of equation (2). Section 3 discusses the details of a loop Monte Carlo algorithm that avoids the freezing phenomenon observed in a MC simulation employing local single spin flip dynamics. Section 4 presents the detailed results from the loop Monte Carlo simulations. The results for the field dependence of the ground state energy and magnetization for fields along [100], [110] and [111] are presented in section 5. We conclude the paper with a brief discussion in section 6. We have included three appendices. Appendix A contains a discussion of the Ewald technique for dipolar interactions in real space, as employed in MC simulations. Appendix B compares the Ewald method to previous results using truncated dipolar interactions. Appendix C discusses some of the effects of a finite demagnetization factor on simulation results.

2. The dipolar spin ice model: the conventional metropolis Monte Carlo approach

In this section we present the results of Monte Carlo simulations of the dipolar spin ice Hamiltonian using a standard single spin flip Metropolis algorithm. All simulations are carried

out on pyrochlore lattices (figure 1) with periodic boundary conditions and of various sizes, which we reference via the linear dimension L. Precisely, L measures the edge of a cubic unit cell of four tetrahedra (figure 4(a)) containing sixteen Ising spins, so the number of spins in each simulation cell is $N = 16 \times L^3$. Using the Metropolis algorithm, we mimic the experimental conditions pertinent to real materials by cooling the simulation sample slowly through a number of temperature steps. At each temperature step, the system is equilibrated carefully, then thermodynamic quantities of interest are calculated in the Monte Carlo 'production' run. Equilibration is typically confirmed by comparing results on several simulation runs of different lengths. In figures in this paper, equilibrium and production times were sufficiently large to give statistical error bars that are smaller than the plotted symbol sizes, unless explicitly illustrated.

In the Hamiltonian equation (2), the dipole interaction strength D can be determined once the crystal field structure of the magnetic ion is known, leaving the nearest neighbour exchange J as the only adjustable parameter in the model. However, the dipole–dipole interaction must be handled with care. A lattice summation of such interactions is conditionally convergent due to its $1/R^3$ nature. In order to properly treat the long range nature of this term, we implement the Ewald method in the simulations, which derives an effective dipole–dipole interaction between spins within the simulation cell [72]. Unlike in dipolar fluid simulations [73], the pyrochlore lattice constrains the positions of the spins in the simulation, allowing the Ewald interactions to be calculated only once, after which a numerical simulation can proceed as normal. Appendix A contains a brief discussion of the Ewald method applied to MC real space simulations.

Simulations carried out on the dipolar spin ice model using the single spin flip Metropolis algorithm are able to map out three different regions (paramagnetic, AF and spin ice) of the phase diagram shown in figure 5. As we shall discuss below, the loop Monte Carlo simulation allows one to explore the development of a long range ordered spin ice state, the $\mathbf{q} = (0, 0, 2\pi/a)$ phase, at $T/D_{nn} \sim 0.08$ (the double horizontal lines in figure 5). Thermodynamic data indicate that when the nearest neighbour exchange is AF and sufficiently large compared to the dipolar interactions ($J_{nn} < 0$ and $|J_{nn}| \gg D_{nn}$), the system undergoes a second order phase transition (in the three-dimensional Ising universality class) to an all-in or all-out $\mathbf{q} = 0$ ground state.

2.1. Magnetic specific heat

We focus now on Monte Carlo data for the magnetic specific heat of the dipolar spin ice model. As illustrated in figure 6, specific heat data for different J_{nn} in the spin ice regime ($J_{eff} = J_{nn} + D_{nn} > 0$) show qualitatively the same broad peak as observed in the nearest neighbour FM exchange model [50], which vanishes at high and low temperatures. The height, C_{peak}, and peak temperature position of this peak, T_{peak}, show very little dependence on system size, for simulation cells of $L = 2, 3, 4, 5$ and 6. However, both C_{peak} and T_{peak} are found to depend strongly on the ratio J_{nn}/D_{nn}, as illustrated in figure 7.

By rescaling the temperature scale for the specific heat corresponding to a number of different interaction parameters J_{nn} and D_{nn}, one can expose more clearly the dependence of the specific heat on the competition between the nearest neighbour exchange J_{nn} and the dipole–dipole interactions. Figure 6 shows that in terms of an effective energy scale, J_{eff}, the medium range to long range effects of the dipolar interactions are in some sense 'screened' by the system, and one recovers qualitatively the short range physics of an effective nearest neighbour spin ice model. As the nearest neighbour exchange interaction becomes AF (see figure 7 for $J_{nn}/D_{nn} < 0$), we find that the approximate collapse onto a single energy scale
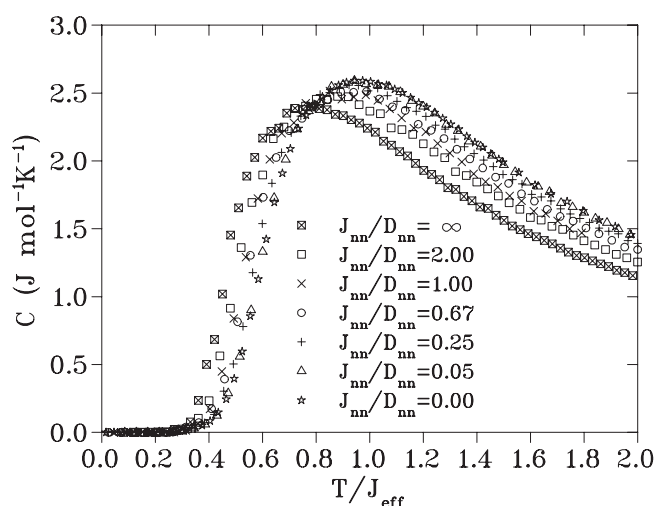
Figure 6. The specific heat for system size $L = 2$, with temperature, T, rescaled into units of the effective nearest neighbour interaction $J_{eff} = J_{nn} + D_{nn}$. $J_{nn}/D_{nn} = 0$ corresponds to purely dipolar interactions, while $J_{nn}/D_{nn} = \infty$ corresponds to nearest neighbour FM exchange only. Simulation runs for $L = 4$ were also performed, but revealed no important finite size effects.
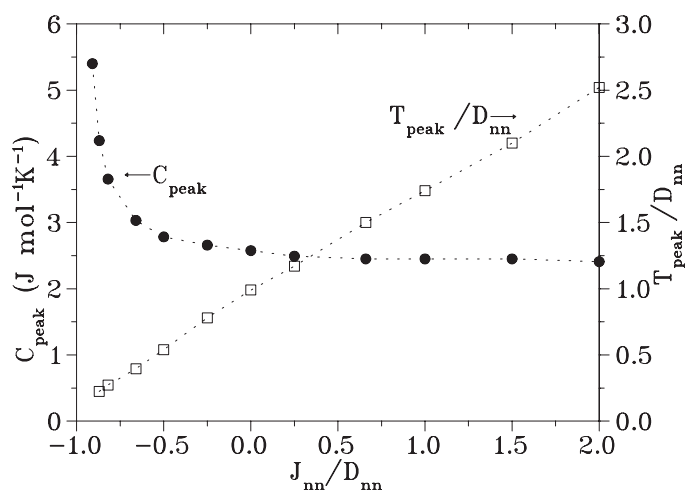
Figure 7. The dependence of the simulated specific heat peak height C_{peak} and temperature location of C_{peak} and T_{peak} on exchange and dipole–dipole interaction parameters. In this figure D_{nn} is set to 2.35 K.

becomes less accurate, with the specific heat becoming dependent on J_{nn}/D_{nn}. It is within this regime that we believe that both $Ho_2Ti_2O_7$ and $Dy_2Ti_2O_7$ are realized, as we now discuss.

Since D_{nn} is calculated from equation (3), J_{nn} must be determined from experimental data. By fitting either the height C_{peak} or the peak temperature T_{peak} of the maximum of the specific heat curves of the Monte Carlo simulation to the experimental results [30] (figure 3(a)), we obtain a value of $J_{nn} = -1.24$ K for $Dy_2Ti_2O_7$ [45]. The results of this fitting are illustrated in the top panel of figure 3. A fitting of the height or peak temperature of the experimental magnetic contribution to the specific heat for $Ho_2Ti_2O_7$ gives $J_{nn} = -0.52$ K for this material [34]. Contrary to what is reported in [33] we therefore conclude that $Ho_2Ti_2O_7$ is
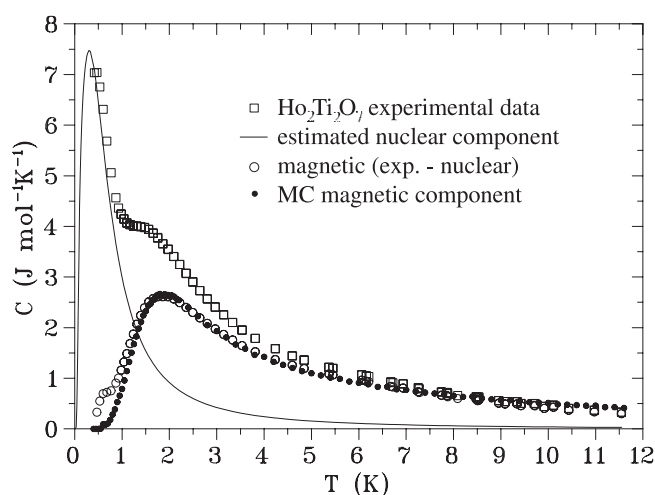
Figure 8. The total experimental specific heat of $Ho_2Ti_2O_7$ is shown by the open squares. The expected nuclear contribution is indicated by the curve, while the resulting magnetic specific heat estimation is shown by the open circles. Near 0.7 K the estimation is prone to a large error. Dipolar spin ice simulation results are indicated by the filled circles.

'deeper' (J_{eff} more positive for $Ho_2Ti_2O_7$ than for $Dy_2Ti_2O_7$) in the spin ice regime (further to the right in figure 5) than $Dy_2Ti_2O_7$. As initially reported in [74], the temperature dependence of the specific heat for $Ho_2Ti_2O_7$ is less straightforward to interpret than for $Dy_2Ti_2O_7$ [30]. In $Ho_2Ti_2O_7$ the specific heat possesses an important contribution from a nuclear component due to a large hyperfine splitting of the nuclear levels well known to occur for Ho^{3+} cations, as discussed in [75] and [76]. This nuclear component was estimated by Blöte *et al* [75] for Ho_2GaSbO_7, another pyrochlore material. By subtracting it off from the (total) experimental specific heat, we can uncover the underlying magnetic contribution and compare to the theoretically calculated Monte Carlo specific heat data, from which T_{peak} or C_{peak} can be determined directly (figure 8). We note here, as recently observed in [77], that for $Dy_2Ti_2O_7$ there should be a hyperfine nuclear contribution to the specific heat manifesting itself at a temperature below $T \sim 0.4$ K in the data of Ramirez *et al* [30] if one uses the typical hyperfine contact interaction expected for a Dy^{3+} insulating salt. The absence of the high temperature $1/T^2$ tail of the nuclear specific heat (on the descending low temperature side of the magnetic specific heat of $Dy_2Ti_2O_7$) below 0.4 K in figure 3(a) is, therefore, somewhat puzzling[7].

The shoulder-like feature in the estimated magnetic contribution to the experimental specific heat data of figure 8 (open circles) near 0.7 K can be entirely eliminated by adjusting the nuclear hyperfine splitting by $\sim$2% around the value estimated by Blöte for Ho_2GaSbO_7, resulting in an exceedingly good agreement with the Monte Carlo results down to $T = 0.4$ K. Such a slight adjustment to account for any small deviations in the hyperfine parameters of 4f rare earth ions (dependent upon electric field gradients, chemical shift etc) would seem reasonable. However, we do not do this, in order to emphasize that the unbiased use of the estimated nuclear specific heat contribution from the isostructural pyrochlore material Ho_2GaSbO_7 [75] already allows a very good agreement with the theoretical magnetic specific heat for $T \geqslant 0.7$ K.

---

[7] We note here as well that, compared to the results of [30], the specific heat measurements reported by Blöte *et al* [75] for $Dy_2Ti_2O_7$ do not suffer this problem: the total specific heat of $Dy_2Ti_2O_7$ in [75] never falls below the expected $1/T^2$ paramagnetic nuclear specific heat contribution.

Having determined J_{nn} and D_{nn} for $Ho_2Ti_2O_7$ from specific heat measurements, we are able to compare the experimental elastic neutron scattering against that determined via the Monte Carlo simulations. The results, reported in [34], show excellent agreement between experiment and simulation. More recent neutron scattering experiments on $Ho_2Sn_2O_7$ show similar results [38]. Such comparison between theory and experiments for $Dy_2Ti_2O_7$ is more difficult due to the large neutron absorption cross section of naturally occurring Dy isotopes. Work in that direction using isotopically enriched samples with the ^{162}Dy isotope is in progress [78, 79].

2.2. Pauling's entropy

Numerical integration of the specific heat divided by temperature can be performed to determine the entropy of the dipolar spin ice model. Specifically, the entropy difference between temperature T_1 and T_2, $S(T_2) - S(T_1)$, can be calculated using the thermodynamic relation

$$S(T_2) - S(T_1) = \int_{T_1}^{T_2} \frac{C(T)}{T} \, dT. \tag{5}$$

The results for $Dy_2Ti_2O_7$ are illustrated in figure 3(b). The entropy recovered from $T = 0.4$ K, where $C(T)$ is very small, up to a temperature $T = 10$ K is $S(T = 10 \text{ K}) - S(T \approx 0) \approx 3.930$ J mol^{-1} K^{-1}. As we can see in figure 3(b), the Monte Carlo data for $S(T)$ at $T = 10$ K are slightly below Pauling's value $R\{\ln(2) - (1/2)\ln(3/2)\}$. To perform the calculation of the recovered entropy from $T = 10$ K up to ∞, we extrapolate the specific heat $C(T)$ for $T > 10$ K by matching the Monte Carlo value of $C(T)$ at $T = 10$ K with the $1/T^2$ high temperature paramagnetic temperature regime, $C(T) = C_\infty/T^2$ for $T > 10$ K. This gives a value $C_\infty = 29.015$ J mol^{-1} K and an extra entropy of $S(T = \infty) - S(T = 10) = 0.145$ J mol^{-1} K^{-1}, and hence a value $S(T = \infty) - S(T \approx 0) = 4.075$ J mol^{-1} K^{-1}, in exceedingly close agreement with Pauling's value, 4.077 J mol^{-1} K^{-1}. Hence, we find that the simulation with the appropriate experimental parameters retains Pauling's entropy (equation (1)), similar to what is found experimentally for $Dy_2Ti_2O_7$ (figure 3(b) and in [30]). A similar experimental procedure was carried out using the magnetic contribution of the specific heat data of $Ho_2Ti_2O_7$, also giving a residual entropy close to Pauling's entropy [35].

While the above conventional Monte Carlo simulations of the model Hamiltonian for the spin ice compounds, equation (2), yields a reasonably successful quantitative theory of spin ice behaviour in Ising pyrochlore materials, there still remains the second question (No 2, section 1.2) as to why dipolar interactions, despite their anisotropic and long range nature, do not (appear to) lift the macroscopic degeneracy associated with the ice rules and select an ordered state. As explained in the next section, in the case of the simulations, the answer lies in the implementation of local single spin flip dynamics in the Metropolis Monte Carlo algorithm.

