

Neutron Scattering and Thermodynamic Studies
of the Magnetic Material $\text{Yb}_2\text{Ti}_2\text{O}_7$ at mK
Temperatures

Jordan Thompson
Lady Margaret Hall



Thesis submitted for the Degree of Doctor of Philosophy
Trinity Term 2016

This thesis is dedicated to my love Vânia Pinto, my family, and all of the friends I
have made during my time at Oxford.

Neutron Scattering and Thermodynamic Studies of the Magnetic Material $\text{Yb}_2\text{Ti}_2\text{O}_7$ at mK Temperatures

Jordan Thompson, Lady Margaret Hall

Thesis submitted for the Degree of Doctor of Philosophy, Trinity Term 2016

The frustrated pyrochlore magnet $\text{Yb}_2\text{Ti}_2\text{O}_7$ has been proposed as a candidate “quantum spin ice” material. This thesis reports extensive inelastic neutron scattering measurements of the spin dynamics and complementary specific heat measurements to map the phase diagram in magnetic fields applied along the cubic [001] direction, not explored experimentally in detail before. The high quality of the single crystals is evidenced by a sharp specific heat anomaly near 215 mK in zero field. In field heat capacity measurements show that the heat capacity peak is completely suppressed by very small magnetic fields ($\mu_0 H_{\text{int}} \sim 0.018$ T), and the opening of an excitation gap upon increasing field in the form of a broad hump in the data. Inelastic neutron scattering measurements at the highest magnetic fields probed (9 T), in addition to dominant one-magnon excitations, also show a broad scattering continuum at higher energies, not detected in previous neutron scattering studies, which is attributed to two-magnon excitations. The evolution of the spectrum upon lowering field is characterised in detail, showing how high-energy magnons decay when they overlap with the high-energy continuum ($B < 3$ T) and the low-energy magnons have their dispersion bandwidth suppressed upon lowering field, and disappear altogether at zero field, where the scattering spectrum consists of a gapless continuum extending over a broad energy range. Through fits of the dispersion relations at high field a re-evaluation of the spin Hamiltonian is proposed to consistently explain quantitatively all existing spin dynamics data for different field directions. A theoretical mean-field calculation indicates that the spin Hamiltonian is very close to a phase boundary line between distinct magnetic orders, and this may be the reason for the unconventional magnetic behaviour, where magnons cannot propagate coherently at any of the energy scales probed.

Acknowledgements

There are many people without whom the completion of this work would have been impossible. The first of these is my supervisor, Radu Coldea, who I would like to thank for all of his hard work, encouragement, and his many patient explanations of the concepts appearing in this thesis. Another such person is Paul McClarty, who I would like to thank for his invaluable collaboration on the theoretical work in this thesis. I would also like to thank both of them for editing chapters of this thesis.

In the Clarendon Laboratory, there is a long list of people deserving of thanks, either for directly contributing to this research, or for making the lab an enjoyable environment: Matt Watson, Arjun Narayanan, and Sam Blake for help in the lab and many trips to the pub; Ivelisse Cabrera and Alun Biffin, for their hard work and company during long beam times; Danielle Kaminski, Liam Gannon, Natasha Perks, and the many other members of the Clarendon tea group, for many interesting and entertaining conversations; Dharmalingam Prabhakaran, for growing the crystals of $\text{Yb}_2\text{Ti}_2\text{O}_7$ used in this work; the Clarendon machine workshop staff for constructing the experimental apparatuses necessary for these experiments; Rob Storey and Tony Hickman, for all of their hard work and assistance; Bob Watkins for his assistance installing and maintaining various lab equipment, and for cutting the crystals used in this work into shape; and finally the summer students who worked on developing the fridge_control software package.

Outside of the Clarendon, I would also like to thank John Dunn and the staff of ICE Oxford for teaching me how to operate a dilution refrigerator; and the ISIS neutron source support staff, and instrument scientists Rob Bewley and Tatiana Guidi, for their invaluable assistance with many neutron scattering experiments. I would also like to thank all of my friends from the Lady Margaret Hall MCR and all of the members of the Oxford University Ice Hockey Club. I need to thank my family for their patience and support and many boxes of cookies over the last five years. Finally, I would like to thank my girlfriend, Vânia Pinto, who has provided invaluable love and support during the final push to complete this work.

Table of Contents

List of Tables	ix
List of Figures	xiii
1 Introduction	1
1.1 The Origin of Magnetic Moments in Insulators	1
1.2 Magnetic Exchange Interactions	2
1.3 Magnetic Order	3
1.4 Mean-Field Theory	5
1.5 Magnetic Frustration	6
1.6 Spin Ice	7
1.7 The Pyrochlore Lattice	10
1.8 Previous Work on $\text{Yb}_2\text{Ti}_2\text{O}_7$	11
1.8.1 The Low Temperature Phase of $\text{Yb}_2\text{Ti}_2\text{O}_7$	12
1.8.2 Sample Dependence in $\text{Yb}_2\text{Ti}_2\text{O}_7$	18
1.8.3 The Exchange Hamiltonian of $\text{Yb}_2\text{Ti}_2\text{O}_7$	23
1.8.4 The Crystal Field of $\text{Yb}_2\text{Ti}_2\text{O}_7$	35
1.8.5 Other Works on $\text{Yb}_2\text{Ti}_2\text{O}_7$	37
2 Heat Capacity	40
2.1 Introduction	40
2.2 Theory	41
2.2.1 The First Law of Thermodynamics	41
2.2.2 The Second Law of Thermodynamics and Entropy	43
2.2.3 Entropy and Statistical Mechanics	44
2.3 Methods	46
2.3.1 AC Calorimetry	46
2.4 Experiment Design	50

2.4.1	Sample Environment	51
2.4.2	Apparatus Design	52
2.4.3	Experimental Electronics and Control	61
2.5	Analysis Procedure and Test Measurement on Cs_2CoCl_4	66
2.5.1	Analysis Procedure	67
2.5.2	Test Measurement: Cs_2CoCl_4	71
2.6	Conclusions	74
3	Low Temperature Heat Capacity of $\text{Yb}_2\text{Ti}_2\text{O}_7$	75
3.1	Introduction	75
3.2	Sample	78
3.3	Zero Field Measurements	78
3.4	Measurements in Applied Field	81
3.5	Schottky-Like Anomaly	84
3.6	Conclusions	88
4	Neutron Scattering	89
4.1	Introduction	90
4.2	Scattering Theory	91
4.2.1	Nuclear Scattering	93
4.2.2	Magnetic Scattering	95
4.3	Time of Flight Neutron Spectroscopy	100
4.4	Neutron Sources	102
4.5	Horace Scans	103
5	Inelastic Neutron Scattering Measurements of the Magnetic Excitations in $\text{Yb}_2\text{Ti}_2\text{O}_7$	105
5.1	Experimental Details	106
5.1.1	Sample Alignment and Mount	106
5.1.2	Sample Environment	108
5.1.3	Instrument Configuration	108
5.2	Neutron Scattering Data	109
5.2.1	Evolution of Magnetic Excitations in $[001]$ Magnetic Field	110
5.2.2	Dynamical Correlations in the Full Brillouin Zone at Representative Magnetic Fields	123
5.3	Conclusions	128

6	Effective Spin Half Linear Spin Wave Description of the One and Two Magnon Excitations in $\text{Yb}_2\text{Ti}_2\text{O}_7$ in Applied Magnetic Field	130
6.1	Linear Spin Wave Theory	131
6.2	One Magnon Dispersion Points	140
6.3	Determining the Optimised Exchange Hamiltonian	143
6.4	One Magnon Dispersions: Data and Spin Wave Calculations	149
6.5	High Energy Two Magnon Continuum Scattering: Data and Calculations	157
6.6	Magnon Interaction and Decay at Intermediate Applied Fields	166
6.7	Comparison of Heat Capacity and Neutron Scattering Energy Gaps	171
6.8	Curie-Weiss Temperature	172
6.9	Mean-Field Phase Diagram	173
6.10	Conclusions	177
7	Conclusions and Future Work	179
7.1	Conclusions	179
7.2	Future Work	185
	APPENDICES	187
A	RuO_2 Thermometer Calibration	188
A.1	Zero Field Calibration	188
A.2	Finite Field Calibration	190
B	Sample Growth	193
C	Laue Alignment	196
D	Magnetisation Measurements and Demagnetisation Corrections	198
D.1	Magnetisation Measurements	198
D.1.1	Introduction	198
D.1.2	The SQUID Magnetometer	198
D.1.3	Sample	199
D.1.4	Demagnetisation Correction	200
D.1.5	Results	202
D.2	Mean-Field Magnetisation for Least Squares Fit	202
D.3	Demagnetisation Corrections for Neutron Scattering and Heat Capacity Measurements	203
D.3.1	Introduction	203
D.3.2	Heat Capacity Demagnetisation Correction	204
D.3.3	Inelastic Neutron Scattering Demagnetisation Correction	205

E	Background Subtraction	209
F	Neutron Absorption Correction	213
	References	216

List of Tables

1.1	Anisotropic $\mathbf{g}$ tensor components of the ground state CEF doublet from various works	38
3.1	Heat capacity features of single crystal samples of $\text{Yb}_2\text{Ti}_2\text{O}_7$, catalogued from the literature, along with the low temperature magnetic behaviour observed in those crystals.	78
5.1	LET Configuration Information	109
6.1	Local quantisation directions for the tetrahedral sublattice sites . . .	134
6.2	The optimised Exchange Hamiltonian parameters compared to those of other works	145
6.3	Components of the total moment sum rule at 9 T	163
D.1	Values of $\mu_0 H_{\text{int}}$ for corresponding values of the applied magnetic field for heat capacity measurements performed in this work	205
D.2	Values of $\mu_0 H_{\text{int}}$ for corresponding values of the applied magnetic field for neutron scattering experiments in this work	208
F.1	The absorption and incoherent scattering cross sections of the elements in $\text{Yb}_2\text{Ti}_2\text{O}_7$	214

List of Figures

1.1	The Néel mean-field ground state of the square lattice antiferromagnet	4
1.2	Illustrations of magnetic frustration on the triangular lattice and the square lattice with next-nearest-neighbour interactions	7
1.3	The local constraints and excitations of spin ice	9
1.4	The pyrochlore lattice, and the lattice of corner sharing tetrahedra formed by the A site	11
1.5	Low temperature magnetic order in $\text{Yb}_2\text{Ti}_2\text{O}_7$ proposed by Chang et al.	17
1.6	Representative heat capacity of powder and single crystal samples of $\text{Yb}_2\text{Ti}_2\text{O}_7$ from the literature	20
1.7	The easy axes and easy planes of the four Yb^{3+} sites in $\text{Yb}_2\text{Ti}_2\text{O}_7$. .	25
1.8	The oxygen environment of a Yb^{3+} ion in $\text{Yb}_2\text{Ti}_2\text{O}_7$	36
2.1	The entropy and heat capacity of first order and continuous phase transitions	44
2.2	Schottky anomaly heat capacity	46
2.3	Schematic diagrams of the basic configuration of AC heat capacity measurements for samples of infinite and finite thermal conductivity, along with the frequency dependent correction factor	49
2.4	Labelled photograph of the coldfinger that supports the heat capacity setup	53
2.5	Schematic diagram of the AC heat capacity apparatus.	54
2.6	Typical resistance curves of Lakeshore RuO_2 thermometers as a function of temperature.	58
2.7	Schematic diagram of the instruments and connections necessary for low temperature AC heat capacity measurements.	62
2.8	Figure showing the raw data collected during a heat capacity experiment and how it is converted into a value for the heat capacity	68
2.9	Frequency sweep and raw heat capacity of Cs_2CoCl_4 compared to the values of the heat capacity from the literature	71
2.10	Corrected frequency sweep and heat capacity of Cs_2CoCl_4 compared to the values of the heat capacity from the literature	73

3.1	The heat capacity of single crystal samples of $\text{Yb}_2\text{Ti}_2\text{O}_7$ reported in the literature	77
3.2	Heat Capacity Measurements on $\text{Yb}_2\text{Ti}_2\text{O}_7$ Single Crystal in Zero Field and Associated Frequency Sweep	80
3.3	Heat Capacity of $\text{Yb}_2\text{Ti}_2\text{O}_7$ in increasing applied magnetic field along [001] and resulting temperature-magnetic field phase diagram	82
3.4	Fits to the Schottky-like anomaly in the intermediate field heat capacity measurements on $\text{Yb}_2\text{Ti}_2\text{O}_7$, along with a comparison to spin wave heat capacity.	85
4.1	Schematic diagram of the geometry of a basic neutron scattering experiment	91
4.2	Schematic diagram of a direct geometry time of flight spectrometer .	101
5.1	$\text{Yb}_2\text{Ti}_2\text{O}_7$ single crystal and mount used for inelastic neutron scattering experiments	107
5.2	Overview of the evolution of magnetic excitations as a function of applied magnetic field	111
5.3	Magnetic inelastic scattering energy scans for applied magnetic fields between 0 T and 3 T	113
5.4	Low field continuum structure	114
5.5	Overview of the field evolution of the two magnon continuum scattering and the $(h, k, 0)$ dependence of the continuum intensity	116
5.6	High energy magnetic continuum scattering energy scans, fitting function, scattering bandwidths, and energy integrated slice in the $(h, k, 0)$ plane	118
5.7	The relationship between one and two magnon scattering	121
5.8	Representative slices through Horace scan data for applied fields of 0, 0.21, 1.5 and 5 T	125
5.9	Slices through 1.34 meV Horace data for applied fields of 0 T, 0.21 T, and 1.5 T	126
5.10	Energy integrated slices through 0 T Horace Scan	127
5.11	Energy scans through 2.5 meV and 1.34 meV Horace data sets	128
6.1	Example 2.5 meV 5 T Magnon dispersion points	141
6.2	Example 4.0 meV 9 T Magnon dispersion points	142
6.3	Experimental and Linear Spin Wave dispersion energies, the χ^2 environment of the least squares fit, correlations of the various free parameters in the effective Hamiltonian, and correlations between the $\mathbf{g}$ tensor components	146

6.4	Linear spin wave calculation of ΔS , the reduction in the ordered spin moment from its maximum possible value of 1/2 as a function of applied field	148
6.5	Slices through 2.5 meV, 5 T Horace data for several directions in reciprocal space, along with the same slices calculated using optimised parameter set, and the model of Ross et al.	151
6.6	Additional slices through 2.5 meV, 5 T Horace data for several directions in reciprocal space, along with the same slices calculated using optimised parameter set, and the model of Ross et al.	152
6.7	A comparison of the 4.0 meV, 9 T magnon dispersion to spin wave calculations	153
6.8	A comparison of the high field inelastic neutron scattering of Ross et al. to LSW calculations	154
6.9	A comparison of 2.5 meV, 5 T single angle inelastic neutron scattering to spin wave calculations including a scan in energy through the experimental and model data	155
6.10	Overview of the magnetic excitations as a function of applied field compared with linear spin wave theory predictions	156
6.11	The one and two magnon scattering bandwidth extracted from experimental data compared to the bandwidth as calculated by linear spin wave theory	157
6.12	Energy scan through 7 T 6.3 meV single angle inelastic neutron scattering data showing both the one and two magnon scattering, along with spin wave calculations of both scattering terms.	160
6.13	Comparison of energy scans through the two magnon continuum to spin wave calculations for fields of $3 \text{ T} \leq B \leq 9 \text{ T}$	162
6.14	Energy integrated slice through the 5 T, 6.3 meV Horace data in the $(h, k, 0)$ plane, showing the high energy magnetic continuum scattering, along with the same slice generated using linear spin wave theory	165
6.15	An energy but through 2.5 meV 0 T inelastic neutron scattering showing the excitation gap in the linear spin spectrum compared to the gapless continuum in the experimental data	168
6.16	A comparison of the gap to the lowest energy excitation extracted from inelastic neutron scattering to the energy gap extracted from heat capacity measurements and linear spin wave calculations of the minimum excitation gap	172
6.17	The semi-classical phase diagram in the space of exchange energies J_e showing the location of the optimised Hamiltonian determined in this work with relation to a nearby phase boundary	174
6.18	Magnetic field dependence of the effective spin-1/2 and projected full moment directions on a single tetrahedron in the splayed ferromagnetic phase for a magnetic field of $0 - 9 \text{ T}$ along the $[001]$ direction	176

A.1	Zero field thermometer calibration.	190
A.2	Thermometer calibration in applied magnetic field.	192
B.1	Structural Bragg peak measured using white beam neutron scattering	195
C.1	X-ray Laue snapshot of a crystal of $\text{Yb}_2\text{Ti}_2\text{O}_7$	197
D.1	Single crystal magnetisation measurements at $T = 2.0$ K	203
D.2	Mean-field magnetisation calculations compared to single crystal magnetisation at 90 mK and 2 K	204
D.3	Low temperature single crystal magnetisation, experimental and extrapolated	207
E.1	Non-magnetic background energy ranges	210
E.2	Non-magnetic background data sets	212

Chapter 1

Introduction

1.1 The Origin of Magnetic Moments in Insulators

In insulators, the electrons of the ions that form the materials exist in localised orbitals with limited spatial extent. Ions with unfilled electronic orbitals can have permanent magnetic moments, and interactions between these moments can give rise to a wide variety of collective magnetic behaviours. The magnetic moments of ions have two sources: the orbital angular momentum of the unpaired electrons of the ion, $\mathbf{L}$, which arises from the motion of the electrons within their orbitals, and the inherent spin angular momentum of the electrons, $\mathbf{S}$. The size of the spin and angular momentum of the free ion is determined by Hund's rules.

In the material of interest in this thesis, $\text{Yb}_2\text{Ti}_2\text{O}_7$, the only ions in the material with magnetic moments are the rare earth ions Yb^{3+} , where the magnetic moment arises from an unfilled $4f$ shell. In Yb^{3+} , strong spin orbit coupling $\lambda \mathbf{L} \cdot \mathbf{S}$, which arises due to interaction between the electron spin and the orbital part of the electron wave function, means that $\mathbf{L}$ and $\mathbf{S}$ are strongly coupled, leading to a total angular momentum $\mathbf{J} = \mathbf{L} + \mathbf{S}$ [1]. The magnetic moment due to the total angular momentum is $\mu_J = -g_J \mu_B \mathbf{J}$, where the $-$ sign comes about because the electron has negative charge, and $g_J = 3/2 + (S(S+1) - L(L+1)) / (2J(J+1))$ is the Landé g factor.