3. Dynamical freezing and loop moves in Monte Carlo simulations

The first attempt to address the question of long range order in dipolar spin ice was an investigation of the soft modes (i.e. critical modes or ordering wavevectors) accessible to the model within the context of mean field theory [52, 53]. This mean field theory suggests that, from the point of view of a strictly equilibrium (statistical mechanics) magnetic ordering phenomenon, the dipolar spin ice model of equation (2) should be a rather conventional system with a single, well defined ordering wavevector, $\mathbf{q} = (0, 0, 2\pi/a)$, and staggered magnetization order parameter. If this were the case, a simple zero-temperature ground state permanent moment structure, where $\langle \mathbf{S}_i \rangle = 1$ at each site, could be constructed using solely
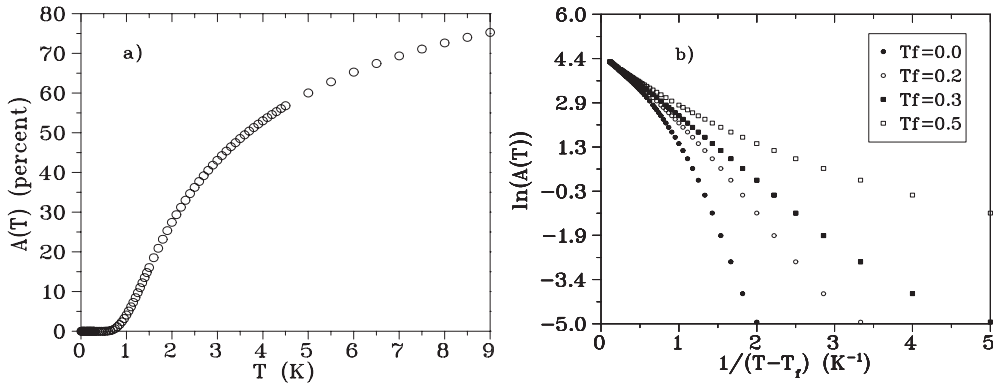
Figure 9. (a) The single spin flip Monte Carlo step acceptance rate $A(T)$ for a simulation of $Dy_2Ti_2O_7$. The simulation becomes frozen when the acceptance rate falls to zero. (b) The logarithm of the acceptance rate plotted versus inverse temperature, minus some freezing temperature T_f, follows a Vogel–Fulcher type of law.

this critical ordering wavevector **q** and critical modes (i.e. eigenvectors of the Fourier transform of the exchange plus dipole interaction matrix). The question, addressed in this section, then becomes: why do Monte Carlo simulations of the dipolar spin ice model, or in fact the real spin ice materials themselves, not develop the long range ordered phase predicted by mean field theory?

3.1. Dynamical freezing in conventional single spin flip Monte Carlo simulations

The problem of dynamical freezing in the above simulation of the dipolar spin ice model turns out to lie in the local single spin flip dynamics employed within the Metropolis algorithm and, similarly, the local spin dynamics at play in the real materials. As we will show below, Monte Carlo simulations of the dipolar spin ice model using single spin flips experience a dynamical freezing at low temperatures. This arises due to the existence of large energy barriers separating distinct quasi-degenerate spin ice configurations, which prevent the simulations (and presumably the real materials) from finding their true energetically favoured long range ordered ground states (see [80] for a related problem).

Observation of the acceptance rate $A(T)$ (percentage of accepted Monte Carlo steps) of the dipolar spin ice simulations makes it immediately apparent that out-of-equilibrium freezing occurs at low temperatures, that is below $T \sim 0.4$ K (as illustrated in figure 9(a)) for the J_{nn} and D_{nn} parameters appropriate for $Dy_2Ti_2O_7$. Figure 9(b) shows that $A(T)$ can be parametrized by a Vogel–Fulcher temperature dependence as found in numerous freezing phenomena: $A(T) \propto \exp(\Delta/(T - T_f))$, where the freezing temperature, T_f, is introduced in an ad hoc fashion. In figure 9(b), $\Delta = 1$ K.

It is clear that in order to investigate the existence of a true energetically favoured ground state in the dipolar spin ice model, a standard Monte Carlo simulation employing local single spin flip dynamics is inadequate. Indeed, as figure 9 shows, these dynamical processes are frozen out at T just slightly below 0.4 K for a model of $Dy_2Ti_2O_7$. For J_{nn} and D_{nn} appropriate to describe $Ho_2Ti_2O_7$ [34], the single spin flip Monte Carlo acceptance rate falls below 10^{-6} at a temperature near 0.6 K. Without ascribing any deep significance to it, it is interesting to note that this freezing out in the simulation at 0.4 and 0.6 K corresponds rather closely to the temperatures at which freezing is found in $Dy_2Ti_2O_7$ [42] and $Ho_2Ti_2O_7$ [41], respectively.

This freezing out occurs due to large free energy barriers separating the (almost) degenerate ice rule states, which develop rapidly at a temperature of order T_{peak} (where the specific heat shows a maximum), and which are associated with introducing a single spin flip to a tetrahedron obeying the ice rules. As discussed above, and further supported by mean field calculations [52, 53], the effective (ferromagnetic) nearest neighbour interaction J_{eff} favours the ice rule configuration. As the temperature drops, the Boltzmann weight $\exp(-4J_{eff}/T)$ suppresses the probability that a single spin flip will take a given tetrahedron into an intermediate, thermally activated configuration not obeying the ice rules. Thus single spin flip Monte Carlo moves are, for all practical purposes, frozen out and dynamically eliminated within the simulation when $T \ll J_{eff}$.

3.2. Loop moves in Monte Carlo simulations

In order to explore the low temperature ordering properties of dipolar spin ice, one needs a Monte Carlo algorithm with non-local updates that effectively bypass the energy barriers that separate nearly degenerate states and allows the simulation to explore the restricted ice rule phase space that prevents ordering in the model [56, 57]. In other words, we employ non-local dynamical processes to restore ergodicity in the Monte Carlo simulation, and then use this new algorithm to explore and characterize the long range ordered state that arises out of the ice rules manifold and which is energetically favoured by the dipolar spin ice model.

We first identify the true zero-energy modes that can take the near neighbour spin ice model from one ice state to another exactly energetically degenerate ice state. An example of these zero modes, or loops, is shown in figure 1. We take as an initial working hypothesis that in the dipolar spin ice model, with interactions beyond nearest neighbour, the system freezes into an ice rule obeying state. This is indeed what we found: in all of the tests performed on the dipolar spin ice model, simulations using conventional single spin flips always froze out into an ice rule obeying state with no 'defects' (by defects we mean violations of the Bernal–Fowler ice rules). With interactions beyond nearest neighbour, these loop moves become quasi-zero modes that can take the dipolar spin ice model from one ice rule state to another without introducing spin defects into tetrahedra in the lattice. This allows all of the quasi-degenerate spin ice states to be sampled ergodically, and facilitates the development of a long range ordered state by the system at low temperatures.

Within the Monte Carlo simulation, we use the Barkema and Newman [54, 55] loop algorithm originally designed for two-dimensional square ice models, and adapt it to work in a similar manner on the three-dimensional pyrochlore lattice. In the context of square ice, we tested two kinds of loop algorithms, the so-called *long* and *short* loop moves. In the square ice model, each vertex on a square lattice has four spins associated with it [57]. The vertices are analogous to tetrahedron centres in the pyrochlore lattice. The ice rules correspond to 'two spins pointing in, two spins pointing out' at each vertex. In the Newman and Barkema algorithm, a loop is formed by tracing a path through ice rule vertices, alternating between spins pointing into and spins pointing out of the vertices. A 'long loop' is completed when the path traced by the loop closes upon the spin from which it started. A 'short loop' is formed whenever the path traced by the loop encounters any other vertex (tetrahedron) already included in the loop—excluding the dangling tail of spins (figure 10).

We now generalize the Barkema and Newman loop algorithm for our study of the three-dimensional pyrochlore lattice spin ice problem. In this system, the smallest complete loop that is a zero mode on the pyrochlore lattice consists of six spins (see figure 1). Such a loop for spin ice was previously identified by Bramwell and Harris [22] and also by Anderson [23] (see footnote 4) in the context of the spinel lattice. However, using the above loop algorithm, much
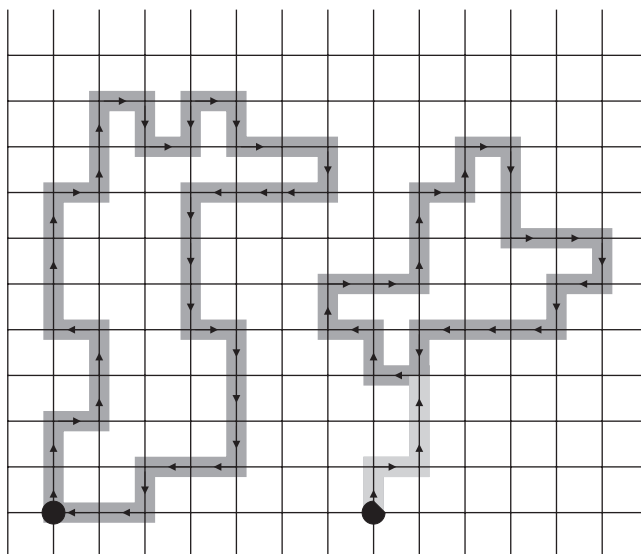
Figure 10. Long and short loops formed by the Newman and Barkema algorithm [54, 55] on a square ice lattice. Vertices are represented by points where lattice lines cross. Each vertex has two spins pointing in and two spins pointing out; however, for clarity, only spins which are included in the loops are shown. Starting vertices are indicated by large black dots. On the left is an example of a long loop, which is completed when it encounters its own starting vertex. On the right is a short loop, which is complete when it crosses itself at any point. Dark grey lines outline completed loops. The excluded tail of the short loop is shown in light grey.

larger loops are possible. When used with the pyrochlore lattice (figure 1), such a loop must pass through two spins on each tetrahedron. A loop always 'enters' a tetrahedron through an inward pointing spin and 'leaves' a tetrahedron through an outward pointing spin. The periodic boundary conditions of the lattice may also be traversed with no adverse consequences. If we form a closed loop in this manner, and each spin is reversed on it, the entire system stays in an ice rule state. However, small dipole–dipole energy gains or losses may be incurred due to configurational differences between the old and the new ice rule state. These small energy changes caused by the loop moves are evaluated via a Metropolis algorithm within the Monte Carlo simulation. Specifically, a loop move that takes the system from one ice rule state to another one of lower energy is automatically accepted, while a loop move that takes the system to a higher energy ice rule state (with energy difference δE between the two states) is accepted with $\exp[-\delta E/(k_B T)] > $ rnd, where rnd is a random number taken from a uniform distribution between 0 and 1 [54, 55].

Before using them in a full-scale Monte Carlo simulation, the long and short loop algorithms were subjected to a variety of characterizing tests on the three-dimensional pyrochlore lattice [57]. The first test is a study of the relative speed (measured by CPU time) of the algorithms for different sized lattices. As reported in [57], it is found that the small loop algorithm creates loops that approach a finite size limit as the system size increases. The long loop algorithm continues creating larger and larger loops, that scale approximately linearly with the number of spins in the simulation cell. This forces the algorithm to become drastically slower for the larger system sizes considered.

Second, tests were carried out to investigate how the two different loop algorithms handle defects that break the ice rules on a tetrahedron. As we know, above the 'spin ice peak' in the
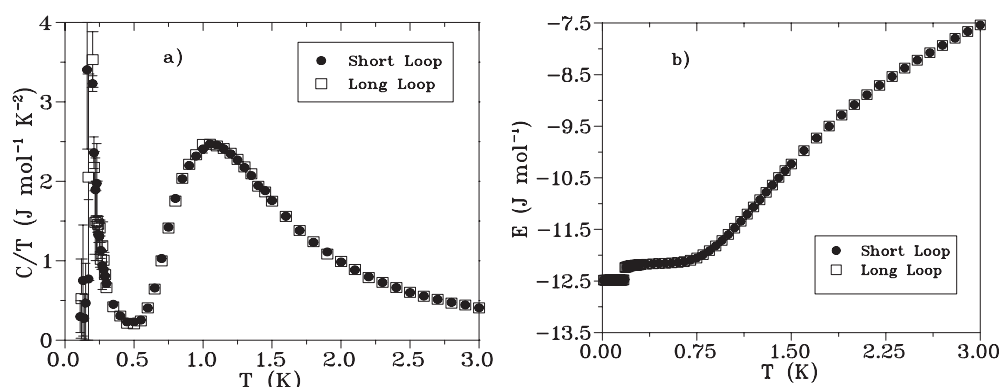
Figure 11. Preliminary data for the low temperature magnetic specific heat (a) and energy (b) of the dipolar spin ice Monte Carlo simulation, system size $L = 3$, with simulation parameters set for $Dy_2Ti_2O_7$. The data represent an average taken over approximately 10^5 production Monte Carlo steps. Closed circles are data from a simulation of the short loop algorithm; open squares are data obtained using the long loop algorithm. Low temperature features are discussed in the next section.

specific heat of the dipolar spin ice model, the ice rules (two spins pointing into a tetrahedron, two spins pointing out) are generally not obeyed. However, to retain detailed balance, the loop algorithm attempts to form loops at temperatures above the onset of spin ice correlations. The attempt to create a loop is simply aborted in the case where the loop path encounters a defect (either a three-in and one-out vertex, or an all-in or all-out vertex). The simulation does not flip any spins on an aborted loop, and the Metropolis algorithm is not employed in this case. As reported in [57], the influence of ice rule defects on loop algorithm performance is significant. Within the long loop algorithm, the inclusion of only one defect spin per one thousand spins in the system causes almost half of all loops which are attempted to be aborted on the grounds that they have encountered the defect, with efficiency decreasing drastically as more defects are included. In contrast, the short loop algorithm remains 87% efficient with the inclusion of one defect in one thousand, and retains an efficiency that is significantly higher than the long loop algorithm as more defects become present in the system [57].

We used both algorithms to perform a true finite temperature Metropolis Monte Carlo simulation of the dipolar spin ice model. The two algorithms were implemented separately, and simulations ran using the regular procedure (cooled slowly from a high temperature, equilibrating carefully at every temperature step). Examples of our preliminary Monte Carlo results for both the short and long loop are given in figure 11. The data in this figure have relatively low statistics, only 10^5 Monte Carlo production steps per spin, where each step consists of a local (single spin) update for each spin and a single loop move. Regardless of this, it is clear that the short and long loops promote roughly the same thermodynamic behaviour in the Monte Carlo simulation. The low temperature features (specific heat peak and energy discontinuity at $T \approx 0.2$ K) of figure 11 are induced in the same manner by both algorithms. Further, additional runs (not illustrated) employing different numbers of equilibrium and production MC steps reproduced the results, albeit with a strong effect on the error bars (figure 11(a)). These features will be discussed in much more detail in section 4. Since the long and short loops display equivalent results in the Monte Carlo approach, we are free to choose between the two algorithms based solely on their performance properties measured above. As alluded to above and detailed in [57], the short loop algorithm works more efficiently within the requirements of our simulation. However, the disadvantage of
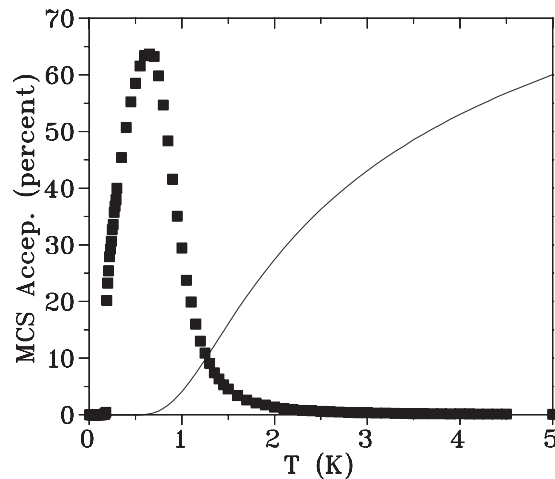
Figure 12. The Monte Carlo acceptance rate for the short loop algorithm (squares) as compared to the single spin flip algorithm (curve), for a simulation using $Dy_2Ti_2O_7$ parameters. Acceptance rates were calculated as percentages of attempted Monte Carlo steps (MCS).

using the short loop algorithm is that each loop does not cover as large of a percentage of spins within the lattice. It is not clear to us, without investigating the computational performances (i.e. autocorrelation times) of both algorithms in much more quantitative detail, whether a small number of long loops is better at bringing the system to equilibrium than a larger number of short loops for fixed available CPU time resources. However, with the additional observation that the long loop algorithm can pass over itself a number of times during its creation, effectively losing additional computational efficiency in this manner, we ultimately choose to perform the majority of the simulations on the dipolar spin ice model using the short loop algorithm.