The electrons in the in partially filled $4f$ shell of Yb^{3+} are strongly influenced by the electrostatic field of the surrounding O^{2-} ions, which splits the 8-fold degenerate free ion $^2F_{7/2}$ ground state into four Kramers doublets. The lowest energy of these doublets is energetically well separated from the next highest energy doublet ($\Delta \sim 600$ K), so that it can be treated as an effective spin-1/2, where the two states of the doublet are treated as the up and down states of the spin-1/2. In order to make this approximation, the anisotropy of the ground state doublet is abstracted into an anisotropic $\mathbf{g}$ tensor, such that in the ground state doublet $\mathbf{J} \equiv \mathbf{g}\mathbf{S}/g_J$. The crystal field of $\text{Yb}_2\text{Ti}_2\text{O}_7$, and the anisotropy of the ground state doublet are discussed in greater detail in Ch. 1.8.4.

1.2 Magnetic Exchange Interactions

In most magnetic materials, the magnetic moments order below a temperature T_c , when the thermal energy drops below the scale of the magnetic interactions. In some of these materials the ordering temperature is much higher than can be reasonably explained if the magnetic moments are interacting only via dipolar interactions, for example elemental Iron has an ordering temperature of $T_c \sim 1000$ K. It has been shown that a combination of electrostatic repulsion and Pauli exclusion between electrons in overlapping orbitals gives rise to an effective interactions between magnetic moments, referred to as magnetic exchange [1]. The exchange interaction between two spins $\mathbf{S}_1$ and $\mathbf{S}_2$ results in an effective interaction Hamiltonian

$$H = E_0 - 2J\mathbf{S}_1 \cdot \mathbf{S}_2, \quad (1.1)$$

where J is an exchange integral, also referred to as an exchange energy. If $J > 0$, the spins tend to align and the interaction is referred to as ferromagnetic, while if $J < 0$, the spins anti-align and the interaction is referred to as antiferromagnetic. This isotropic form of the exchange interaction involving the dot product of neighbouring

spins is referred to as Heisenberg exchange. Magnetic exchange is inherently short ranged, as it relies on the overlap of electron orbitals, which are limited in extent. The dominant form the exchange interaction expected between rare earth ions with unpaired $4f$ electrons is superexchange, which relies on electrons hopping between the orbitals of magnetic ions via an orbital of an intermediate non-magnetic ion. In the case of $\text{Yb}_2\text{Ti}_2\text{O}_7$, the intermediate ion is O^{2-} .

Evidence of magnetic interactions in materials appears in the high temperature magnetic susceptibility χ , where the presence of interactions modifies the susceptibility from the free ion Curie form $\chi = C/T$ into the form $\chi = \frac{C}{T - \theta_{\text{CW}}}$. Here θ_{CW} is referred to as the Curie-Weiss temperature, and reflects the strength and sign of the magnetic interactions. If θ_{CW} is positive, then the interactions between the magnetic moments are net ferromagnetic, while if θ_{CW} is negative, the interactions are net antiferromagnetic. Take for example the 3D cubic ferromagnet, where the magnetic moments occupy a the sites of a cubic lattice. In this case, the magnetic susceptibility diverges as T tends towards θ_{CW} from above, and at $T = \theta_{\text{CW}}$ the system undergoes a phase transition from a paramagnetic phase to a phase with spontaneous magnetic order, where the magnetic moments break the rotational symmetry and spontaneously select a direction along which they all align, so that the system develops a net magnetic moment. In this case the ordering temperature is $\theta_{\text{CW}} = 2zJS(S + 1)/3k_B$ [1], where S is the size of the spins on the lattice, z is the number of nearest neighbours (in this case 6), and k_B is Boltzmann's constant. The direct proportionality of θ_{CW} to J reflects how the magnetic interactions affect the ordering temperature of the system.

1.3 Magnetic Order

As discussed in the previous section, as the temperature decreases, interactions between the magnetic moments can lead to the onset of spontaneous magnetic long range order below a temperature on the scale of the magnetic interactions. A long

range ordered state is one where the magnetic moments adopt a repeating pattern that minimises the interactions. Figure 1.1 shows the repeating pattern that minimises the magnetic interactions of the square lattice antiferromagnet, described by the exchange term $-JS_i \cdot S_j$ for $J < 0$, where neighbouring spins S_i and S_j want to point anti-parallel. This state is referred to as the Néel state. In practise the sublattice magnetization of the Néel state is reduced from the full saturated value S , by an amount ΔS , due to zero point quantum fluctuations.




Figure 1.1: The Néel mean-field ground state of the square lattice antiferromagnet, where neighbouring magnetic moments are anti-aligned. The shaded square shows the magnetic unit cell.

Magnetic long range order is characterised by an ordering wave vector $\mathbf{Q}$ that describes the periodicity of the repeating pattern made by the magnetic moments in the ordered phase. The most common experimental method for studying magnetic order is neutron diffraction, which probes the magnetic order in reciprocal space. In such experiments magnetic long range order is revealed by the presence of intense Bragg peaks in the neutron diffraction pattern, which reflect the periodicity of the magnetic order. For example, the mean-field ground state on the square lattice antiferromagnet shown in Fig. 1.1, has ordering wave vector $\mathbf{Q} = (\frac{1}{2}, \frac{1}{2})$, as the ordered pattern of the magnetic moments has double the periodicity of the structural lattice. This means the Bragg peaks will appear in elastic neutron scattering experiments at fractional

coordinates in the first Brillouin zone. Ferromagnetic order on the other hand has an ordering wave vector $\mathbf{Q} = \mathbf{0}$, because it has the same periodicity as the structural cell. This means that the magnetic Bragg peaks will occur below T_c , on top of the structural Bragg peaks already present above T_c .

1.4 Mean-Field Theory

The transition to spontaneous magnetic order below a temperature T_c on the scale of the interaction strength can be understood within the mean-field approximation. It assumes that the effect of interactions between a central spin and its neighbours is treated using an interaction field, the strength of which is based on the thermally averaged values of the neighbouring magnetic moments [1]. This simplifies the problem from one with many body interactions to an effective single spin problem. The average field is used to compute the thermal average of the moment of the central spin, and the resulting problem is solved self consistently to determine the average magnetic moments of all of the ions in the system. At $T = T_c$ the net moment of the system will be zero, but as the temperature decreases to zero, there are no longer any thermal fluctuations, and the magnetic moments of the ions will saturate. The resulting zero temperature mean field state of the magnetic moments will minimise the magnetic interactions between the moments, thus providing a systematic method for determining the lowest energy configuration of the magnetic moment, neglecting the presence of any quantum fluctuations. One flaw in mean-field theory is that it will always yield a finite value of T_c , due to the neglect of fluctuations in the magnetic moments, which in some materials can completely suppress magnetic long-range order.

1.5 Magnetic Frustration

In some materials it is impossible for the magnetic moments to simultaneously minimise all of their pairwise interactions. This is referred to as magnetic frustration. Frustration can arise for different reasons, including competition between different interactions present in the system, and combinations of the interactions and lattice geometry of the system. The simplest example of the latter form of frustration is the Ising antiferromagnet on the triangular lattice. On the square lattice, Ising antiferromagnetism is not frustrated, as shown in Fig. 1.1, but by placing the spins on the corners of a triangle, it becomes impossible to minimise the interaction energy of all of the pairs of spins simultaneously. Figure 1.2(a) shows a single triangle from the lattice, with three Ising spins located on its corners. The Ising nature of the spins means they must point either up or down, and the presence of antiferromagnetic exchange between the spins means that in order to minimise the interaction energy neighbouring spins must point anti-parallel to each other. Starting with spin 1 in Fig. 1.2(a), it is possible to go around the triangle, aligning the spins such that they are anti-parallel, until the bond between sites 3 and 1 is reached, where both spins' orientations have already been determined by other pairwise interactions, and they are aligned parallel to each other. This will lead to the exchange interaction on that bond being maximised, not minimised. If spin 1 is flipped, the unsatisfied bond is simply moved, showing that it is impossible to simultaneously satisfy all of the pairwise magnetic interactions. This leads to a 6-fold degeneracy of the possible spin arrangements on the triangle that minimise the energy of the triangle. If this triangle is then tiled to form a triangular lattice, the degeneracy of the ground state increases even further. Fig. 1.2(b) shows the case of frustration due to competing interactions, through the example of Ising antiferromagnetism on the square lattice (shown by the solid bonds), which is not inherently frustrated, with next-nearest-neighbour antiferromagnetic interactions. These additional interactions (indicated by dashed bonds) are only satisfied if next-nearest-neighbour spins point in opposite directions, which would break the Néel state, causing the nearest-neighbour interactions to be

unsatisfied.

The massive degeneracy of the ground state of a system that can arise due to magnetic frustration leads to the breakdown of mean-field theory, as the system cannot choose between these energetically equivalent states, so that the system never orders at low temperature. Instead, the system enters a correlated paramagnetic phase, where the system samples all of the energetically equivalent ground states.




Figure 1.2: (a) An illustration of Ising antiferromagnetism on a triangle. The arrows represent the direction of the magnetic moments on each corner of the triangle. The line styles of the bonds linking the corners of the triangular indicate whether the energy of Ising antiferromagnetic exchange is minimised (solid) or maximised (dashed) on that bond. (b) The square lattice antiferromagnet with next-nearest-neighbour interactions. The Néel phase which minimises nearest-neighbour antiferromagnetic exchange, with neighbouring spins pointing in opposite directions. The square lattice bonds are solid to reflect this energetic minimisation. Antiferromagnetic next-nearest-neighbour interactions (dashed bonds) are not satisfied by the Néel state, leading to frustration of the interactions.

1.6 Spin Ice

Another example of a model where magnetic frustration leads to exotic physics is spin ice. In spin ice, the magnetic moments are local Ising-like, and sit on a lattice of corner sharing tetrahedra (shown in Fig. 1.4(b)), where the local Ising directions are the $\langle 111 \rangle$ directions joining the centres of neighbouring tetrahedra. The local Ising nature

of the magnetic moments combined with ferromagnetic nearest-neighbour exchange is frustrated, and gives rise to a local energetic constraint on the orientations of the magnetic moments on each tetrahedra, that the four moments on each tetrahedron are configured such that two moments point into each tetrahedron and two moments point out (see Fig. 1.3(a)). On a single tetrahedron, this gives rise to six energetically equivalent configurations for the magnetic moments, when this constraint is applied to a lattice of corner sharing tetrahedra, the ground state become extensively degenerate, with a finite entropy at zero temperature. This set of constraints and the resulting entropy is equivalent to the case of proton disorder in water ice, hence why this magnetic phase is referred to as spin ice. Even though the ground state of spin ice is not long range ordered, the local constraints on the magnetic moments give rise to correlations, which are predicted to be manifested in the presence of pinch points at certain wave vectors in quasi-elastic neutron scattering measurements on these materials. The extensively degenerate spin ice ground state is also expected to host exotic excitations. If a single spin in the spin ice state is flipped, it gives rise to two tetrahedra that violate the two-in two-out energetic constraint, as shown in Fig. 1.3(b). If we think of each magnetic moment as a dipole, with positive magnetic charge on one of the moment, and a negative charge on the other, then in the ground state of spin ice each tetrahedron has neutral charge (see Fig. 1.3(c)), and the two tetrahedra the violate the spin ice constraints have net positive and net negative charge respectively. These tetrahedra with net magnetic charge behave as two fractional quasi-particles. Flipping additional spins allows the two defects to separate spatially, as shown in Fig. 1.3(e), connected by a “Dirac string” of flipped spins.

In classical spin ice materials, the monopole excitations do not propagate, they separate only by thermal diffusion or move under the influence of external magnetic fields. It has been theoretically proposed that if perturbations are added to the Hamiltonian such that they preserve the key properties of the ground state manifold, namely the two-in two-out constraint, but allow for quantum spin flips between the in and out




Figure 1.3: This figure is taken from Castelnovo et al. [2]. (a) The ground state of spin ice on two neighbouring tetrahedra, two magnetic moments point into, and two magnetic moments point out of both tetrahedra. (b) A spin flip excitation in the ground state of spin ice, now one tetrahedron has three spins pointing in and one out, and the other tetrahedron has three spins pointing out and one in. (c) and (d) show the same magnetic configurations as (a) and (b), with the magnetic moments replaced by dipoles with a negative magnetic charge on one end (blue) and a positive magnetic charge (red) on the other. When the spin in (b) is flipped it gives rise to a net magnetic moment in each tetrahedron. It is these net moments that behave like magnetic monopoles. Flipping further spins allow the two monopoles to separate spatially, being connected by a Dirac string of flipped spins, shown in (e).

orientations of the magnetic moments, then an even more exotic magnetic phase could arise, so called “quantum spin ice” [3]. In quantum spin ice the spin flip terms create a quantum superposition of the degenerate classical spin ice states, and allow the magnetic monopoles to coherently hop (propagate) through the lattice and acquire a dispersion relation. Other even more exotic excitations have also been proposed emerge in quantum spin ice: photon-like excitations that arise from ring fluctuations between the degenerate spin configurations, which are expected to be gapless with linear dispersion; and electric monopole or vison excitations [3]. Experimental observation of propagating magnetic monopoles and photon excitations would be very exciting results, in particular evidence for propagating monopoles would constitute the first physical observation of fractionalised quantum excitations in a fully three dimensional real material.

The potential for the existence of such an exotic state in a real material is part of the motivation for studying in detail the magnetism in $\text{Yb}_2\text{Ti}_2\text{O}_7$. In the classical spin ice materials, $\text{Ho}_2\text{Ti}_2\text{O}_7$, $\text{Ho}_2\text{Sn}_2\text{O}_7$, and $\text{Dy}_2\text{Ti}_2\text{O}_7$ [4], the ground state doublet of the crystal field does not allow for quantum spin flips between the up and down states of the local Ising moments, these materials are too close to the Ising limit. In $\text{Yb}_2\text{Ti}_2\text{O}_7$ the ground state doublet of the crystal field is not pure Ising, but rather XY-like [5–8], allowing for quantum spin flips to occur. Additionally, the symmetry allowed exchange Hamiltonian for $\text{Yb}_2\text{Ti}_2\text{O}_7$ also contains terms that will drive quantum spin flips (see Eqn. (6.3) in Ch. 6.1), making $\text{Yb}_2\text{Ti}_2\text{O}_7$ one of the strongest candidates in the literature to be a “real world” quantum spin ice.

1.7 The Pyrochlore Lattice

$\text{Yb}_2\text{Ti}_2\text{O}_7$ belongs to a family of materials known as the rare earth pyrochlores, which also includes the spin ice materials, described by the generic chemical formula $\text{A}_2\text{B}_2\text{O}_7$, where the A and B sites form interpenetrating lattices of corner sharing tetrahedra, as shown in Fig. 1.4(a). The crystal structure is described by the space group

$Fd\bar{3}m$ In the case of $\text{Yb}_2\text{Ti}_2\text{O}_7$, Yb^{3+} occupies the A site, while the B site is occupied by non-magnetic Ti^{4+} . Other members of this family substitute other rare earth ions on to the A site, and different non-magnetic ions on the B site (e.g. Sn, Ge). The lattice of corner sharing tetrahedra occupied by the A sites is shown in Fig. 1.4(b), and can be constructed by attaching a tetrahedral sublattice to each site of an FCC Bravais lattice. The positions of the sublattice sites relative to the FCC lattice sites are $r_1 = \frac{a}{8}(1, 1, 1)$, $r_2 = \frac{a}{8}(-1, -1, 1)$, $r_3 = \frac{a}{8}(-1, 1, -1)$, and $r_4 = \frac{a}{8}(1, -1, -1)$.




Figure 1.4: (a) The pyrochlore lattice formed by compounds with the chemical formula $\text{A}_2\text{B}_2\text{O}_7$. The A sites (red spheres) and B sites (blue spheres) form interpenetrating lattices of corner sharing tetrahedra. The green spheres show the position of O^{2-} ions. (b) The lattice of corner sharing tetrahedra formed by the A site, occupied by the Yb^{3+} ions in $\text{Yb}_2\text{Ti}_2\text{O}_7$. The numbers show the sublattice indices used in Chapter 6.

1.8 Previous Work on $\text{Yb}_2\text{Ti}_2\text{O}_7$

$\text{Yb}_2\text{Ti}_2\text{O}_7$ is a material that has been studied starting with the work of Blöte et al. in 1969 [9]. This early work [9] reported the heat capacity and susceptibility of powder samples, shown in Fig. 1.6(a), showing evidence of a first order phase transition at $T_c = 214$ mK and a broad feature at $T = 2$ K, of unknown origin. This work began a long line of papers [5–42] setting out to study the magnetic behaviour of this material,

and asking questions about the low temperature phase of this material that have still not been completely answered. This section seeks to give the reader an overview of the current state of the study of $\text{Yb}_2\text{Ti}_2\text{O}_7$. Rather than a chronological paper by paper review, a themed overview will be provided, discussing various aspects of the work that has been done on $\text{Yb}_2\text{Ti}_2\text{O}_7$, and providing all of the references relevant to that aspect of the investigation.

1.8.1 The Low Temperature Phase of $\text{Yb}_2\text{Ti}_2\text{O}_7$

Early work on $\text{Yb}_2\text{Ti}_2\text{O}_7$ by Blöte et al. [9] using heat capacity and magnetic susceptibility measurements on powder samples of $\text{Yb}_2\text{Ti}_2\text{O}_7$ indicated the presence of a sharp phase transition at $T_c = 214$ mK, and proposes net ferromagnetic interactions between the Yb^{3+} ions based on a positive Curie-Weiss temperature (θ_{CW}) of 0.4 K. Estimating the magnitude of the exchange energies from both the heat capacity and susceptibility they found that the predicted energy scales differed by an order of magnitude ($T_c \ll \theta_{\text{CW}}$), and thus provide the first indications that magnetic frustration may be at play in this material. Ever since this work, researchers have been attempting to explain the nature of the low temperature phase.