With the short loop algorithm chosen for the simulations, we re-investigate the Monte Carlo Metropolis algorithm acceptance rate. Since each loop successfully created (i.e., not aborted by encountering an ice rule defect) by the short loop algorithm is still subject to rejection by the Metropolis condition on the basis of its change in system energy, we expect the maximum acceptance rate to be somewhat less than the maximum efficiency of the algorithm given above. Results for the loop acceptance rate are shown in figure 12. Clearly, the loop algorithm becomes effective in the temperature range where the spin ice manifold is well developed, and the single spin flip algorithm loses its ability to explore all possible configurations of the system. Above approximately 1 K, the number of accepted loops is very low, due to the fact that the system is not entirely in an ice rule configuration. As the simulation is slowly cooled, ice rule constraints begin to develop, and the loop algorithm begins to work efficiently, moving the system between different ice rule states. In figure 12, a sharp drop is observed in the loop acceptance rate at approximately 0.18 K. As discussed in the next section, this corresponds to the temperature where a phase transition develops in the system, which locks the system into a long range ordered state.

4. Loop Monte Carlo investigation of the transition to long range order

As suggested by the results above, the short loop algorithm is successful in restoring ergodicity in the simulation. As a consequence of this, we observe a low temperature phase transition in the model. The most familiar and robust indicator of a phase transition is a finite size remnant
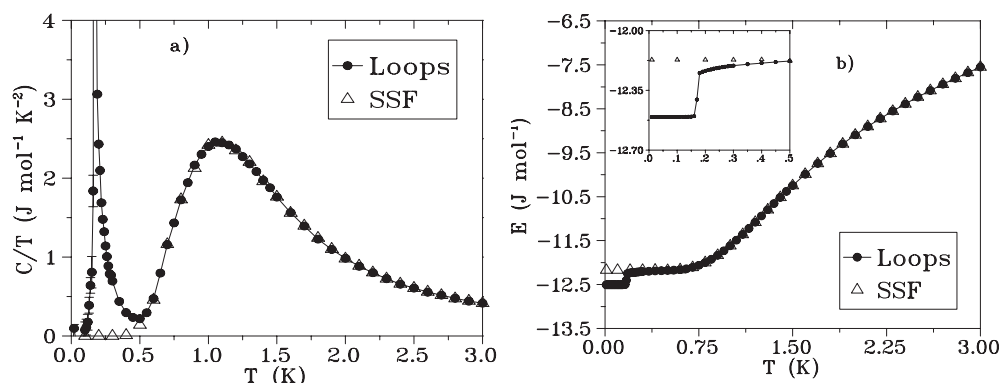
Figure 13. The low temperature magnetic specific heat (a) and energy (b) of the dipolar spin ice Monte Carlo simulation, system size $L = 4$, with simulation parameters set for $Dy_2Ti_2O_7$. Closed circles are simulation data run with the short loop algorithm; open triangles are data obtained using the single spin flip Metropolis algorithm. In the inset of (b), the energy shows an apparent discontinuity at a critical temperature $T_c \approx 0.18$ K. The broad feature in the specific heat at $T \approx 1$ K indicates the rapid development of the spin ice rule obeying states. The sharp feature at T_c is the appearance of a phase transition to a ground state being made (dynamically) accessible via the non-local loop dynamics. Note that these results have higher statistics than those for figure 11, specifically, 1×10^5 equilibration and 3×10^5 production Monte Carlo steps were used. In addition, the system size has been increased from $L = 3$ to 4. Also, note that the location of the specific heat peak is at roughly the same temperature, but is narrower than for $L = 3$, indicating a finite size effect on the singular behaviour of $C(T)$.

of a singularity (divergence or discontinuity) in some thermodynamic quantity at the transition temperature T_c. For example, in the simulation of $Dy_2Ti_2O_7$, a sharp cusp in the specific heat is observed at a temperature below the spin ice peak (see figure 13(a)). The feature in the specific heat and the abrupt drop in energy at the same temperature (figure 13(b)) give good preliminary evidence that the loop algorithm is successful in allowing a phase transition to occur at a temperature of $T_c = 0.18$ K. This is the same temperature as the loop algorithm acceptance rate goes to zero at in figure 12. The energy curve shows a discontinuous drop at T_c (i.e. a latent heat) for large lattice sizes (figure 13(b)), suggesting a first order phase transition. The long range ordered state that develops is observed to be independent of the system size simulated. In the remainder of this section, we attempt to characterize this state and the phase transition that leads to it. The first step is to identify the order parameter associated with the low temperature ordered state.

4.1. The order parameter

Direct inspection of the spin directions at $T < 0.18$ K reveals that the ordered state is a long range ice rule obeying state with zero magnetic moment per unit cell and commensurate with the pyrochlore cubic unit cell (see figure 4). This state corresponds to the critical mode found in the mean field calculation of [52, 53]. It is also the same as the ground state found in a previous Monte Carlo simulation using truncated dipolar interactions [33]. There are twelve symmetrically equivalent spin configurations for the ground state as explained below, two for each cubic axis direction and their spin reversed states. The ordering wavevector $\mathbf{q}$ lies parallel to one of the cubic axis directions, specifically $\mathbf{q} = (0, 0, 2\pi/a)$ or one of its symmetrically equivalent directions. To construct the ordered state, first consider a starting tetrahedron with its six possible ice rule states. For a given ordering wavevector $\mathbf{q}$, this tetrahedron selects one of
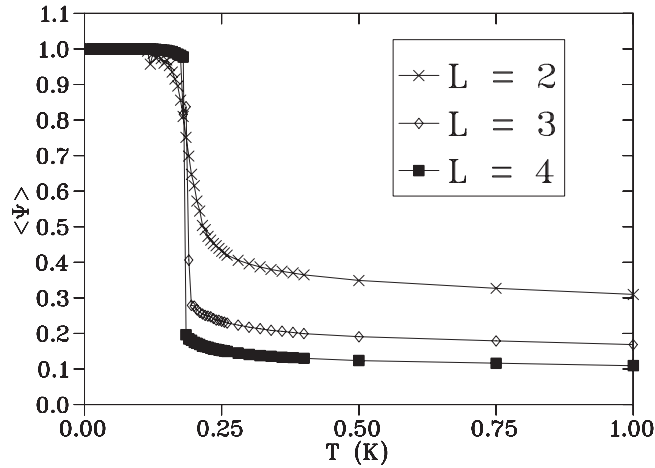
Figure 14. The $\mathbf{q} = (0, 0, 2\pi/a)$ order parameter. Curves are shown for system size $L = 2, 3$ and 4.

the four possible spin configurations (two independent configurations and their spin reversals, $\mathbf{S}_i^a \to -\mathbf{S}_i^a$) with a total magnetic moment for the tetrahedron perpendicular to $\mathbf{q}$. The entire ordered state may then be described by planes (perpendicular to $\mathbf{q}$) of such tetrahedra. The wavelength defined by this $\mathbf{q}$ physically corresponds to antiferromagnetically stacked planes of tetrahedra, which means that a given plane has tetrahedra of reverse spin configuration to the planes above and below it. Hence, we construct the multicomponent order parameter

$$\mathbf{\Psi}_\alpha^m = \frac{1}{N} \left| \sum_{j=1}^{N/4} \sum_{a=1}^{4} \sigma_j^a e^{(i\phi_a^m + i\mathbf{q}_\alpha \cdot \mathbf{R}_j)} \right|. \tag{6}$$

This type of labelling is natural given that the pyrochlore lattice can be viewed as an FCC lattice with a 'downward' tetrahedral basis (figure 1). Thus j labels the FCC lattice points of the pyrochlore lattice and the index a sums over the four spins comprising the basis connected to each j. The index α labels the three possible symmetry related $\mathbf{q}$ ordering wavevectors. For a given $\mathbf{q}_\alpha$, as described above, there are two ice rule configurations and their reversals which can each form a ground state. Thus $m = 1, 2$ labels these possibilities with the phase factors $\{\phi_a^m\}$, describing the given configurations m. Each Ising variable σ_j^a has a value +1 or −1 when a spin points into or out of its downward tetrahedron j, respectively.

As written in equation (6), $\mathbf{\Psi}_\alpha^m$ has six degenerate components, each of which can take on a value between 0 and 1. Upon cooling through the transition, the system selects a unique ordered configuration, causing the corresponding component of $\mathbf{\Psi}_\alpha^m$ to rise to unity and the other five to fall to zero (provided that the finite size system is simulated over a timescale less than the ergodic timescale where full spin symmetry is restored). The component which rises to unity is equally likely to be any one of the six, selected at random through spontaneous symmetry breaking.

Figure 14 is a plot of $\langle \Psi \rangle$ for three system sizes, where

$$\langle \Psi \rangle = \sqrt{\sum_{m=1}^{2} \sum_{\alpha=1}^{3} \left(\mathbf{\Psi}_\alpha^m \right)^2} \tag{7}$$

is the magnitude of the multicomponent order parameter. These three curves illustrate important finite size effects for $\langle \Psi \rangle$. For $T < T_c$ the different lattice sizes produce identical

order parameters. By contrast, $\langle\Psi\rangle$ for the smaller lattice size displays pronounced rounding near T_c and an increased residual value for large T. The larger lattice size produces an order parameter with a clear discontinuity at T_c. This discontinuity in the order parameter combined with the discontinuity of the total energy in figure 13(b) can be viewed as strong preliminary evidence for a first order transition.

We now argue the need to study this phase transition with greater numerical accuracy. This is necessary in order to confirm its first order nature. More importantly, once this is done, we want to use the data to confirm the full recovery of Pauling's entropy through an estimation of the latent heat at the transition.

4.2. Evidence for a first order transition

To begin, we note that there are a number of criteria at one's disposal for demonstrating the occurrence of a first order transition in a Monte Carlo simulation. In particular:

 (i) The order parameter $\langle\Psi\rangle$ should have a clear discontinuity at T_c.
 (ii) The energy probability histogram, $H(E)$, should have a double peak at T_c, which identifies the coexistence of two distinct phases at T_c.
(iii) There should be a latent heat at the transition, identifiable by a discontinuity in the internal energy for large system sizes.
(iv) In the Monte Carlo simulation, the height (maximum) of the specific heat, C_{peak}, and the magnetic susceptibility, χ_{peak}, should be proportional to the simulation volume:

$$C_{\text{peak}}, \chi_{\text{peak}} \propto a + bL^d \tag{8}$$

where a and b are constants and d is the dimension of the lattice, in our case equal to three ($d = 3$).
 (v) The minimum of the fourth order energy cumulant [81],

$$V = 1 - \frac{\langle E^4 \rangle}{3\langle E^2 \rangle^2}, \tag{9}$$

should vary as

$$V_{\min} = V_0 + cL^{-d} \tag{10}$$

where $V_0 \neq 2/3$.
(vi) The temperature $T_{\text{peak}}(L)$ at which C_{peak} or χ_{peak} has a maximum should vary with the simulation volume as

$$T_{\text{peak}}(L) = T_c + cL^{-d} \tag{11}$$

where c is a constant and $T_c = T_{\text{peak}}(L \to \infty)$.

We have already confirmed the first condition of our list. To check for the second condition the energy probability histogram was calculated by binning the simulation energy values for every Monte Carlo step as the system passed through the transition from higher to lower temperatures (figure 15). Above T_c, we observe a single peak Gaussian-like distribution of energies. At T_c, the energy probability distribution shows a double peak, characteristic of the coexistence of two phases found at a first order phase transition. Below T_c we would normally expect to see a Gaussian peak. However, in our case, the histograms below T_c are distorted by the accumulation of energies into the lowest bins due to the proximity of the transition to the ground state. At zero temperature, we expect all of the energies to lie in the bin for the lowest energy.
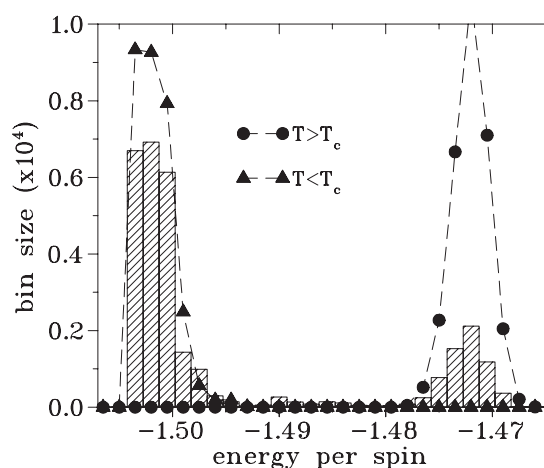
Figure 15. The energy probability distribution histogram for three temperatures: $T = 0.178$ K, $T = 0.180$ K (T_c) and $T = 0.182$ K. For $T > T_c$ (filled circles), a single peaked Gaussian histogram is present. At the transition temperature (hashed rectangles), a second peak appears which has a lower mean energy. As the temperature falls below T_c (filled triangles), the peak with the higher mean energy disappears and the system energy eventually gathers in the lowest bin.
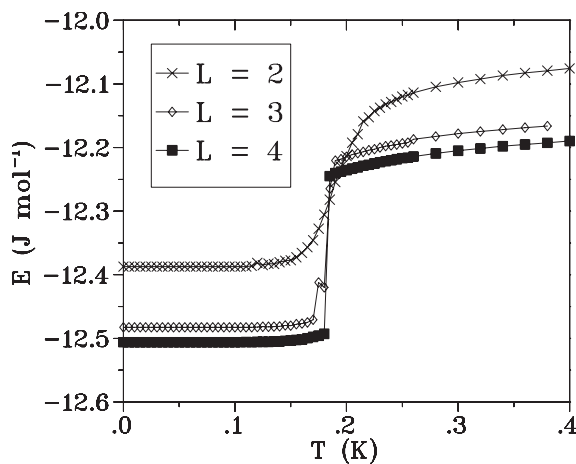
Figure 16. Details of the simulation energy near the transition for different system sizes.

The next condition on our list is the observation of a latent heat at the transition. Figure 16 shows the energy near the transition for three different system sizes. A clear discontinuity develops as we increase the system size. The energy discontinuity, ΔE, as read off from this graph for $L = 4$, is

$$\Delta E \approx 0.248 \text{ J mol}^{-1}. \tag{12}$$

This behaviour is also consistent with the transition being first order. Below, we compare this ΔE to the latent heat released from the entropy jump at the transition calculated through other means. Another calculation of the latent heat at the transition comes from the finite size scaling of condition (iv) above [81]:

$$C_{\text{peak}} \approx \frac{(\Delta E)^2}{4k_B T_c^2} L^d + a \tag{13}$$

where L^d is the system volume as before and a is the intercept of the graph of $C_{\text{peak}}(L)$ versus L to test for condition (iv), which we now discuss.