The first works on the low temperature phase by Hodges et al. [5, 15, 16] used muon spin resonance (μSR), Mössbauer absorption, and neutron scattering to probe the magnetism in $\text{Yb}_2\text{Ti}_2\text{O}_7$ as a function of temperature, finding that at a temperature near the observed peak in the heat capacity, the fluctuation rate of the magnetic moment observed slowed dramatically, but the spins did not long range order, as indicated by the shape of μSR asymmetry measurements, with no detectable additional neutron intensity at Bragg peak positions in their elastic neutron scattering measurements. This was surprising as it indicated that the low temperature phase may be more exotic than a long range ordered magnetic state. Shortly after this work, Yasui et al. [17] performed polarised neutron diffraction measurements on a single crystal of $\text{Yb}_2\text{Ti}_2\text{O}_7$, finding evidence of an increase in the temperature dependence

of the intensity of several structural Bragg peaks at a temperature consistent with a peak in the AC magnetic susceptibility, indicating that the transition was to some kind of long range ordered state, which they suggested was ferromagnetic in nature, though not a simple collinear ferromagnet. They also found that the magnetic intensity did not achieve its maximum at low temperatures unless the sample was left to equilibrate for times between 30 and 120 minutes, indicating that there was still significant dynamics in this ordered phase. The difference between the observations of Hodges et al. [5, 15, 16] and Yasui et al. [17] was the first example of what has become a complex issue in the physics of $\text{Yb}_2\text{Ti}_2\text{O}_7$, significant sample dependence in the low temperature physics of this material. This issue will be addressed in Section 1.8.2. For now, the various findings on the low temperature phase of $\text{Yb}_2\text{Ti}_2\text{O}_7$ will be discussed.

Subsequent to the works of Hodges et al. [5, 15, 16] and Yasui et al. [17], many other works have been published attempting to directly address the nature of the low temperature phase of $\text{Yb}_2\text{Ti}_2\text{O}_7$. These works can be split into two camps, those who find evidence of persistent magnetic dynamics, and those that find evidence of magnetic long range order.

Works Finding Evidence of No Magnetic Long Range Order at Low Temperature

In the persistent dynamics camp, in additions to Refs. [5, 15, 16], we find the work of Bonville et al. [18, 20], Gardner et al. [19], Dalmas de R  otier et al. [21], Ross et al. [22, 25], and D’Ortenzio et al. [36].

Bonville et al. [18, 20] provide a review of several of the experiments performed on $\text{Yb}_2\text{Ti}_2\text{O}_7$, along with some new AC susceptibility measurements which show evidence of a glass-like irreversibility in the field cooled and zero-field cooled susceptibility. They also present quasi-elastic neutron scattering measurements, which show evidence of the development of anisotropic correlations upon cooling their sample below $T = 25$ K, with no development of magnetic Bragg peaks, or indeed any fundamental change

in the quasi-elastic scattering at $T_c \sim 240$ mK. Based on their review of previous work, and these new measurements, they conclude that the low temperature phase of $\text{Yb}_2\text{Ti}_2\text{O}_7$ must be some kind of spin-liquid.

Inspired by the work of Yasui et al. [17], Gardner et al. [19] performed elastic neutron scattering and neutron spin echo (NSE) experiments on polycrystalline samples of $\text{Yb}_2\text{Ti}_2\text{O}_7$, finding evidence of ferromagnetic correlation in the form of increasing $S(Q)$ as Q approaches zero, and increased intensity of the (111) Bragg peak below T_c , but at the same time finding NSE signals, beam depolarisation effects, and lack of increased intensity on other Bragg peaks, inconsistent with the presence of long range ferromagnetic order.

Dalmas de R  otier et al. [21] present additional μSR measurements that show evidence of a magnetic environment that is very disordered and not obviously consistent with magnetic long range order. The spectrum they measure is described by a Gaussian broadened Gaussian model. In this model, the magnitude of the magnetic field experienced by the implanted muons is described by a series of gaussians whose mean values are described by a Gaussian distribution. This implies that the magnetic moments in the system are static on the time scale that the muon probes, but not long range ordered, as the muons do not experience a discrete number of fields due to different implantation sites as would be expected for magnetic long range order.

The works of Ross et al. [22, 25] present quasi-elastic neutron scattering showing rods of diffuse scattering intensity along the $\langle 111 \rangle$ directions in reciprocal space. They attribute these rods to two dimensional correlations existing in the Kagome planes that make up the lattice of corner sharing tetrahedra occupied by the Yb^{3+} ions in $\text{Yb}_2\text{Ti}_2\text{O}_7$. At temperatures below both of the observed peaks in the heat capacity of their sample, they find that these rods have not completely condensed into Bragg peaks, but have only coalesced into broad features around the Bragg peak locations, indicating three dimensional correlations, but not magnetic long range order. Ross et al. [25] fit cuts through the rods of scattering both along and across the rods, extracting from these cuts correlations lengths, which show that below the temperature

of the phase transition the magnetic correlations in the system go from correlations within a single Kagome star in one of the Kagome planes, to correlations over multiple Kagome stars in one plane, and between planes. Ross et al. [25] also present heat capacity measurements on powder and single crystal samples, shown in Fig. 1.6(b), showing that the heat capacity of the single crystal is very different from that of the stoichiometric powder. The single crystal used for the inelastic neutron scattering shows a small peak (compared the the size of the peak found in the data collected in the stoichiometric powder) at the same temperature ($T_c = 265$ mK) as the peak found in measurements on the powder sample, along with a broad feature at $T \sim 190$ mK. Another single crystal measured shows only the broad feature. This is the first indication in the literature that single crystals may behave very differently from powder samples of $\text{Yb}_2\text{Ti}_2\text{O}_7$ at low temperatures, presenting a potential explanation for the existence of so many contradictory set of data on $\text{Yb}_2\text{Ti}_2\text{O}_7$. This issue will be discussed further in Section 1.8.2.

D’Ortenzio et al. [36] present μSR measurements on both polycrystalline and single crystal samples of $\text{Yb}_2\text{Ti}_2\text{O}_7$. They find that even though the two types of samples show very different behaviours in their heat capacities at mK temperatures, they show similar features in their μSR signals in longitudinal field and transverse fields, though they are not identical and the powder data shows more distinct features. None of the observed signals are consistent with the presence of magnetic long range order, but are interpreted to mean that there are no static magnetic moments on the μSR time scale. It should be noted here that the existence of persistent low temperature dynamics in frustrated magnetic systems, of which $\text{Yb}_2\text{Ti}_2\text{O}_7$ is one, is an outstanding issue in the field, as they appear in a great deal of these types of materials. It is unclear to researchers whether these dynamics are the result of low lying excitation that only μSR is capable of detecting, or whether this is some kind of muon induced effect. This is an active field of investigation [43] so conclusions based on μSR data alone should not be taken as evidence of no long range order. However, when combined with all of the measurements discussed above, they form a large body

of evidence that in some samples of $\text{Yb}_2\text{Ti}_2\text{O}_7$ the low temperature phase may not be a long range ordered phase.

Works Finding Evidence of Magnetic Long Range Order at Low Temperature

In the long range ordered camp, in addition to the work of Yasui et al. [17], and using the same and similarly grown crystals, we find the recent works of Chang and collaborators [33, 37, 39], and the work of Bhattacharjee et al. [41]. The first group of works present evidence from polarised neutron scattering, magnetisation, susceptibility, heat capacity, and μSR measurements on powders and single crystals of $\text{Yb}_2\text{Ti}_2\text{O}_7$, that the transition observed in their samples is a first order phase transition to a state that is ferromagnetic in nature, while the other work uses ultrasound measurements to indirectly probe the magnetism.

Based on the observation of a definite onset of magnetic scattering at Bragg peak positions observed in polarised neutron scattering experiments that shows thermal hysteresis, and the onset of neutron depolarisation near $T_c \sim 250$ mK in a single crystal of $\text{Yb}_2\text{Ti}_2\text{O}_7$, Chang et al. [33] propose that their sample of $\text{Yb}_2\text{Ti}_2\text{O}_7$ is ferromagnetically ordered at temperatures below the observed peak in the heat capacity. Figure 1.5 shows the magnetic order Chang et al. [33] propose for this phase based on mean-field calculations using the exchange Hamiltonian they refine based on quasi-elastic neutron scattering. The supplemental material of Chang et al. [33] includes heat capacity measurements, showing that their sample, which displays evidence of magnetic long range order at low temperatures, also shows a single sharp peak in the heat capacity, indicating the presence of a single phase transition in the sample at $T_c \sim 195$ mK (see Fig. 1.6(d)).