When attempting to quantify the relationships in conditions (iv)–(vi) on our list, we notice a problem. The extremely sharp nature of the transition makes accurate estimations for these quantities almost impossible using a traditional temperature cooled MC simulation of the hybrid single spin flip–loop algorithm. The reason is that the transition temperature region is narrow, and first order metastability effects are strong, so obtaining accurate data for quantities such as C_{peak} and V_{min} very near to T_c is extremely difficult. As shown in figure 15, the energy probability histogram near a first order transition displays a double hump. The energies that occur between these humps correspond to system configurations that are strongly suppressed by the Boltzmann probability distribution near the transition. We call these 'interface configurations' [82]. Traditional Monte Carlo simulations try to 'avoid' these interface configurations as the system is cooled through the transition, because of their suppression by the Boltzmann factor which is the basis of the Metropolis condition. Therefore, the simulation often behaves poorly in this region, moving quickly through interface configurations to find more favourable configurations nearby in configuration space. This can lead to erratic behaviour and poor statistics in thermodynamic quantities of interest near the phase transition, thereby reducing the numerical accuracy of the quantities used in finite size scaling.

To overcome this problem, Berg and Neuhaus [82] proposed the *multicanonical* method, which is designed to enhance configurations that have energies which occur between the two humps of the probability distribution. If these interface configurations are artificially enhanced, the simulation does not avoid this energy range as strongly and better statistics can be obtained. The version of the multicanonical Monte Carlo algorithm that we use is that proposed by Hansmann and Okamoto [83], originally developed to be used in the context of protein folding simulations. The core of the method is:

Perform Monte Carlo simulations in a multicanonical ensemble instead of the usual canonical ensemble. Then, obtain the relevant canonical distribution by using the histogram reweighting techniques of Ferrenberg and Swendsen [84]. From this, calculate the thermodynamic quantities of interest.

In the multicanonical ensemble, we define the probability distribution by

$$p_{\text{mu}}(E) = \frac{g(E)w_{\text{mu}}(E)}{Z_{\text{mu}}} = \text{constant} \tag{14}$$

where $g(E)$ is the density of states, $w_{\text{mu}}(E)$ is the multicanonical weight factor (not temperature dependent) and Z_{mu} is the associated partition function. The distribution is constant, meaning that all energies have equal weight, which sometimes leads to the name 'flat histogram' method. This flatness is important because it ensures that configurations in the interface region of the transition are not suppressed.

Unlike for the canonical ensemble, the multicanonical weight factor w_{mu} is not *a priori* known. This turns out to be the crucial step of this scheme: finding an accurate estimator of w_{mu} that makes the distribution $p_{\text{mu}}(E)$ flat over the energy range of interest. The details of how to do this are somewhat involved, and will not be explicitly outlined here. The reader is referred to the relevant technical references for details [57, 83]. Our procedure follows that of [83] very closely.

Assuming that we can find a good estimator of w_{mu}, our method proceeds as follows:

(i) We find an accurate estimator of the multicanonical weight factor such that $p_{\text{mu}}(E)$ is reasonably flat over an energy range that includes the transition interface.
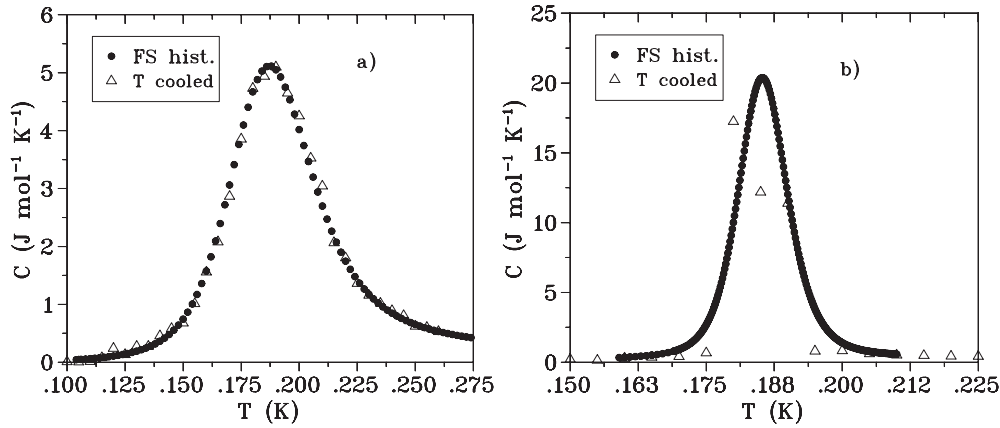
Figure 17. Specific heat curves over the transition temperature, for $L = 2$ (a) and $L = 3$ (b) system sizes. Closed circles represent data obtained with Ferrenberg and Swendsen's (FS) histogram reweighting technique. Open triangles represent data taken using a traditional temperature cooled Monte Carlo simulation, and are plotted without error bars in (b) for clarity.

 (ii) With this weight factor we perform a multicanonical simulation at one given temperature T slightly higher than T_c.

(iii) During this simulation run, we gather statistics for the physical variables of choice (for example, the energy E). These variables are weighted according to the multicanonical distribution.

(iv) From this single simulation, we then obtain the Boltzmann-distributed variables at any temperature for a wide range of temperatures using a reweighting technique.

We use the reweighting technique proposed by Ferrenberg and Swendsen [84], which allows us to transform, or reweight, data obtained from another distribution (in our case the multicanonical distribution) to the relevant Boltzmann distribution, at some inverse temperature β. We use this to obtain an estimate for a given physical quantity in the canonical distribution.

We collect data within the multicanonical distribution and use them to calculate the specific heat for the dipolar spin ice model. For the smallest system size considered, we accurately reproduce the specific heat over the transition using the flat histogram method. Figure 17(a) shows a comparison between the specific heat of an $L = 2$ system obtained using the histogram method at one temperature, and the traditional Monte Carlo procedure with 8×10^5 equilibration steps and 2×10^6 data production steps for every temperature point. The CPU time that it took to get the histogram data was a small fraction of the time it took to obtain the regular Monte Carlo data. Figure 17(b) is a similar result for a larger system size, $L = 3$. The traditional Monte Carlo data were taken with 5×10^5 equilibration steps and 1×10^6 data production steps. The histogram data were obtained in the same amount of CPU time as for the $L = 2$ histogram data, and it was only slightly more difficult to find a good estimate for $w_{mu}(E)$. The poor quality of the traditional Monte Carlo Metropolis data for $L = 3$ stands in stark contrast to the smooth data obtained using the multicanonical simulation.

Unfortunately, one difficulty with the multicanonical algorithm used here is that, in general, as the system size is increased, it becomes increasingly difficult to obtain a good estimate for a $w_{mu}(E)$ that would give a flat $p_{mu}(E)$. The critical temperature, T_c, of the transition seems to be the quantity most sensitive to variations in the flatness of $p_{mu}(E)$. In contrast, the height of specific heat peak is fairly accurately determined for simulation sizes $L = 2, 3, 4$ and 5 (figure 18), showing only a weak sensitivity to the flatness of $p_{mu}(E)$.
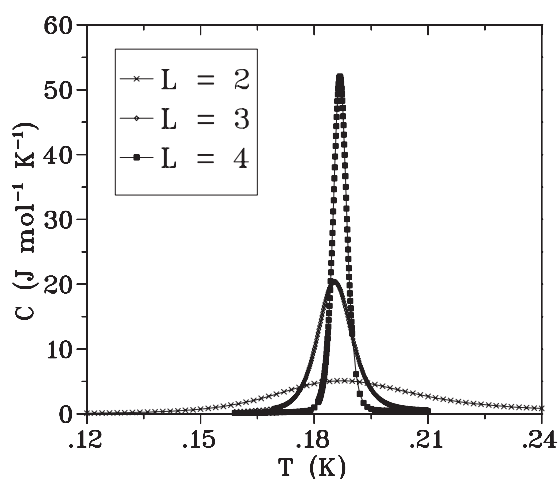
Figure 18. The specific heat of the transition to long range order, for system sizes $L = 2$, 3 and 4. It was found that peak heights for these system sizes did not vary significantly with the flatness of $p_{mu}(E)$. Critical temperatures (T_c) were more sensitive to the flatness of the multicanonical distribution, and hence were more difficult to estimate than C_{peak}.
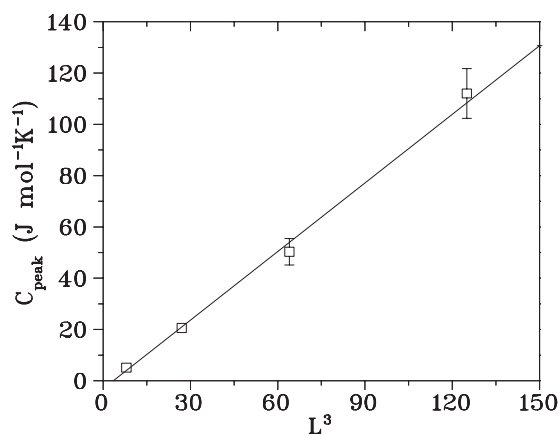
Figure 19. The finite size scaling fit for the specific peak heights of the ordering transition. Data points represent the mean C_{peak} value for a given L. Error bars show one standard deviation.

The aforementioned error associated with T_c for the $L = 4$ peak, as determined from simulations, is of the order of 0.04 K, and becomes increasingly more drastic for the larger system sizes. The variation in the height of the specific heat was found to be much less. Nevertheless, to combat any minor variation in peak height and obtain accurate finite size scaling results, a statistical averaging was done on several ($\sim$10) multicanonical weighting factors to obtain values for C_{peak}. These results are plotted in figure 19. A straight line fit to the data using linear regression gives

$$C_{peak} = 0.8924L^3 - 3.149. \tag{15}$$

The L^3 dependence of $C_{peak}(L)$ shows that the finite size scaling is consistent with that expected for a first order transition. Also, as a second estimator of the latent heat, we can use the slope

of this line and equation (13) to extract ΔE. Doing so we deduce a latent heat of

$$\Delta E \approx 0.245 \text{ J mol}^{-1}, \tag{16}$$

consistent with the value obtained in figure 16 (for the $L = 4$ system) from reading directly off from the energy graph (see equation (12)).

   This completes our study of the nature of the ordering transition in dipolar spin ice. As we have shown, the discontinuity in the order parameter, the release of latent heat, the double peaked energy probability distribution and the finite size scaling of the specific heat peak all give consistent and compelling evidence for the transition being first order. As the technical details concerning this transition are understood, we can proceed to study where it, and the long range ordered state which results from it, stand in our broader picture of ground state entropy found in experiments and in standard single spin flip simulations of the dipolar spin ice model.

4.3. Recovery of Pauling's entropy

Since we have confirmed the first order nature of the transition, the configuration of the ordered state, calculated the latent heat and obtained reliable data for the specific heat through the transition, we are in a position to recalculate the total entropy that the dipolar spin ice model releases as it is cooled to low temperatures. This calculation must be done carefully. We know that in an infinite system, a first order transition is characterized by a cusp in the specific heat. If the transition is temperature driven, as in our case, this first order singularity is the latent heat. For an infinite system going through a first order transition, thermodynamics gives

$$\Delta S = \int_0^{T_c^-} \frac{C_<}{T} \, \mathrm{d}T + \int_{T_c^+}^{\infty} \frac{C_>}{T} \, \mathrm{d}T + \frac{\Delta E}{T}, \tag{17}$$

where $\Delta E / T$ is the latent heat contribution to the entropy (see figure 16), and T_c^- and T_c^+ are the temperature limits asymptotically close to T_c, below and above T_c, respectively (see figure 16).

   To estimate a value for the entropy, we consider the system size $L = 4$ which has good statistical data for the widest temperature range. We integrate the low temperature data for the specific heat in figure 18 divided by temperature obtained from the histogram reweighting technique (up to $T \approx 0.21$). For $T > 0.21$ K we use our regular temperature cooled Monte Carlo data (canonical loop + single spin flip) for the integration above this point, and up to 10 K, giving $S(T = 10) - S(T \approx 0) = 5.530 \text{ J mol}^{-1} \text{ K}^{-1}$. To integrate up to $T = \infty$, we follow the same high temperature extrapolation procedure as was described in section 2, giving $S(T = \infty) - S(T = 10) = 0.145 \text{ J mol}^{-1} \text{ K}^{-1}$. Doing this simple calculation, we find a total recovered entropy of

$$S(T = \infty) - S(T \approx 0) = 5.675 \text{ J mol}^{-1} \text{ K}^{-1}, \tag{18}$$

less than 2% below the expected maximum entropy value of $R \ln 2 = 5.764 \text{ J mol}^{-1} \text{ K}^{-1}$. The inset of figure 20 clearly shows the entropy recovered near the low temperature transition.

   By considering the entropy recovered by the integration of the finite size system specific heat over the transition (inset, figure 20), we confirm that it is approximately equal to the value of the entropy that we would expect to recover from the latent heat of an infinite system. Using our latent heat calculations above (equations (12) and (16)), this value is approximately $\Delta S_{T_c} = \Delta E / T_c \approx 0.2465/0.180 = 1.37 \text{ J mol}^{-1} \text{ K}^{-1}$, in good agreement with the jump in $S(T)$ in the inset of figure 20.
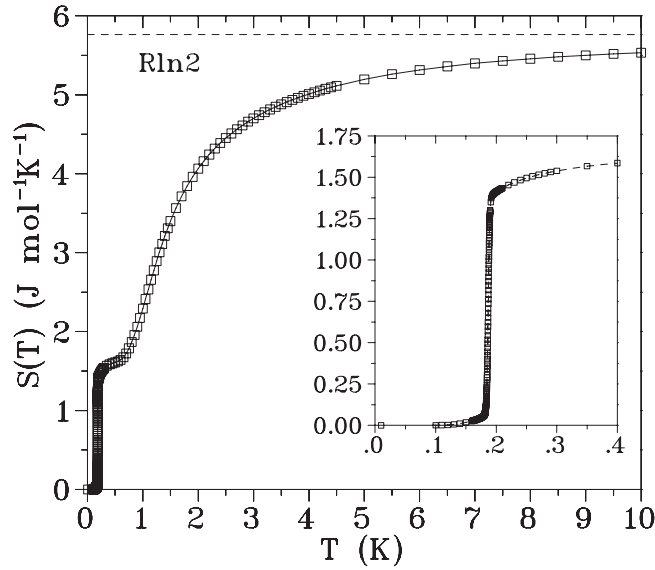
Figure 20. The entropy calculated from integrating the simulated specific heat, as explained in the text. The entire value of the entropy for the system ($R \ln 2$) is recovered in the high temperature limit. The inset shows the details of the entropy recovered by the transition to long range order.

4.4. Summary

Taking all indicators together, we have demonstrated here that the transition to long range order at 180 mK recovers all residual Pauling entropy of the dipolar spin ice model. Thus we can assert that the degeneracy associated with the spin ice model, and the corresponding value of the zero-point entropy, is lifted due to perturbations beyond nearest neighbour dipole–dipole interactions, if equilibrium can be maintained at sufficiently low temperatures.

To summarize our results for this section, we refer the reader to the dipolar spin ice Monte Carlo phase diagram, figure 5. As illustrated there, the transition between the spin ice phase (which retains Pauling's entropy) and the $\mathbf{q} = (0, 0, 2\pi/a)$ ordered phase is independent of the strength of J_{nn}. This is consistent with our understanding that the long range order results from perturbative interactions beyond nearest neighbour, caused by the long range dipolar interaction. This is also what mean field theory [52, 53] finds in the spin ice regime ($J_{nn}/D_{nn} > -0.905$). We find that this first order line also runs slightly up the boundary between the AF $\mathbf{q} = 0$ ordered phase and the higher temperature paramagnetic phase, and that a tricritical point separates these two regions of the line, occurring near the value $J_{nn}/D_{nn} \sim -1.1$.

Due to the near vertical nature of the phase boundaries in this region, simulations run at a finite T and varying J_{nn} help better map out the low temperature phase lines of interest. However, using this method, we observed that the simulations could easily get 'stuck' in the previous spin configuration (either spin ice disordered, $\mathbf{q} = (0, 0, 2\pi/a)$ or AF $\mathbf{q} = 0$) when crossing the vertical phase boundary. This history dependence is illustrated in the phase diagram as hysteresis at low temperatures, mainly between the long range ordered $\mathbf{q} = (0, 0, 2\pi/a)$ and antiferromagnetic $\mathbf{q} = 0$ phases. Regardless of this difficulty, we have confirmed from direct Ewald energy calculations at zero temperature that the true zero-temperature phase boundary between the $\mathbf{q} = (0, 0, 2\pi/a)$ and the AF $\mathbf{q} = 0$ phases lies at
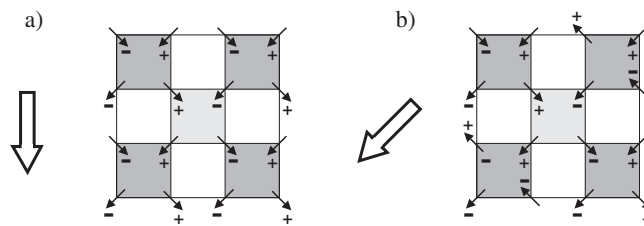
Figure 21. The pyrochlore unit cell projected down the z axis. The symbols $+$ and $-$ represent the z component of the spin 'head'. The configurations are (a) $\mathbf{q} = 0$ and (b) $\mathbf{q} = X$. The large arrows point in the (a) [100] and (b) [110] directions, and are included to aid in the discussion of ground states in the next section. Note that $\mathbf{q} = X$, although similar, is a spin ice state distinct from the $\mathbf{q} = (0, 0, 2\pi/a)$ ordered state shown in figure 4. Specifically, the chains of spins parallel to the [100] direction are staggered antiferromagnetically in the zero-field $\mathbf{q} = (0, 0, 2\pi/a)$ ground state (figure 4), while they are ferromagnetically correlated and parallel to the field in the $\mathbf{q} = X$ state of (b) above.

$J_{\mathrm{nn}}/D_{\mathrm{nn}} = -0.905$, in agreement with the results found in mean field [52, 53] and exact [85] calculations.

5. Dipolar spin ice in a magnetic field

A very interesting problem that pertains to dipolar spin ice materials is their behaviour in an external magnetic field, $\mathbf{h}$. A number of recent experiments [58–60, 62, 78] have shown a rich variety of new behaviour when spin ice materials are subjected to such a field, which warrants some theoretical investigation [61, 86]. Although not all of the relevant experiments can be described in this short section, we briefly outline some of the most important, referring the reader to the reference list for further details on methods and results.

5.1. Experimental overview

The first experiments on spin ice materials in an applied magnetic field were performed by Harris *et al* [21]. In a neutron scattering experiment, they applied a magnetic field of strength 2 T along the [110] direction of a single crystal of $Ho_2Ti_2O_7$, and looked for signs of ordering. They found scattering intensity features which suggest evidence of two ordered magnetic structures, the so-called $\mathbf{q} = 0$ and X phases (figure 21). As we will see below, these two ordered structures are of importance in our study of the ground states of the dipolar spin ice model.

An interesting set of early experiments were performed by Ramirez *et al* [30] and Higashinaka *et al* [87], who subjected polycrystalline samples of $Dy_2Ti_2O_7$ to a variety of different field strengths. Ramirez *et al* presented evidence of phase transitions in a powder sample, manifested as distinct features in the specific heat at

  (i) 0.34 K for $h > 1$ T,
 (ii) 0.47 K for $1\ \mathrm{T} \lesssim h \lesssim 3\ \mathrm{T}$,
(iii) 1.12 K for $h \neq 0$,

where we have used h to represent the magnitude of the applied field $\mathbf{h}$. Higashinaka and co-workers reproduced the basic features of the Ramirez *et al* results down to $T \cong 0.38$ K [87], confirming the existence of the two higher temperature peaks only. The search for a microscopic explanation of these three peaks has been a driving force behind much of the experimental and

theoretical work in this area over the past few years, and we discuss it further in the work that follows below.

Another significant experimental study is the measurement of the single crystal magnetization curves (M versus $\mathbf{h}$) for the spin ice materials. Fukazawa *et al* performed a number of experiments on single crystals of $Dy_2Ti_2O_7$, obtaining magnetization curves for the different applied field directions and a range of temperatures [58]. They showed that magnetization data at 2 K were consistent with the behaviour predicted by the spin ice model, in particular the limiting field (large $\mathbf{h}$) values of the magnetic 'anisotropy' (which we illustrate below). Very recently, measurements [59, 62] of the magnetization curves for $\mathbf{h} \parallel$ [111] (read '$\mathbf{h}$ parallel to the [111] crystal direction') have uncovered a novel macroscopically degenerate state corresponding to ice-like behaviour on the kagomé planes in the pyrochlore lattice [61, 86].

5.2. The Hamiltonian coupling to a magnetic field

We take into account the applied magnetic field $\mathbf{h}$ in the dipolar spin ice model with a simple term added to the Hamiltonian (equation (2)),

$$H' = -\sum_i \mathbf{h} \cdot \mathbf{S}_i^a = -\sum_i (\mathbf{h} \cdot \hat{z}^a)\sigma_i^a. \tag{19}$$

We work strictly with a classical Ising model and neglect any quantum transverse field effects and perturbative changes to the moments arising in a strong field. In this classical approximation, the field $\mathbf{h}$ couples to the spins through the simple scalar product in equation (19). That is, we neglect small corrections to the energy coming from the very small, though finite, local susceptibility perpendicular to the $\langle 111 \rangle$ direction. Also, we neglect the quantum mechanical transverse field effect that would arise from admixing the doublet ground state wavefunctions with that of the excited crystal field levels. For $Ho_2Ti_2O_7$ and $Dy_2Ti_2O_7$, the first excitation gap is (very roughly) $\Delta \sim 300$ K [29]. For the magnetic moments of approximately 10 μ_B for both Ho^{3+} and Dy^{3+}, this means a ground state Zeeman energy splitting of 12.8 K T^{-1}. One can therefore initially neglect magnetic field, exchange and dipole–dipole induced admixing for fields less than 10 T, assuming the worse case scenario where the excited doublet was also split by about 10 K T^{-1}.

To gain a theoretical understanding of the experimental behaviour mentioned above, several insightful calculations are possible, using only this simple classical Hamiltonian and a knowledge of the possible ground states of figure 21. First, a geometrical understanding of how the magnetic field couples to classical spins on the pyrochlore lattice is desirable. We expect that application of a magnetic field along the three principal symmetry axes of the crystal will result in different spin–field coupling behaviour. To explore this, we begin by considering the non-interacting limit ($h \rightarrow \infty$ or J_{nn}, $D_{nn} \rightarrow 0$). In this case, the only constraints on the spins are the local $\langle 111 \rangle$ anisotropy and the coupling with the magnetic field. We can gain more insight by viewing a projection of a tetrahedron down the cubic z axis as in figure 22.

For $\mathbf{h} \parallel$ [100], all four spins on a given tetrahedron are coupled with the field (i.e. all four have a non-zero dot product in equation (19)). The expected lowest energy configuration in the absence of spin–spin interaction is the one where all spins have their [100] components aligned with the field. Knowing this, we can calculate the $h \rightarrow \infty$ value for the M versus $\mathbf{h}$ curves by considering the average moment M (in units of Bohr magneton per rare earth ion, μ_B/R^{3+}), in the direction of the field $\mathbf{h} \parallel$ [100]:

$$M(\mathbf{h} \rightarrow \infty) = \frac{1}{4\sqrt{3}} \left([111] + [1\bar{1}\bar{1}] + [1\bar{1}1] + [11\bar{1}]\right) \cdot [100] \cdot \mu$$

$$= \frac{1}{\sqrt{3}} \cdot \mu \cong 0.5774 \cdot \mu. \tag{20}$$
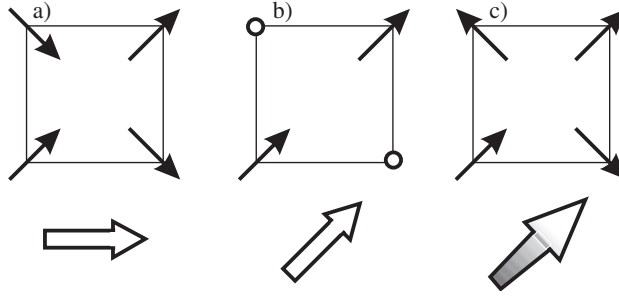
Figure 22. A single tetrahedron projected down the z axis. The field directions are (a) $\mathbf{h} \parallel$ [100], (b) $\mathbf{h} \parallel$ [110], (c) $\mathbf{h} \parallel$ [111], indicated by the large arrow outline. Small arrows represent dipole moments coupled to the field. Empty circles represent decoupled spins.

Also, to calculate M in the appropriate units (μ_B/R^{3+}, measured in experiments) one must include the factor μ, which is the magnetic moment of the appropriate rare earth ion. If R is Dy^{3+} or Ho^{3+}, $\mu \approx 10\ \mu_B$. Note that, from figure 22(a), this lowest energy spin–field coupled state is compatible with the ice rules. If we decorate the entire lattice with tetrahedra such as this, we recover the $\mathbf{q} = 0$ state of figure 21(a). This suggests that this ordered state should be one of the ground states for the interacting dipolar spin ice model, with a sufficiently strong external field $\mathbf{h} \parallel$ [100]. Indeed, this order has been observed experimentally [78] for samples of $Dy_2Ti_2O_7$.
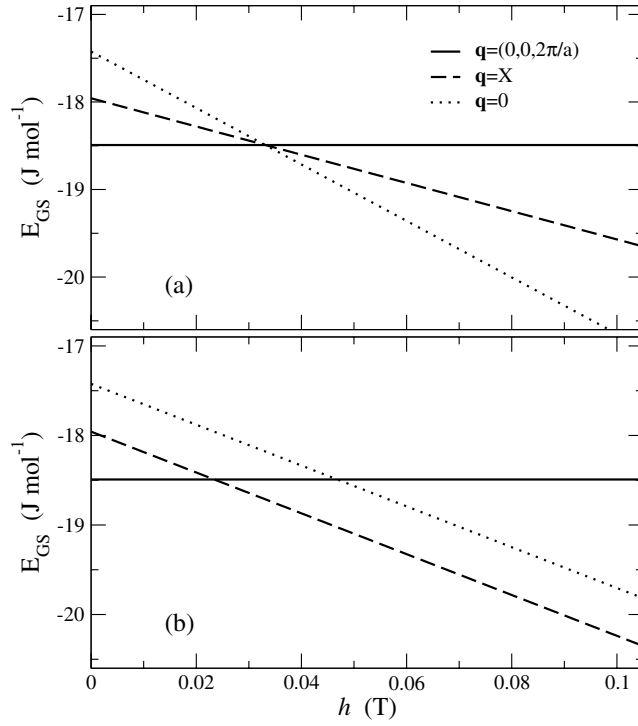
For $\mathbf{h} \parallel$ [110], only two of the four spins on a tetrahedron couple to the field. One expects that, with *precise* alignment of the sample, these other two spins would remain decoupled even in the application of high magnetic fields. These decoupled spins are thus free to choose an ordering pattern that satisfies their dipolar interaction. Because of the complexity of the dipolar interaction, the ground state spin configuration is not immediately obvious from studying the geometry. However, one expects any zero-temperature phase to be consistent with the ice rules (see figure 22(b)). In the limit of very high applied field and perfect sample alignment, one expects the magnetization to approach

$$M(\mathbf{h} \to \infty) = \frac{1}{2\sqrt{3}} \left([111] + \left[11\overline{1}\right]\right) \cdot \frac{1}{\sqrt{2}} [110] \cdot \mu$$

$$= \frac{1}{\sqrt{6}} \cdot \mu \cong 0.4082 \cdot \mu. \qquad (21)$$

Finally, for $\mathbf{h} \parallel$ [111], all four spins on a tetrahedron are coupled to the field. An interesting complication arises in this case due to crystal geometry; any high field phase of the material will be inconsistent with the ice rules, and the spins will form a 'three-in and one-out' (or its spin reverse) tetrahedral configuration (figure 22(c)). For zero temperature, both the long range ordered ice rule state and the three-in and one-out state will exist for different field strengths. For low magnitudes of $\mathbf{h}$, we expect a competition between the exchange, dipolar and magnetic field parts of the Hamiltonian. At low enough temperatures, one predicts [57, 59, 61, 62, 86] a plateau to develop in the magnetization curve due to the tendency of each tetrahedron to stay in the ice rules up to a critical field. If we couple three of the spins to the magnetic field, and leave one to oppose the field but obey the ice rules, we find a magnetization of

$$M(\mathbf{h} = \text{'small'}) = \frac{1}{4\sqrt{3}} \left([111] + \left[\overline{1}11\right] + \left[1\overline{1}1\right] + \left[\overline{1}\overline{1}1\right]\right) \cdot \frac{1}{\sqrt{3}} [111] \cdot \mu$$

$$= \tfrac{1}{3} \cdot \mu \cong 0.3333 \cdot \mu. \qquad (22)$$

Figure 23. The $T = 0$ energies per spin of the three ice rule ordered states of the dipolar spin ice model, as a function of applied internal field $h = |\mathbf{h}|$ along the (a) [100] and (b) [110] directions for $J_{nn} = -0.52$ K and $D_{nn} = 2.35$ K (i.e. Ho$_2$Ti$_2$O$_7$ parameters). The $\mathbf{q} = 0$ and X structures are illustrated in figure 21.

In the limit of very high applied field, we expect the spin that is coupled anti-parallel to the field (in the case above, the last spin vector $[\bar{1}\bar{1}1]$) to break the ice rules, favouring minimizing its energy with respect to the field. In this case the high field magnetization is

$$M(\mathbf{h} \to \infty) = \frac{1}{4\sqrt{3}} \left([111] + [\bar{1}11] + [1\bar{1}1] + [11\bar{1}] \right) \cdot \frac{1}{\sqrt{3}} [111] \cdot \mu$$
$$= \tfrac{1}{2} \cdot \mu = 0.5 \cdot \mu. \tag{23}$$

We find that our Monte Carlo simulation is successful in reproducing the high field limiting values of the experimental M versus $\mathbf{h}$ curves [58]. In addition, we find that the Monte Carlo simulation also reproduces the plateau expected for $\mathbf{h} \parallel [111]$ and low temperatures [57]. However, these large $\mathbf{h}$ results are easily obtainable for a nearest neighbour spin ice model ($D_{nn} \to 0$). The reader is referred to [57] and [58] for the detailed results of this study.