Magnetisation measurements on single crystals and powders of $\text{Yb}_2\text{Ti}_2\text{O}_7$ by L’Hotel et al. [37] show thermal hysteresis and low temperature saturation at temperatures slightly below an observed sharp peak in the heat capacity of the samples. AC susceptibility measurements on powder samples show a sharp peak in both the in-phase

and out-of-phase components at the same temperature as the transition observed in the magnetisation [37]. It is notable that that transition temperature observed in the AC susceptibility and magnetisation are slightly below that of the peak in the heat capacity, corresponding to the temperature of a small shoulder observed in the data. The peak in the heat capacity corresponds to a small shoulder in the magnetisation, which they propose is due to sample inhomogeneity. The reason for this apparent disconnect between the temperature of the transition visible in the heat capacity, and the transition temperature inferred from magnetisation data is unclear.

The μ SR measurements of Chang et al. [39] are consistent with static (though not necessarily long range ordered [43]) magnetic moments in the low temperature phase.




Figure 1.5: The nature of the ferromagnetic order in $\text{Yb}_2\text{Ti}_2\text{O}_7$ proposed by Chang et al. [33] based on mean-field calculations using their exchange Hamiltonian. (a) The large arrows show the directions of the pseudo-spin- $1/2$ s. (b) The large arrows show the directions of the full magnetic moments. The thin red, green, and blue arrows show the local z , x , and y directions respectively. Figure is taken from Ref. [33].

Bhattacharjee et al. [41] present results of ultrasound velocity and attenuation measurement on single crystals of $\text{Yb}_2\text{Ti}_2\text{O}_7$. The major outcome of this work is that there is evidence of significant coupling between acoustic modes (phonons) and magnetism in $\text{Yb}_2\text{Ti}_2\text{O}_7$, suggesting that finite magnetoelastic coupling exists in this material. This may have some impact on the results of work on $\text{Yb}_2\text{Ti}_2\text{O}_7$ performed in high magnetic fields. Additionally they find that below (above) the temperature

of the broad peak in the heat capacity of their sample, the sound velocity and attenuation are monotonic (non-monotonic) as a function of applied magnetic field. The high temperature non-monotonic behaviour is consistent with the in field polarisation of a paramagnetic phase, where the observed minimum (maximum) in the ultrasound velocity (attenuation) indicates the field at which the system crosses over from disordered to polarised ferromagnet. The lack of any evidence of the onset of polarisation below the temperature of the transition observed in the heat capacity is interpreted to mean that below this temperature the system is already ferromagnetically ordered and cannot be further polarised. Both quantities are also hysteretic below T_c , consistent with the existence of a first order phase transition. Bhattacharjee et al. [41] also report heat capacity measurements on their sample of $\text{Yb}_2\text{Ti}_2\text{O}_7$, showing a peak with similar magnitude to that of Chang et al. [33], but broader in temperature.

Overall, works that find evidence of magnetic long range order tend to be more recent than works that find evidence of persistent magnetic disorder in the low temperature phase of $\text{Yb}_2\text{Ti}_2\text{O}_7$. More importantly, most of these works have been performed on crystals grown in a consistent manner [33, 37, 39], or that have consistent features in low temperature heat capacity measurements, which we shall see shortly, is important in understanding the low temperature phase of $\text{Yb}_2\text{Ti}_2\text{O}_7$.

1.8.2 Sample Dependence in $\text{Yb}_2\text{Ti}_2\text{O}_7$

As discussed in the Section 1.8.1, Ross et al. [25] presented the first direct evidence of significant sample dependence in the low temperature behaviour of $\text{Yb}_2\text{Ti}_2\text{O}_7$. Given the wide variety of experimental evidence supporting both ferromagnetically ordered and disordered phases below the temperature of observed peaks in the heat capacity of samples, some underlying cause must exist, and differences between the samples appears to be a logical cause. This section will discuss the many works which have subsequently explored this issue and the conclusions that have been drawn from them.

In addition to quasi-elastic neutron scattering and heat capacity results, Ross et al.

[25] also show results of elastic neutron scattering on powder and crushed single crystal samples. The data shows no indication of structural distortions, but the single crystal data shows a higher background and broader Bragg peaks, as well as a larger lattice parameter, indicating the presence of disorder and strain. This is the first indication in the literature that disorder issues in single crystal samples can be associated with differences in the low temperature behaviour of single crystal samples, specifically, multiple peaks, or broad features at very low temperature in the heat capacity, as shown in Fig. 1.6(b).

The work of Yaouanc et al. [26] presents further heat capacity measurements on single crystal and powder samples of $\text{Yb}_2\text{Ti}_2\text{O}_7$ shown in Fig. 1.6(c). One of the single crystals used in the heat capacity measurements was previously used for the measurements of Bonville et al. [18, 20], that found evidence of persistent magnetic dynamics and no long range order in the low temperature phase. They also show the heat capacity of a single crystal grown using the travelling flux method. All of the single crystal samples measured by Yaouanc et al. [26] show broad features at $T_c \sim 170$ mK, as opposed to their powder sample, which shows a sharp peak at $T_c \sim 250$ mK. Based on these results, Yaouanc et al. suggest that any results from measurements on single crystal samples of $\text{Yb}_2\text{Ti}_2\text{O}_7$ should be examined cautiously. This work adds further to the idea that samples which show broad features in the heat capacity at mK temperatures show evidence of persistent magnetic dynamics and no long range order at low temperatures.

Chang et al. [33], who as discussed previously, find evidence of long range order in the low temperature phase of a single crystal sample of $\text{Yb}_2\text{Ti}_2\text{O}_7$ present heat capacity measurements on various single crystals in the supplementary material to their work, as shown in Fig. 1.6(d). They find that the heat capacity of the single crystal used in their neutron scattering measurements (the same sample used in Yasui et al. [17]) shows a single sharp peak (though not as sharp as that of powders) at $T_c \sim 195$ mK. They also show that crystals grown using the same optical floating zone (OFZ) method but with faster growth rates (the sample showing a single sharp peak




Figure 1.6: Heat capacity measurements on various powder and single crystal samples of $\text{Yb}_2\text{Ti}_2\text{O}_7$ showing the degree of variation in results across the literature. (a) The original results of Blöte et al. [9] where they found a single sharp peak at $T_c = 214$ mK. (b) The results of Ross et al. [25] showing the marked difference in the heat capacity between powder and single crystal samples. (c) The results of Yaouanc et al. [26] who found that all of their single crystal samples had much smaller and broader peaks in their heat capacity compared to powder samples. (d) shows the results of Chang et al. [33] showing the first single crystal samples to display a single sharp peak in their heat capacity, along with samples that displayed no such peak. (e) shows the results of Ross et al. [34] showing the heat capacity of as-grown single crystal samples, and single crystal samples annealed under oxygen overpressure as compared to that of a powder sample. (f) shows the results of D'Ortenzio et al. [36] who found that their single crystal sample showed a broad peak at the same temperature as the sharp peak in the crystal of Chang et al. [33]. All figures have been taken directly from the cited works. Not included are heat capacity measurements reported by Robert et al. [40], and L'Hotel et al. [37]. The results of Robert et al. [40] are similar to those of Yaouanc et al. [26], while the results of L'Hotel et al. [37] are from measurements on the same crystal as Chang et al. [33].

was grown at the exceptionally slow rate of 1.5 mm/hr) or unstable melt regions do not show sharp peaks, suggesting that growth conditions and method are important to the cooperative magnetic behaviour, and showing that they have grown the first single crystal in the literature to show behaviour similar to that of stoichiometric powders. They also associated the heat capacity results with EXAFS (Extended X-ray absorption fine structure) measurements which show that samples that have either a reduced broadened feature (compared to the sharp peak observed in their primary sample) or no feature in the heat capacity show evidence of increased vacancies, primarily of the O^{2-} ions in the centre of the Yb^{3+} tetrahedra, but also of Yb^{3+} and Ti^{4+} in order to achieve charge neutrality.

Ross et al. [34] present the results of heat capacity measurements and of detailed analysis of elastic neutron scattering measurements on powder and crushed single crystal samples of $Yb_2Ti_2O_7$, in an attempt to determine the nature of any structural or chemical differences between the two. Their heat capacity measurements are shown in Fig. 1.6(e), which shows that their powder sample displays a sharp single peak at $T_c \sim 250$ mK while the single crystals studied (one as grown using OFZ methods, the other annealed in oxygen) show two peaks, one at the same temperature as the peak in the powder sample heat capacity, the other at lower T . These results show that oxygen annealing has little effect on the heat capacity of single crystal samples. Reitveld refinement of the neutron scattering measurements shows that both the powder and single crystal samples are described by the $Fd\bar{3}m$ pyrochlore structure, but the single crystal has a slightly larger lattice parameter than the powder, and shows evidence of Yb substitution on the Ti site (but not vice versa) at the level of 2.3%. Based on the results of all of their measurements, Ross et al. conclude that the difference between stoichiometric powder and single crystals of $Yb_2Ti_2O_7$ that show different low temperature heat capacities is not Oxygen off-stoichiometry, as had been assumed in the past, but substitution of Yb on to the Ti sites in the crystals. This is not consistent with the findings of Chang et al. [33], but the sample of Chang et al. [33] was grown in standard atmospheric conditions, while the sample of Ross et al.

[34] was grown in an overpressure of oxygen, which may lead to a reduction in O^{2-} deficiencies.