5.3. Ground state energies

A numerical calculation of interest that is easily performed is that of the Ewald energies of the various ground state configurations that we have encountered so far in the dipolar spin ice model. The spin ice configurations that we consider are both the $\mathbf{q} = 0$ and X phases identified by Harris [21], and the $\mathbf{q} = (0, 0, 2\pi/a)$ ground state identified previously in this work (figure 4). In addition, we expect the 'three-in and one-out' state to become the lowest energy state for some critical field along the $\mathbf{h} \parallel [111]$ direction. Figures 23 and 24 are the results
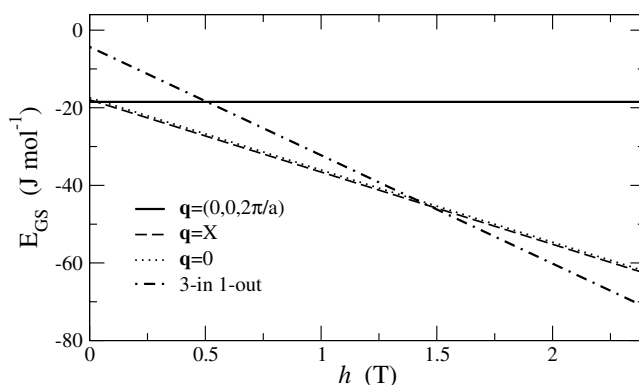
Figure 24. The $T = 0$ energies per spin of the three ice rule ordered states and the three-in and one-out spin state of the dipolar spin ice model as a function of applied internal field **h** along the [111] direction. The **q** = X state becomes the ground state at 0.029 T. The three-in and one-out configuration becomes the ground state at around 1.46 T, breaking the ice rules for each tetrahedron. The **q** = X line is always 0.534 J mol^{-1} below the **q** = 0 line at any given field.

of these ground state energy calculations for a system size $L = 4$ and parameters appropriate for Ho$_2$Ti$_2$O$_7$ ($J_{nn} = -0.52$ K, $D_{nn} = 2.35$ K). As expected, we find the same qualitative behaviour for calculations involving Dy$_2$Ti$_2$O$_7$ parameters. Also, finite size corrections to the energy calculations in figures 23 and 24 are very small, scaling away as $1/L^3$ (see appendix A).

Figure 23(a) confirms that the **q** = 0 configuration becomes the lowest energy state for large field strength ($h > 0.033$ T) for **h** ∥ [100], as expected from the simple geometrical considerations outlined above. Recall that in the **h** ∥ [110] case, there exist two decoupled spins per tetrahedron, and subsequently no lowest energy configuration is obvious from the geometric field coupling equation (19). However, one may anticipate that for **h** ∥ [110], the decoupled spins (figure 22(b)) will order in 'chains' perpendicular to the [110] direction, arranged in such a way as to partially satisfy the dipolar interaction. We find that this is precisely the **q** = X state, which figure 23(b) shows to be the lowest energy state for $h > 0.023$ T.

Consequently, figure 23(b) has direct relevance to experiments by Fennel *et al* and Hiroi *et al* that were performed on Dy$_2$Ti$_2$O$_7$ with a magnetic field **h** ∥ [110] [60, 78]. Fennel *et al* observed neutron diffraction patterns that showed Bragg scattering at **q** = 0 'points' and diffuse scattering at **q** = X 'points', but no true **q** = X long range order. They suggested that this behaviour would arise from long range ferromagnetic order occurring in field coupled spin chains (called α chains by Hiroi [60]), and short range 'antiferromagnetic' order occurring in the field decoupled spin chains (β chains) [78]. In this argument, the true ground state is a **q** = X structure—however, one that is possibly dynamically inhibited from being accessed on experimental timescales.

Specific heat measurements by Hiroi *et al* were used to extract the specific heat contributions of both the α chains and the β chains [60]. They suggest that the specific heat due to the β chains resembles that which one would expect for a low dimensional spin system without long range order. They also argue for the presence of geometrical frustration in the triangular sublattice that contains the β chains. If such a frustration exists, it might be expected to destabilize the 'antiferromagnetic' correlations between these chains that would otherwise lead to **q** = X order. Therefore, Hiroi *et al* argued against the development of true long range order for the system—rather that the β chains become effectively isolated and behave as one-dimensional ferromagnetic systems without long range order in the ground state.

At this point, figure 23(b) is consistent with the idea of long range $\mathbf{q} = X$ order for the dipolar spin ice model with a magnetic field $\mathbf{h} \parallel [110]$. As discussed elsewhere [88], finite temperature Monte Carlo calculations on the dipolar spin ice model also support the idea that, similarly to the development of $\mathbf{q} = (0, 0, 2\pi/a)$ order in the zero-field case, the development of $\mathbf{q} = X$ order for $\mathbf{h} \parallel [110]$ may in some cases be dynamically inhibited in experimental systems with local spin dynamics. A similar result has recently been obtained by Yoshida *et al* in [63]. The failure of the 'frustration' of the β chain sublattice invoked by Hiroi *et al* [60] to destroy the long range $\mathbf{q} = X$ order may be another example of the small energy scale left over by the infinite range dipole–dipole energy (equation (2)), and why such interactions must be handled carefully using techniques such as the Ewald method.

It should also be noted here that the $\mathbf{q} = 0$ and X lines are parallel in figure 23(b) *only* for samples that are perfectly aligned with $\mathbf{h}$ along the [110] crystal axis. This is an important phenomenon that one must consider when comparing theory and experiment, as only a small crystal misalignment will partially couple spins on the β chains to the field. Because *precise* alignment of a crystal is often very difficult, the possibility of misalignment of the order of a degree must be taken into consideration when studying single crystal data with $\mathbf{h} \parallel [110]$. Repeating our ground state energy calculation for misalignment of one degree toward the [100] direction, one would find a crossing of the $\mathbf{q} = X$ and 0 lines (not illustrated) at a field of about 1.3 T, the $\mathbf{q} = 0$ configuration being of lowest energy above this field strength [57, 58].

We turn finally to the case of a field parallel to the [111] crystal axis. As figure 24 confirms, the three-in and one-out spin configuration becomes the lowest energy state for large $\mathbf{h} \parallel [111]$. Interestingly, the $\mathbf{q} = X$ state is the ground state for 0.029 T $< h <$ 1.46 T. This is consistent with the idea that the z component of the field for the [111] direction gives a zero net Zeeman contribution to the unit cell for both the $\mathbf{q} = 0$ and X spin configurations. Therefore, the only energy scale difference left over between the two states comes from the [110] component of the field, meaning that $\mathbf{q} = X$ will be slightly lower in energy than $\mathbf{q} = 0$ (as in figure 23(b)) for all $\mathbf{h} \parallel [111]$. Most other features of the $\mathbf{h} \parallel [111]$ field (such as the intermediate field plateau and the high field breaking of the ice rules) are readily explainable in a nearest neighbour spin ice model without dipolar interaction; hence we will discuss them no further in this work [86].

5.4. Summary and open issues

As mentioned previously, the desire to explain the three field independent specific heat peaks [30, 87] of polycrystalline $Dy_2Ti_2O_7$ has provided recent impetus in this field. Ramirez *et al* were the first to suggest [30] that some of these peaks can be attributed to a fraction of crystallites whose [110] axes happen to align (closely) with the applied magnetic field. More generally, we can interpret this argument as saying that magnetic moments which are not strongly coupled to the magnetic field through equation (19) are free to contribute to a dipole induced phase transition (and therefore sharp peaks in the specific heat) at low temperatures. Historically, the $\mathbf{h} \parallel [110]$ coupling of figure 22(b) was considered to be the most likely scenario for providing these field decoupled spins in a finite number of crystallites in the polycrystalline sample [30]. However, as discussed above, very precise alignment of the field with the [110] crystal axis is probably needed for $\mathbf{q} = X$ order to be maintained to infinite field strength, a condition that is not obviously satisfied for a macroscopic number of crystallites in a polycrystalline sample. One may therefore argue that crystallites with only *one* field decoupled spin would occur in much greater number in a real polycrystalline sample. This is because, for a given crystallite orientation, there is a rotational degree of freedom in the choice of the applied magnetic field direction for which a given sublattice remains decoupled (one of these [59] being $\mathbf{h} \parallel [112]$), which does not exist in the two-spin field decoupled

case ($\mathbf{h} \parallel$ [110]). However, an additional subtlety occurs, as for any field perfectly aligned perpendicular to a local $\langle 111 \rangle$ sublattice (e.g. for $\mathbf{h} \parallel$ [112]), the three sublattices that are coupled to the field will create an *internal* local exchange plus dipolar field that forces the field-perpendicular spin to obey the ice rules. This effect causes a Schottky anomaly in the specific heat at a temperature scale of approximately 2 K. Work to investigate the cancellation of the local internal field with a slight detuning of the external magnetic field *away* from the plane perpendicular to the $\langle 111 \rangle$ direction, which may possibly promote a new low temperature phase transition, is currently in progress [88].

One of the ultimate goals of any finite temperature Monte Carlo study of the dipolar spin ice model in an external magnetic field is to reproduce the general features [30, 87] of the specific heat of polycrystalline $Dy_2Ti_2O_7$. Simulations are currently under way using both single spin flips and loop move Monte Carlo simulation for various field directions to look for signs of an ordering transition in the specific heat. In particular, in the cases of $\mathbf{h} \parallel$ [100] and [110], the ground state is known and hence an order parameter can be constructed, facilitating the identification of any possible phase transitions. Preliminary results [88] indicate the existence of a sharp feature in the specific heat of $Dy_2Ti_2O_7$ for $\mathbf{h} \parallel$ [110], corresponding to a phase transition to the $\mathbf{q} = X$ ground state (see also [63]). However, as this result is very recent it lies outside the scope of this review, and will be left for future discussion [88]. Finally, we should note that in any finite temperature Monte Carlo simulation, the simulation boundary will induce significant demagnetization effects, which must be considered if comparisons are to be made to any experimental results. A brief discussion of demagnetization effects is relegated to appendix C.

To summarize this section, we have performed calculations of the properties of the dipolar spin ice model in an external magnetic field. Ewald energy calculations of various spin configurations reveal the preferred $T = 0$ ordering for various field directions. In particular, for $\mathbf{h} \parallel$ [100] the $\mathbf{q} = 0$ structure is the ground state, while for $\mathbf{h} \parallel$ [110] we find the $\mathbf{q} = X$ state becomes energetically favoured. Finite temperature Monte Carlo simulations of the dipolar spin ice model find [88] that the system settles into these ground states for the respective field directions.

6. Conclusion

We have reviewed much of the early experimental and theoretical work on the static magnetic properties of spin ice. We have also clarified our point of view that long range dipolar interactions are consistent with and responsible for the physics observed in spin ice materials based on Dy^{3+} and Ho^{3+} rare earth ions. Support for this perspective resides in the detailed Monte Carlo calculations presented in this paper, and in mean field studies presented elsewhere [52, 53].

Monte Carlo simulations were performed on the dipolar spin ice model with the long range dipole–dipole interactions treated via the Ewald method. Using a single spin flip Monte Carlo method, we were able to study the development of the spin ice manifold. We found that spins freeze out at temperatures O(1 K) with a macroscopic degeneracy characterized by a residual Pauling entropy. We also found that single spin flip dynamics is not effective at equilibrating the system, thus making it impossible to determine the ordered state of spin ice using this technique.

Mean field theory [52, 53], applied to the same dipolar spin ice Hamiltonian with the dipolar interactions treated via the Ewald method in $\mathbf{q}$-space, showed that an ordering wavevector may be selected and that a proper treatment of the long range dipoles is crucial to achieving a picture consistent with the experiments. A key point is that the symmetry of self-screening is

not exact for the dipolar Hamiltonian. In the end, a quasi-degenerate spectrum emerges with a commensurate ordering wavevector ($\mathbf{q} = (0, 0, 2\pi/a)$) and a two-in and two-out spin ice structure.

In order to find the ordered state of spin ice in a Monte Carlo simulation, we developed a non-local algorithm that employs loop moves (or updates) within the spin ice manifold. These loop moves represent the 'nearly' zero-energy collective dynamics that allows our model to sample the highly degenerate phase space of spin ice. Application of this method at temperatures within the spin ice manifold, i.e., $T \lesssim 1$ K, leads to the selection of a single spin ice ground state configuration with $\mathbf{q} = (0, 0, 2\pi/a)$. The loop Monte Carlo simulation and mean field results agree. In addition, we find a first order transition to the ground state at $T_c \approx 0.08\, D_{nn}$, independent of J_{nn}, which recovers all of the residual Pauling entropy of the spin ice manifold. Our physical understanding of spin ice is aided by the picture that any collective dynamics in real spin ice materials is inhibited by a freezing process as the system enters the temperature range where the ice rule fulfilling manifold develops, i.e., $T_{\text{freeze}} \approx 0.4$ K for $Dy_2Ti_2O_7$ and $T_{\text{freeze}} \approx 0.6$ K for $Ho_2Ti_2O_7$, compared to the predicted $T_c \approx 180$ mK for both materials.

On the strength of the experimental evidence and the success of the dipolar spin ice model, we assert that both $Ho_2Ti_2O_7$ and $Dy_2Ti_2O_7$ are spin ice materials.

Finally, we have discussed the application of a magnetic field to the spin ice materials as a means of exploring the possible structures of the spin ice manifold and to further characterize the interactions present in these intriguing systems. We find excellent agreement between the dipolar spin ice model and many experimental studies to date. In addition, we have made important predictions regarding the ground state configurations of the model for different cases of applied magnetic fields. The consequence of this study in relation to known experimental results is intriguing, arguing for more theoretical, numerical and experimental work, to resolve all the perplexing issues at stake.

Acknowledgments

Appendix A. The Ewald method

We give only a brief overview of the Ewald [89] technique as it applies to dipole–dipole interactions in Monte Carlo simulations. A more detailed discussion of the method can be found in [72] and [90]. The mean field case as it applies to moments on the pyrochlore lattice is treated in detail in [53].

The dipole–dipole interaction is an infinite sum that falls off as the inverse cube of the separation distance between dipoles, $1/|\mathbf{R}_{ij}^{ab}|^3$. Hence it is a conditionally convergent series. The point of the Ewald method is to convert this slowly converging lattice sum into of two

absolutely (rapidly) converging series, one in real space and the other in Fourier space. The general lattice sum for $\langle 111 \rangle$ Ising dipoles on the pyrochlore lattice is

$$
\mathcal{A} = \sum_{i,j} \sum_{a,b} \left(\frac{\hat{z}^a \cdot \hat{z}^b}{|\mathbf{R}_{ij}^{ab}|^3} - \frac{3(\hat{z}^a \cdot \mathbf{R}_{ij}^{ab})(\hat{z}^b \cdot \mathbf{R}_{ij}^{ab})}{|\mathbf{R}_{ij}^{ab}|^5} \right),
$$

$$
= -(\hat{z}^a \cdot \nabla_x)(\hat{z}^b \cdot \nabla_x) \left\{ \sum_{i,j} \sum_{a,b} \frac{1}{|\mathbf{R}_{ij}^{ab} - \mathbf{x}|} \right\}_{\mathbf{x}=0} \tag{A.1}
$$

where the spin variables σ_i^a have been dropped for notational convenience. The dipole sum excludes terms with $R_{ij}^{ab} = 0$. Absolute convergence is forced on the sum inside the curly brackets of equation (A.1) by use of a convergence factor. The form of this convergence factor differs depending on whether the dipolar sum is performed on N particles in real space (e.g., Monte Carlo and molecular dynamic simulations) or in the thermodynamic limit in momentum space (mean field theory).