They suggest that the substitution issues they uncover may explain the differences observed between various samples of $Yb_2Ti_2O_7$. At first glance a change in the local structure of the crystal on the order of 2% does not seem like a reason for a large change in the low temperature behaviour of $Yb_2Ti_2O_7$ from one where magnetic long range order is present to one where there is no obvious long range order, but as we shall see later, work on the exchange Hamiltonian of $Yb_2Ti_2O_7$ puts it very close to a Quantum Phase Transition achieved by varying the symmetry allowed exchange parameters, which could allow for a 2% change in the local structure of the crystal to drive the system across this transition, or create strong quantum fluctuations.

D’Ortenzio et al. [36] present heat capacity measurements on single crystal and polycrystalline samples of $Yb_2Ti_2O_7$. The polycrystalline samples were previously used in the μ SR experiments of Dalmas de R  otier et al. [21]. Figure 1.6(f) shows the results of their heat capacity measurements; their polycrystalline samples show a single sharp peak at $T_c \sim 265$ mK, consistent with other powder samples of $Yb_2Ti_2O_7$, while their single crystal data shows a broad peak near $T_c \sim 185$ mK. Given that they state that their μ SR measurements on these samples are not consistent with magnetic long range order at low temperatures, the existence of a broad feature in the single crystal heat capacity is consistent with other works.

Finally, Chang et al. [39] state that the heat capacity of their single crystal samples is smaller than those of Yaouanc et al. [26] and D’Ortenzio et al. [36] below the peak in the heat capacity, perhaps indicating that their sample is less disordered magnetically at low temperatures. This is consistent with the finding of that work, that μ SR measurements on single crystal samples show evidence consistent with magnetic long range order. The single crystal heat capacity results in this work are not new, they are taken from L’Hotel et al. [37], but they compare them to the results of Yaouanc et al. [26] and D’Ortenzio et al. [36].

Taken as a body, all of the works on sample dependence issues seem to indicate

that samples with measurable levels of defects are more likely to show evidence of persistent magnetic dynamics and a lack of magnetic long range order in the low temperature phase. These samples show broad features in their low temperature heat capacity, such features tend to be at lower temperatures compared to the sharp peaks observed in polycrystalline (powder) samples, and some single crystals. Conversely, single crystal samples with reduced defects compared to samples with broad features in their heat capacity, as confirmed by techniques like EXAFS [33], tend to show evidence of magnetic long range order at mK temperatures, and also show sharp peaks in their low temperature heat capacity, consistent with a single first order transition. This is important to note, as in this thesis, low temperature heat capacity is used to position the sample used with respect to the rest of the literature.

1.8.3 The Exchange Hamiltonian of $\text{Yb}_2\text{Ti}_2\text{O}_7$

In an effort to understand the physics behind the observed phase transition to the low temperature phase of $\text{Yb}_2\text{Ti}_2\text{O}_7$, several works have attempted to establish the nature of the magnetic exchange Hamiltonian. This section will provide a brief summary of the very early work on determining the nature of the magnetic interactions in $\text{Yb}_2\text{Ti}_2\text{O}_7$, and a more detailed discussion of recent works on the subject.

Early work on the heat capacity and susceptibility of $\text{Yb}_2\text{Ti}_2\text{O}_7$ used fits to these quantities to estimate the nature and magnitude of the exchange interactions in $\text{Yb}_2\text{Ti}_2\text{O}_7$. Blöte et al. [9] estimated that if the exchange interactions in $\text{Yb}_2\text{Ti}_2\text{O}_7$ were isotropic, then the nearest-neighbour exchange would be ferromagnetic, with an exchange energy of either $J = 0.026$ meV, estimated from susceptibility, or $J = 0.1 - 0.12$ meV, estimated from the heat capacity. This was backed up by a ferromagnetic Curie-Weiss temperature (the temperature axis intercept of a linear fit to the inverse magnetic susceptibility) of $\theta_{\text{CW}} = 0.4$ K, which indicated the presence of net ferromagnetic interactions. Since this work, many others have reported values of θ_{CW} [7, 14, 15, 17, 18, 33, 36, 40], varying between $\theta_{\text{CW}} = 0.4$ K and $\theta_{\text{CW}} = 0.8$

K, confirming that the interactions in $\text{Yb}_2\text{Ti}_2\text{O}_7$ are on the whole ferromagnetic in nature.

Heisenberg ferromagnetism on the pyrochlore lattice is not frustrated, and so could not explain the suppression of T_c compared to θ_{CW} . Furthermore, Blöte et al. [9] suggested that the exchange interactions in $\text{Yb}_2\text{Ti}_2\text{O}_7$ may be anisotropic, as they could not obtain consistent values of J by fitting heat capacity and magnetic susceptibility measurements. One aspect this early work neglected was the local symmetry of the Yb^{3+} site in $\text{Yb}_2\text{Ti}_2\text{O}_7$. As discussed in Section 1.8.4, the local environment of the Yb^{3+} site leads to the magnetic moments of the Yb^{3+} ions preferring to lie in the local easy plane, perpendicular to the local $\langle 111 \rangle$ directions at each corner of the tetrahedra that make up the lattice of Yb^{3+} ions in $\text{Yb}_2\text{Ti}_2\text{O}_7$. This tendency is reflected in the anisotropic $\mathbf{g}$ tensor with $g_{\perp} \gg g_{\parallel}$ (see Eqn. (1.4)).

Hodges et al. [5] and Yasui et al. [17] attempted to gain some insight into the exchange interactions in $\text{Yb}_2\text{Ti}_2\text{O}_7$ (while accommodating the local planar nature of the ground state crystal field doublet), by fitting bulk magnetisation measurements. Hodges et al. [5] found that taking into account crystal field effects was not sufficient to fit their magnetisation measurements, so they incorporated an “interionic coupling tensor” to their magnetisation model. This tensor is simply an estimate of the mean magnetic field experienced by one Yb^{3+} ion due to interactions with its neighbours. They state that it is likely that the coupling tensor should be anisotropic, but the magnetisation does not provide sufficient information to determine the anisotropy, so they assume it is isotropic. Both Hodges et al. [5] and Yasui et al. [17] include this isotropic tensor, concluding that it is ferromagnetic in nature.

The first works to attempt to take into account anisotropic interaction were by Cao et al. [7, 23]. The work was based on measurements of the local susceptibility of single crystal samples, and bulk susceptibility of powder samples. The local susceptibility is a quantity extracted from polarised neutron scattering measurements. This is a local probe of the magnetic susceptibility, which avoids the problem that bulk measurements of the magnetic susceptibility lose access to any anisotropy in

the susceptibility due to averaging over the four $\langle 111 \rangle$ directions in $\text{Yb}_2\text{Ti}_2\text{O}_7$ (shown in Fig. 1.7), which define the local crystal field environments, and thus the local magnetic environment. Using this data Cao et al. [7, 23] attempted to recursively fit an anisotropic inter-ionic interaction tensor to the magnetic susceptibility using a mean-field model. They found that the resulting tensor is anisotropic, but they use a single sublattice model, which fails to respect the symmetry of the lattice [30]. Refs. [7, 23] do find that the exchange is local Ising-like, that is, the component of the interionic-coupling tensor parallel to the local $\langle 111 \rangle$ direction is much larger than the component perpendicular to the local $\langle 111 \rangle$ direction, the reversed anisotropy compared to the $\mathbf{g}$ tensor of the ground state crystal field doublet. Malkin et al. [24] attempted to use a model to fit the local susceptibility that respects the symmetry of the lattice, but they found that the resulting exchange couplings were unphysical.




Figure 1.7: Figure depicting the four local $\langle 111 \rangle$ axes on a single tetrahedron in $\text{Yb}_2\text{Ti}_2\text{O}_7$. The blue spheres represent the four Yb^{3+} ions that make up the tetrahedron. The red arrows show the directions of the four $\langle 111 \rangle$ axes, while the green planes show the local XY-planes. The lattice of corner sharing tetrahedra that the Yb^{3+} ions in $\text{Yb}_2\text{Ti}_2\text{O}_7$ occupy is constructed by placing tetrahedra like this at each site of an FCC lattice.

All of the previous works to determine the magnetic interactions in $\text{Yb}_2\text{Ti}_2\text{O}_7$ assumed a simple anisotropy in the exchange interaction, assuming different exchanges

for the components parallel and perpendicular to the local $\langle 111 \rangle$ directions. S. Curnoe [44] introduced the concept of symmetry allowed nearest-neighbour interactions in rare earth pyrochlore materials, showing that the most general exchange matrix between nearest-neighbour sites has four independent interaction terms. To fully describe the magnetic exchange interactions in $\text{Yb}_2\text{Ti}_2\text{O}_7$, it is necessary to use a model that accounts for all four of the symmetry allowed terms. These interactions can be expressed in many linearly independent forms [27, 29, 45]. The currently most popular form of expressing these interactions comes from McClarty et al. [46], and is expressed as:

$$\begin{aligned} \mathcal{H}_{\text{Exchange}} = \sum_{\langle i,j \rangle} & \left(J_{zz} \mathbf{S}_i^z \mathbf{S}_j^z - J_{\pm} (\mathbf{S}_i^+ \mathbf{S}_j^- + \mathbf{S}_i^- \mathbf{S}_j^+) + J_{\pm\pm} (\gamma_{ij} \mathbf{S}_i^+ \mathbf{S}_j^+ + \gamma_{ij}^* \mathbf{S}_i^- \mathbf{S}_j^-) \right. \\ & \left. + J_{z\pm} [\mathbf{S}_i^z (\zeta_{ij} \mathbf{S}_j^+ + \zeta_{ij}^* \mathbf{S}_j^-) + H.C.] \right), \end{aligned} \quad (1.2)$$

A detailed explanation of this Hamiltonian can be found in Chapter 6.1. Using this form of the magnetic exchange interactions, the exchange Hamiltonian can be completely specified by the four exchange terms $J_e = \{J_{zz}, J_{\pm}, J_{\pm\pm}, J_{z\pm}\}$.