In our work, MC simulations are performed on three-dimensional lattices of $L \times L \times L$ cubic cells of the pyrochlore lattice under periodic boundary conditions; thus there are $N = 16 \times L^3$ spins in the simulation cell. The separation of moments within a simulation cell is given by $\mathbf{R}_{ij}^{ab}$. The dipolar energy for any pairwise interaction is calculated within the minimum image convention by summing replicas of the N-site simulation cell over spherical shells of radii $\mathbf{n} = L(n_x, n_y, n_z)$ (n_x, n_y, n_z are integers) with the inclusion of a spherical convergence factor $e^{-s|\mathbf{n}|^2}$. The effect of the convergence factor is removed from the final form of the Ewald equations by imposing the limit $s \rightarrow 0$. Therefore, the starting point for the Ewald method is the dipole–dipole pair interaction,

$$
\mathcal{A}_{ij}^{ab}(s) = -(\hat{z}^a \cdot \nabla_x)(\hat{z}^b \cdot \nabla_x) \left\{ \sum_{\mathbf{n}'} \frac{e^{-s|\mathbf{n}|^2}}{|\mathbf{n} + \mathbf{R}_{ij}^{ab} - \mathbf{x}|} \right\}_{\mathbf{x}=0}, \tag{A.2}
$$

where $\sum_{\mathbf{n}'}$ means that $\mathbf{n} = 0$ is omitted whenever $\mathbf{R}_{ij}^{ab} = 0$. The point charge distribution, $1/|\mathbf{n} + \mathbf{R}_{ij}^{ab} - \mathbf{x}|$, is rewritten with the aid of the Γ function identities

$$
\frac{1}{|\mathbf{X}|} = \frac{1}{\sqrt{\pi}} \int_0^\infty t^{-1/2} e^{-t|\mathbf{X}|^2} \, dt \tag{A.3}
$$

$$
= \frac{2}{\sqrt{\pi}} \int_0^\infty e^{-t^2|\mathbf{X}|^2} \, dt. \tag{A.4}
$$

Using equation (A.3), the pairwise interaction becomes

$$
\mathcal{A}_{ij}^{ab}(s) = -(\hat{z}^a \cdot \nabla_x)(\hat{z}^b \cdot \nabla_x) \frac{1}{\sqrt{\pi}} \int_0^\infty dt \left\{ \sum_{\mathbf{n}'} t^{-1/2} e^{-t|\mathbf{n} + \mathbf{R}_{ij}^{ab} - \mathbf{x}|^2 - s|\mathbf{n}|^2} \right\}_{\mathbf{x}=0}, \tag{A.5}
$$

and the remainder of the Ewald calculation for $\mathcal{A}_{ij}^{ab}$ follows arguments outlined in [90]. The Ewald equations for a Monte Carlo simulation can also be found in appendix A of [67]. $\mathcal{A}_{ij}^{ab}$ is calculated for each pairwise interaction, $\{(i, a), (j, b)\}$, in the simulation cell, but this need be done only once because the spins are fixed to the lattice points. These pair interactions are stored in a look-up table and used in the stochastic sampling and measurement procedures of a Monte Carlo simulation.

Finally, recall that the Ewald technique re-sums an infinite number of dipole images, and with the periodic boundary conditions used in the MC simulation cell, the following two constrains are notable: (i) a slight bias toward an ordering wavevector commensurate with the inverse of the simulation cell will develop due to a small enhancement of the relevant

paramagnetic correlations; and (ii) a small finite size scaling, of order $1/L^3$, exists in the Ewald interaction representation of the dipolar interaction. However, both of these effects can be easily systematically investigated in a Monte Carlo simulation by comparing results obtained on a variety of different size lattices.

Appendix B. Ewald method versus truncated dipolar interactions

In this appendix, we compare the results obtained from our loop Monte Carlo simulation with Ewald summation methods to the Monte Carlo results obtained by truncating the dipolar interaction in [32, 33].

Historically, the first Monte Carlo simulation on a model of spin ice was carried out by Ramirez *et al* [30], who used these numerics to supplement their experimental findings on $Dy_2Ti_2O_7$. In their model, spin–spin interactions were assumed to be purely dipole–dipole, but with a g-factor reduced by 25% compared from the expected value for Dy^{3+}. It was commented [30] that this reduction in the strength of the dipole–dipole interaction was most probably the result of a compensating effect due to a small admixture of superexchange interaction. This hypothesis was later confirmed [32] in a series of Monte Carlo simulations complementing DC susceptibility (χ) and specific heat measurements of $Ho_2Ti_2O_7$. In this work, Siddharthan *et al* derived a limiting high temperature series expansion of χ for $\langle 111 \rangle$ Ising spins on the pyrochlore coupled by both nearest neighbour exchange and long range dipole–dipole interactions with the aim of determining the nearest neighbour exchange J_{nn} (assuming a known dipolar coupling strength $D_{nn} \equiv 5D/3 = 2.35$ K). Specifically, by fitting $T\chi(T)$ versus $1/T$ in the temperature range $T = 2$–10 K, they obtained $J_{nn} = -1.92$ K. We believe that this determination of J_{nn} for $Ho_2Ti_2O_7$ in that temperature range is erroneous. A comparison of calculated $\chi(T)$ in the parametric regime, $T > 20$ K, with experimental data for $Ho_2Ti_2O_7$ gives $J_{nn} = -0.55$ [91]. However, this value of $J_{nn} = -1.92$ K was used in Monte Carlo simulations with the dipole–dipole interactions truncated at the fifth nearest neighbour. A freezing into a partially ordered state was found at a temperature depending on the ratio $D_{nn}/(D_{nn} + J_{nn})$ (where $J_{nn} < 0$ for antiferromagnetic exchange) [32].

A different work systematically studied the competition in the spin ice model between exchange and long range dipole–dipole interactions handled via Ewald summation methods [45]. There, using conventional single spin flip Monte Carlo simulations, it was found that for all $J_{nn}/D_{nn} > -0.91$ the system develops spin ice behaviour, as opposed to long range order. Further, for $J_{nn} > 0$ (i.e. ferromagnetic nearest neighbour exchange), the temperature at which the specific heat peak occurs is $T_{peak} \sim 1.1(J_{nn} + D_{nn})$, apparently in contradiction with the results in the inset of figure 4 in [32]. It was argued in [45] that truncation of the dipole–dipole interactions can lead to spurious results, an observation criticized in [33]. In [33], Siddharthan *et al* put forward several compelling arguments for a specific long range ordered ground state for Ising spins on the pyrochlore lattice. A different perspective [52] employed mean field theory calculations of long range dipole coupled $\langle 111 \rangle$ Ising spins on the pyrochlore lattice. There it was argued that important screening effects exist in these systems, and that shells of dipoles at distances larger than the cut-off distance, $R > R_c$, largely cancel the tendency to order created by the shells at smaller $R < R_c$ [52]. Furthermore, a unique commensurate soft mode at $\mathbf{q} = (0, 0, 2\pi/a)$ was predicted. The Monte Carlo simulations presented in the current work, employing loop moves and an Ewald summation formulation for dipolar interactions, confirm the ordering wavevector of [52] and show that it gives rise to a ground state magnetic structure identical to that found in [33].

The above perspective allows us to comment that the difficulties arising with dipolar interactions occur when the nearest neighbour exchange is *antiferromagnetic* and close to the

critical point between a ground state with ordering wavevector at $\mathbf{q} = (0, 0, 2\pi/a)$, and the all-in or all-out state with $\mathbf{q} = 0$ order, which occurs at $J_{\text{nn}}/D_{\text{nn}} = -0.91$. For the region of parameters $-1 < J_{\text{nn}}/D_{\text{nn}} < 0$, it is very important to realize that there is a cancellation (or near cancellation) of the effective nearest neighbour coupling. In that parameter region, the long range part of the dipolar interactions must be calculated very carefully, since the dipolar interactions are to a large extent 'self-screened'. In other words, the development of the paramagnetic correlations (i.e. at temperatures above the ultimate development of long range order) is quantitatively very sensitive to the cut-off distance in this $J_{\text{nn}}/D_{\text{nn}}$ parameter window. Indeed, in that range, the ordering wavevector is typically incommensurate and an irrational function of $J_{\text{nn}}/D_{\text{nn}}$ that depends on relatively large cut-off distances (of order 100 neighbours or more). However, for an Ising system, it is not possible to have a $T = 0$ ground state with a permanent moment structure, where the static ordered moment is equal to unity for all sites. The best candidate ordered structure in the range $-1 < J_{\text{nn}}/D_{\text{nn}} < 0$ (and for moderately large cut-off distance) appears to be the 'same' $\mathbf{q} = (0, 0, 2\pi/a)$ ordered state as we find in this work.

In summary, in the coupling parameter window $-0.91 < J_{\text{nn}}/D_{\text{nn}} < 0$, nearest neighbour antiferromagnetic interactions strongly compete with the dipolar coupling at the nearest neighbour level. In such a case, the dipolar interactions must be summed up to very large distances to allow the long distance self-screening to be manifest. However, the upper limit of the 'validity of truncated dipolar interactions', $J_{\text{nn}}/D_{\text{nn}} = 0$, is not well defined. For example, the present work finds a transition to the $\mathbf{q} = (0, 0, 2\pi/a)$ order at $T_c \approx 0.08 D_{\text{nn}} \sim 0.18$ K and a non-singular specific heat peak at $T_{\text{peak}} \approx 2.3$ K for $D_{\text{nn}} = 2.35$ K. However, Siddharthan *et al* [33] finds a rounded specific heat maximum at $T \sim 0.6$ K and no indication of a feature at $T \geqslant 2.3$ K. As $J_{\text{nn}}/D_{\text{nn}} \to \infty$, the spin ice manifold forms at a temperature $T_{\text{peak}} \sim 0.8 J_{\text{nn}}$ (figure 6). The small dipole induced correlations, which ultimately lead to the $\mathbf{q} = (0, 0, 2\pi/a)$ at $T_c/D_{\text{nn}} = 0.08$, cannot develop via conventional local spin flip dynamics in the frozen ice rule obeying state established at temperatures below $0.8 J_{\text{nn}}$. The equilibrium correlations ultimately developing at $D_{\text{nn}} < T \ll 0.8 J_{\text{nn}}$ can only be well described quantitatively by a proper handling of the long range dipole–dipole interactions.

Appendix C. Demagnetization effects

When doing finite temperature Monte Carlo simulations on magnetic materials in an applied magnetic field, the effect of the boundary of the simulation cell must be carefully considered. For systems of interest, the dipolar spin ice Hamiltonian is augmented with a field dependent term, equation (19). The inclusion of this term in our Monte Carlo simulations leads to subtle effects. In a microscopic Hamiltonian, the field $\mathbf{h}$ referred to in equation (19) is the sample *internal* field, i.e., the magnetic field that directly couples to each magnetic dipole moment. However, in real materials, bulk demagnetization effects alter the magnitude of the internal field in a complicated manner that depends on sample size, shape, alignment and surrounding medium. In general, experimentalists define three separate quantities (the magnetic flux density $\mathbf{B}$, the magnetic field strength $\mathbf{H}$ and the magnetization $\mathbf{M}$) to account for these effects. In a macroscopic material, these quantities are related by

$$\mathbf{B} = \mu_0(\mathbf{H} + \mathbf{M}), \tag{C.1}$$

where $\mathbf{B}$ is the independent quantity controlled in the experiment, but $\mathbf{H}$ is the field strength that couples to the spins (through which the bulk susceptibility is defined). In order to benchmark an experiment to a theory such as ours, an attempt must be made to relate the *external* experimental (applied) $\mathbf{B}$ controlled by external sources of current to the internal $\mathbf{h}$ of our Hamiltonian. From

the experimental side, this amounts to knowing the internal $\mathbf{M}$ associated with the specific sample being measured. This $\mathbf{M}$ is in general not easily deduced—although, for certain sample shapes (e.g., ellipsoids of revolution) it is at least uniform, and fairly accurate estimates can be made. The procedure of correcting for $\mathbf{M}$ to obtain $\mathbf{H}$ is called making a demagnetization correction.

Theoretically, demagnetization effects are incorporated into a Monte Carlo simulation by imposing certain boundary conditions on the microscopic Hamiltonian in question. As described earlier, we use the Ewald summation method to calculate the long range dipolar interactions of our model. We follow the standard approach in which the pairwise interactions are evaluated by summing over periodic copies of the N-site simulation cell until convergence is obtained [90], effectively simulating the infinite range nature of the dipoles. A consequence of this technique is that the finite size nature of the simulation cell is suppressed. We are, therefore, faced with the question of how to interpret an 'infinite boundary'. If one wishes to simulate materials with no net magnetic moment, or materials with no internal demagnetizing field, no correction due to sample boundary is needed, and the simple Ewald sum results may be used. This is equivalent to simulating a long thin 'needle' of the bulk material. However, if one wishes to simulate a material in which the unit cell has a net magnetic moment or internal demagnetizing fields, then we must modify the Ewald sum to take into account the necessary boundary effects. This is especially important in our simulation because the long range nature of the dipole–dipole interactions greatly accentuates these effects.

The approach described by de Leeuw $et\ al$ [90] is to include a $boundary\ term$ in the Ewald sum of the form

$$\left(\frac{4\pi}{2\mu' + 1}\right)\frac{\mu_i \cdot \mu_j}{L^3} \tag{C.2}$$

where μ_i is the magnetic dipole moment of a spin and L is the system linear dimension. Physically, the inclusion of this term corresponds to the consideration of a region external to the spherical Ewald boundary (see the discussion in [90]). This external region is a continuum with the magnetic permeability constant ranging from $\mu' = 1$ to infinity. Because of the nature of our Ewald sum, the $\mu' = 1$ case will in effect simulate the bulk of a spherical sample surrounded by a vacuum. The $\mu' = \infty$ case corresponds to simulating a bulk sample which is needle-like and parallel to an applied $\mathbf{B}$, and hence contains no internal demagnetizing field.

In summary, to make a meaningful comparison between simulation and experiment within the dipolar spin ice model in an applied magnetic field, one of two scenarios must arise:

 (i) Simulations are performed using the regular Ewald summation method, and bulk demagnetization effects are accounted for by experimentalists.
(ii) Simulations are performed with the inclusion of a boundary term, equation (C.2). Experimentalists are restricted to measurements on spherical samples to make quantitative comparisons. However, measurements on other sample shapes (with approximately constant internal fields) may allow some qualitative comparison.

The inclusion of the boundary term is a non-trivial matter in many simulations. For example, it will promote effects such as domain formation in simulations of global Ising ferromagnets. It is therefore always important to check ground state configurations of systems where the term of equation (C.2) is absent against those where it has been included, in order not to miss any important secondary effects.

As a means of addressing some of these issues, we present some preliminary results on Monte Carlo simulations of the dipolar spin ice model. Figure C.1 shows specific heat curves for $Dy_2Ti_2O_7$ in an applied magnetic field, with and without the inclusion of the boundary
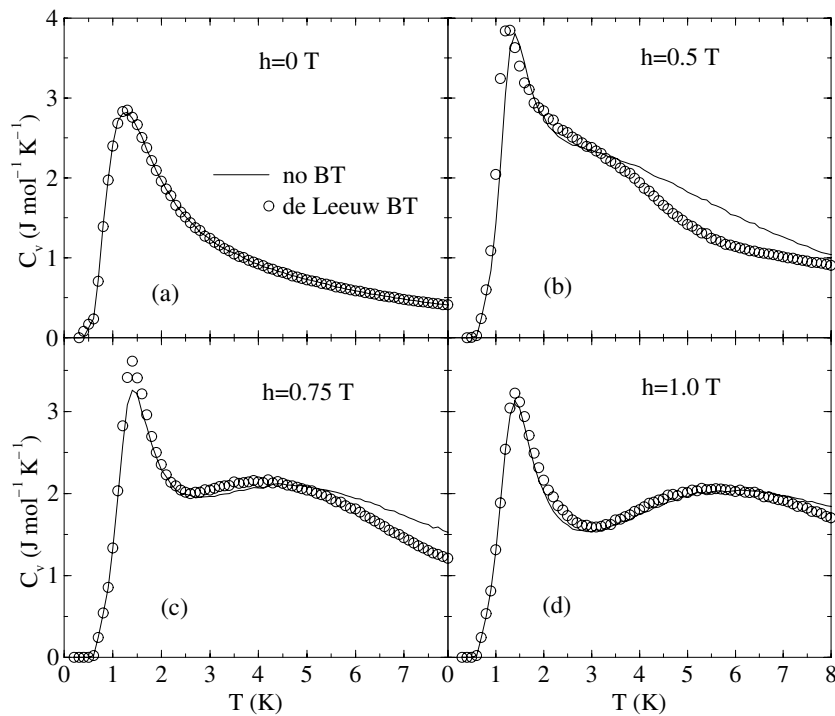
Figure C.1. Specific heat curves for $Dy_2Ti_2O_7$ in an applied field **h** $\parallel$ [110], in simulations without (curve) and with (circles) the boundary term (BT) in the Ewald energy summation. These simulations employed single spin flip dynamics in the Monte Carlo simulation. The significance of these data in the context of experimental measurements on $Dy_2Ti_2O_7$ is discussed in detail in section 5.

term, equation (C.2), with $\mu' = 1$. The results presented in figure C.1 were performed using single spin flip dynamics on a system of size $L = 3$. As we see, the boundary term does not affect the $h = 0$ specific heat curve. This is consistent with the understanding that the boundary term is only necessary in Monte Carlo programs where the simulation cell has a net magnetic moment (which is not true in general for the spin ice manifold). For moderate fields, we see that the boundary term significantly alters the shape of the specific heat curve, as expected, since the simulation cell is acquiring a net magnetic moment. For very large fields, the boundary term begins to lose its effect, as the field term (equation (19)) becomes dominant in the Hamiltonian.

References

[1] Toulouse G 1977 *Commun. Phys.* **2** 115
[2] Villain J 1979 *Z. Phys.* B **33** 31
[3] Diep H T (ed) 1994 *Magnetic Systems with Competing Interactions* (Singapore: World Scientific)
[4] See 2001 *Proc. Waterloo Conf. on Highly Frustrated Magnetism*; *Can. J. Phys.* **79** 1283–596
[5] Ramirez A 2001 *Handbook of Magnetic Materials* vol 13, ed K J H Buschow (New York: New Holland) pp 423–520
[6] Schiffer A P and Ramirez A P 1996 *Comments Condens. Matter Phys.* **18** 21
[7] Greedan J E 2001 *J. Mater. Chem.* **11** 37
[8] Moessner R 2001 *Can. J. Phys.* **79** 1283

[9] Henley C L 2001 *Can. J. Phys.* **79** 1307

[10] Bramwell S T and Gingras M J P 2001 *Science* **294** 1405

[11] Moessner R and Chalker J T 1998 *Phys. Rev. Lett.* **80** 2929
     Moessner R and Chalker J T 1998 *Phys. Rev.* B **58** 12049

[12] Shastry B S 2003 *Physica* B **329** 1024

[13] Lhuillier C and Misguich G 2001 unpublished *Preprint cond-mat/0109146*

[14] Misguich G, Serban D and Pasquier V 2002 *Phys. Rev. Lett.* **89** 137202

[15] Lhuillier C, Sindzingre P and Fouet J-B 2001 *Can. J. Phys.* **79** 1525

[16] Sindzingre P, Fouet J-B and Lhuillier C 2002 *Phys. Rev.* B **66** 174424

[17] Moessner R and Sondhi S L 2001 *Phys. Rev. Lett.* **86** 1881
     Hermele M, Fisher M P A and Balents L 2004 *Phys. Rev.* B **69** 064404

[18] Büttgen N, Kaps H, Trinkl W, Loidl A, Klemm M and Horn S 2001 *Can. J. Phys.* **79** 1487

[19] Lacroix C 2001 *Can. J. Phys.* **79** 1469

[20] Ballou R 2001 *Can. J. Phys.* **79** 1475

[21] Harris M J, Bramwell S T, McMorrow D F, Zeiske T and Godfrey K W 1997 *Phys. Rev. Lett.* **79** 2554

[22] Bramwell S T and Harris M J 1998 *J. Phys.: Condens. Matter* **10** L215

[23] Anderson P W 1956 *Phys. Rev.* **102** 1008

[24] Bramwell S T, Gingras M J P and Reimers J N 1994 *J. Appl. Phys.* **75** 5523

[25] Bernal J D and Fowler R H 1933 *J. Chem. Phys.* **1** 515

[26] Pauling L 1935 *J. Am. Chem. Soc.* **57** 2680

[27] Pauling L 1960 *The Nature of the Chemical Bond* (Ithaca, NY: Cornell University Press) pp 464–9

[28] Giauque W F and Ashley M F 1933 *Phys. Rev.* **43** 81
     Giauque W F and Stout J W 1936 *J. Am. Chem. Soc.* **58** 1144

[29] Rosenkranz S, Ramirez A P, Hayashi A, Cava R J, Siddharthan R and Shastry B S 2000 *J. Appl. Phys.* **87** 5914

[30] Ramirez A P, Hayashi A, Cava R J, Siddharthan R and Shastry B S 1999 *Nature* **399** 333

[31] Harris M 1999 *Nature* **399** 311

[32] Siddharthan R, Shastry B S, Ramirez A P, Hayashi A, Cava R J and Rosenkranz S 1999 *Phys. Rev. Lett.* **83** 1854

[33] Siddharthan R, Shastry B S and Ramirez A P 2001 *Phys. Rev.* B **63** 184412

[34] Bramwell S T, Harris M J, den Hertog B C, Gingras M J P, Gardner J S, McMorrow D F, Wildes A R,
     Cornelius A L, Champion J D M, Melko R G and Fennell T 2001 *Phys. Rev. Lett.* **87** 047205

[35] Cornelius A L and Gardner J S 2001 *Phys. Rev.* B **64** 060406

[36] Petrenko O A, Lees M R and Balakrishnan G 2003 *Phys. Rev.* B **68** 012406

[37] Bondah-Jagalu V and Bramwell S T 2001 *Can. J. Phys.* **79** 1381

[38] Kadowaki H, Ishii Y, Matsuhira K and Hinatsu Y 2002 *Phys. Rev.* B **65** 144421

[39] Matsuhira K, Hinatsu Y, Tenya K, Amitsuka H and Sakakibara T 2002 *J. Phys. Soc. Japan* **71** 1576

[40] Bansal C, Kawanaka H, Bando H and Nishihara Y 2002 *Phys. Rev.* B **66** 052406

[41] Matsuhira K, Hinatsu Y, Tenya K and Sakakibara T 2000 *J. Phys.: Condens. Matter* **12** L649

[42] Matsuhira K, Hinatsu Y and Sakakibara T 2001 *J. Phys.: Condens. Matter* **13** L737

[43] Snyder J, Slusky J S, Cava R J and Schiffer P 2001 *Nature* **413** 48
     Snyder J, Slusky J S, Cava R J and Schiffer P 2002 *Phys. Rev.* B **66** 064432
     Snyder J, Ueland B G, Slusky J S, Karunadasa H, Cava R J, Mizel A and Schiffer P 2003 *Phys. Rev. Lett.* **91**
     107201
     Snyder J, Ueland B G, Slusky J S, Karunadasa H, Cava R J and Schiffer P 2004 *Phys. Rev.* B **69** 064414
     Snyder J, Ueland B G, Mizel A, Slusky J S, Karunadasa H, Cava R J and Schiffer P 2004 *Preprint* cond-
     mat/0405233

[44] Ehlers G, Cornelius A L, Orendác M, Kajnaková M, Fennell T, Bramwell S T and Gardner J S 2003 *J. Phys.:
     Condens. Matter* **15** L9
     Gardner J S, Ehlers G, Bramwell S T and Gaulin B B 2004 *J. Phys.: Condens. Matter* **16** S643
     Ehlers G, Cornelius A L, Fennell T, Koza M, Bramwell S T and Gardner J S 2004 *J. Phys.: Condens. Matter*
     16 S635

[45] den Hertog B C and Gingras M J P 2000 *Phys. Rev. Lett.* **84** 3430

[46] Harris A B, Kallin C and Berlinsky A J 1992 *Phys. Rev.* B **45** 2899

[47] Reimers J N, Berlinsky A J and Shi A-C 1991 *Phys. Rev.* B **43** 865

[48] Greedan J E, Raju N P, Wills A S, Morin C, Shaw S M and Reimers J N 1998 *Chem. Mater.* **10** 3058

[49] Kanada M, Yasui Y, Kondo Y, Iikubo S, Ito M, Harashina H, Sato M, Okumura H, Kakurai K and
     Kadowaki H 2002 *J. Phys. Soc. Japan* **71** 313

[50] Harris M J, Bramwell S T, Holdsworth P C W and Champion J D M 1998 *Phys. Rev. Lett.* **81** 4496

[51] Champion J D M, Bramwell S T, Holdsworth P C W and Harris M J 2002 *Europhys. Lett.* **57** 93

[52] Gingras M J P and den Hertog B C 2001 *Can. J. Phys.* **79** 1339
[53] Enjalran M and Gingras M J P 2003 *Preprint* cond-mat/0307151
[54] Barkema G T and Newman M E J 1998 *Phys. Rev.* E **57** 1155
[55] Newman M E J and Barkema G T 1999 *Monte Carlo Methods in Statistical Physics* (Oxford: Clarendon)
[56] Melko R G, den Hertog B C and Gingras M J P 2001 *Phys. Rev. Lett.* **87** 067203
[57] Melko R G 2001 *MSc Thesis* University of Waterloo
[58] Fukazawa H, Melko R G, Higashinaka R, Maeno Y and Gingras M J P 2002 *Phys. Rev.* B **65** 054410
[59] Higashinaka R, Fukazawa H and Maeno Y 2003 *Phys. Rev.* B **68** 014415
[60] Hiroi Z, Matsuhira K and Ogata M 2003 *J. Phys. Soc. Japan* **72** 3045
[61] Wills A S, Ballou R and Lacroix C 2002 *Phys. Rev.* B **66** 144407
[62] Matsuhira K, Hiroi Z, Tayama T, Takagi S and Sakakibara T 2002 *J. Phys.: Condens. Matter* **14** L559
     Hiroi Z, Matsuhira K, Takagi S, Tayama T and Sakakibara T 2003 *J. Phys. Soc. Japan* **72** 411
[63] Yoshida S-I, Nemoto K and Wada K 2004 *J. Phys. Soc. Japan* **73** 1619
[64] Bramwell S T, Gingras M J P and Reimers J N 1994 *J. Appl. Phys.* **75** 5523
[65] Moessner R 1998 *Phys. Rev.* B **57** R5587
[66] Gardner J S, Dunsiger S R, Gaulin B D, Gingras M J P, Greedan J E, Kiefl R F, Lumsden M D, MacFarlane W A,
     Raju N P, Sonier J E, Swainson I and Tun Z 1999 *Phys. Rev. Lett.* **82** 1012
[67] Gingras M J P, den Hertog B C, Faucher M, Gardner J S, Dunsiger S R, Chang L J, Gaulin B D, Raju N P and
     Greedan J E 2000 *Phys. Rev.* B **62** 6496
[68] Gardner J S, Gaulin B D, Berlinsky A J, Waldron P, Dunsiger S R, Raju N P and Greedan J E 2001 *Phys. Rev.*
     B **64** 224416
[69] Yasui Y, Kanda M, Ito M, Harashina H, Sato M, Okumura H, Kakurai K and Kadowaki H 2002 *J. Phys. Soc.*
     Japan **71** 599
[70] Kao Y-J, Enjalran M, Del Maestro A, Molavian H R and Gingras M J P 2003 *Phys. Rev.* B **68** 172407
[71] Mirebeau I, Goncharenko I N, Cadavez-Peres P, Bramwell S T, Gingras M J P and Gardner J S 2002 *Nature*
     420 54
[72] Ziman J M 1972 *Principles of the Theory of Solids* 2nd edn (Cambridge: Cambridge University Press)
     Born M and Huang K 1968 *Dynamical Theory of Crystal Lattices* Oxford
[73] Camp P J, Shelley J C and Patey G N 2000 *Phys. Rev. Lett.* **84** 115
     Ayton G, Gingras M J P and Patey G N 1995 *Phys. Rev. Lett.* **75** 2360
[74] den Hertog B C, Gingras M J P, Bramwell S T and Harris M J 1999 unpublished *Preprint cond-mat/9912220*
[75] Blote H W J, Wielinga R F and Huiskamp W J 1969 *Physica* **43** 549
[76] Ramirez A P and Jensen J 1994 *J. Phys.: Condens. Matter* **6** L215
[77] Bertin E, Bonville P, Bouchaud J-P and Hodges J A 2002 *Eur. Phys. J.* B **27** 347
[78] Fennell T, Petrenko O A, Balakrishnan G, Bramwell S T, Champion J D M, Fåk B, Harris M J
     and Paul D McK 2002 *Appl. Phys.* A **74** S889
[79] Fennell T, Petrenko O A, Fak B, Bramwell S T, Enjalran M, Yavors'kii T, Gingras M J P, Melko R G
     and Balakrishnan G 2004 *Preprint* cond-mat/0404594
[80] Shore J D, Holzer M and Sethna J P 1992 *Phys. Rev.* B **46** 11376
[81] Binder K and Heermann D W 1997 *Monte Carlo Simulation in Statistical Physics* 3rd edn (Berlin: Springer)
[82] Berg B A and Neuhaus T 1992 *Phys. Rev. Lett.* **68** 9
[83] Hansmann U H E and Okamoto Y 1994 *Physica* A **212** 415
[84] Ferrenberg A M and Swendsen R H 1988 *Phys. Rev. Lett.* **61** 2635
[85] Chamberlain S L, Hess S T and Corruccini L R 2004 *Phys. Lett.* A **323** 310
[86] Moessner R and Sondhi S L 2003 *Phys. Rev.* B **68** 064411
     Isakov S V, Raman K S, Moessner R and Sondhi S L 2004 *Preprint* cond-mat/0404417
[87] Higashinaka R, Fukazawa H, Yanagishima D and Maeno Y 2002 *J. Phys. Chem. Solids* **63** 1043
[88] Ruff J P C, Melko R G and Gingras M J P, unpublished
[89] Ewald P P 1921 *Ann. Phys., Lpz.* **64** 253
[90] de Leeuw S W, Perram J W and Smith E R 1986 *Annu. Rev. Phys. Chem.* **37** 245
[91] Molavian H R, Higashinaka R, Maeno Y and Gingras M J P, unpublished