Thompson et al. [29] was the first work to attempt to determine the strength of the four symmetry allowed interaction terms in $\text{Yb}_2\text{Ti}_2\text{O}_7$. They fit anisotropic quasi-elastic neutron scattering originally presented by Bonville et al. [18] to Random Phase Approximation calculations and demonstrated that the exchange interactions in $\text{Yb}_2\text{Ti}_2\text{O}_7$ are highly anisotropic. This work overestimated the magnitude of the exchange interactions, as it failed to account for the presence of an Onsager reaction field [47, 48], which leads to a significant correction to any quantity computed near the mean-field ordering temperature of the Hamiltonian. The exchange energies determined in Thompson et al. [29] were also used to calculate the local susceptibility for comparison to the data of Cao et al. [7, 23], achieving sufficient agreement with the experimental data [30], suggesting that local susceptibility is not a sufficiently sensitive quantity to constrain the magnetic exchange interactions.

The next work to extract values for the four symmetry allowed exchange energies was that of Ross et al. [27]. This work fitted spin wave dispersion relations observed

in inelastic neutron scattering measurements on a single crystal in high magnetic fields ("B" = 2 and 5 T) applied along the $[1\bar{1}0]$ direction, in order to extract values of the exchange energies. They found that the application of high magnetic fields was necessary as there are no sharp excitations of any kind visible in the zero field inelastic neutron scattering of $\text{Yb}_2\text{Ti}_2\text{O}_7$ [22, 27, 40]. The use of high magnetic fields in the experiment necessitates the inclusion of a Zeeman term in the Hamiltonian used to fit the data of the form:

$$\mathcal{H}_{\text{Zeeman}} = -\mu_0\mu_B \sum_{\mu} H^{\mu} \sum_{i,a} g_a^{\mu\nu} S_{ia}^{\nu}, \quad (1.3)$$

where H^{μ} is the x, y, z component of the magnetic field, $g_a^{\mu\nu}$ are the global Cartesian components of the $\mathbf{g}$ tensor described in Section 1.8.4, which characterises the ground state doublet of the crystal field states of the Yb^{3+} ion. The label a specifies the tetrahedral sublattice site that the effective spin-1/2 operator S_{ia} occupies. The label i specifies the FCC lattice site to which the sublattice is attached. By fitting the spin waves observed in an applied field of "B" = 5 T, assuming fixed $g_{\perp}/g_{\parallel} = 2.4$, Ross et al. [27] extracted exchange energies of $J_{zz} = 0.17(4)$ meV, $J_{\pm} = 0.05(1)$ meV, $J_{\pm\pm} = 0.05(1)$ meV, and $J_{z\pm} = -0.14(1)$ meV, along with a value for $g_{\parallel}$ of 1.8. The two dominant exchange terms are J_{zz} and $J_{z\pm}$. The mean-field phase diagram of this Hamiltonian as a function of temperature and applied magnetic field has a large region of ferromagnetic order (with a finite magnetisation along one the $\langle 100 \rangle$ directions), with $T_c = 3.2$ K in zero applied field. This mean-field phase diagram does not agree with the observed physics of $\text{Yb}_2\text{Ti}_2\text{O}_7$ at low temperatures, as discussed previously. Ross et al. [27] then discuss the potential impact of significant quantum fluctuations, driven by the XY nature of the ground state CEF doublet, and the other three terms in the exchange Hamiltonian. The exchange Hamiltonian of Eqn. 1.2 is known to support a Quantum Spin Liquid state when $J_{\pm}$ and $J_{\pm\pm}$ are small and $J_{z\pm} = 0$ [49]. This state is a quantum analog of classical spin ice, a state that supports magnetic monopole like excitations in the form of single spin flips. This Quantum

Spin Ice phase has been theoretically proposed to support rather exotic excitations, such as so called magnetic monopoles (like classical spin ice, see Fig. 1.3), along with dual electric monopoles, and photon-like excitations [27]. The exact exchange parameters of Ross et al. [27] are outside the region of the $\{J_{zz}, J_{\pm}, J_{\pm\pm}, J_{z\pm}\}$ phase diagram that can be expected to support a quantum spin liquid phase based on a perturbative approach ($J_{zz} \gg J_{\pm}, J_{\pm\pm}, J_z$), and rather would indicate the system should be in a state in which the magnetic moments order with a net moment along one of the $\langle 100 \rangle$ directions [27]. Ross et al. [27] suggest that the term $J_{z\pm}$ may help to stabilise the quantum spin liquid phase over larger areas of the phase diagram, but this idea is not explored quantitatively. The fact that J_{zz} is among the strongest terms in the exchange Hamiltonian of Ross et al. [27] has led to the proposal that $\text{Yb}_2\text{Ti}_2\text{O}_7$ could be a candidate “Quantum Spin Ice” material with exotic excitations [31, 32, 38, 41, 50–52].

The ability of the exchange Hamiltonian of Ross et al. [27] to successfully compute experimentally measured properties of $\text{Yb}_2\text{Ti}_2\text{O}_7$ has been tested by two recent works: Applegate et al. [32] and Hayre et al. [31]. These works use non-local cluster methods to calculate the heat capacity, entropy, and magnetisation of a material with the exchange Hamiltonian of Ross et al. [27]. Applegate et al. [32] computed the heat capacity and entropy at high temperatures $T > 700$ mK, and found that the exchange Hamiltonian of Ross et al. [27] gave a high temperature heat capacity in good agreement with the data of Blöte et al. [9], yielding a broad feature centred around $T = 2$ K. This would indicate that if $\text{Yb}_2\text{Ti}_2\text{O}_7$ is a Quantum Spin Ice, the high temperature feature in the heat capacity is similar in origin to the broad feature observed in the classical spin ice materials [53], which is well fitted by calculation based on classical magnetic monopoles. Applegate et al. [32] also compute the high temperature heat capacity based on the exchange Hamiltonian of Thompson et al. [29], finding that it yields a broad feature in the heat capacity at $T \sim 1$ K, which does not fit well with experiments. Hayre et al. [31] compute the magnetisation, heat capacity, and entropy of the exchange Hamiltonian of Ross et al. [27] using the same method as in

Applegate et al. [32], adding in the ability to compute these quantities in an applied magnetic field. Their results show that the exchange Hamiltonian of Ross et al. [27] explains the high temperature ($T > 1.8$ K) magnetisation of $\text{Yb}_2\text{Ti}_2\text{O}_7$ in zero field, and in applied fields between $0.2 \leq B < 5$ T. They also performed calculations of the heat capacity of $\text{Yb}_2\text{Ti}_2\text{O}_7$ in applied magnetic fields in the range $0.2 \leq B < 5$ T. These show a broad feature, the peak temperature of which shifts up in temperature upon increasing the applied field.

Other aspects of the Hamiltonian model proposed by Ross et al. [27] have also been explored theoretically. Savary et al. [54] explored the phase diagram of the exchange Hamiltonian starting from a dominant J_{zz} interaction, exploring the effects of $J_{\pm}$ and $J_{z\pm}$ using gauge mean-field theory. They found that the phase diagram as a function of these two additional interaction terms contains four phases: a ferromagnetic phase, an antiferromagnetic phase, a “Coulomb Ferromagnet” phase, and a Quantum Spin Liquid phase. The Coulomb Ferromagnet is an unconventional magnetic phase that displays ferromagnetic long range order, but has deconfined spinon and photon-like excitations. The spin liquid phase is that discussed in Ross et al. [27], which has magnetic and electric monopole and photon excitations, but no long range magnetic order. These two exotic phases persist only when $J_{z\pm} < 0.5J_{zz}$ and $J_{\pm} < 0.3J_{zz}$. Lee et al. [55] expanded on the range of parameters explored in Savary et al. [54], including finite values of $J_{\pm\pm}$, which they find creates interactions between spinons. In their analysis, Lee et al. [55] focused on the case of a non-Kramers ion, making most of the investigation not particularly relevant to $\text{Yb}_2\text{Ti}_2\text{O}_7$, as Yb^{3+} is a Kramers ion.

Wan et al. [56] explored the nature of potential excitations in a system where J_{zz} is the dominant term in the exchange Hamiltonian, and the other terms are treated perturbatively, assuming simplified lattice geometries. They find that in a system with a checker board lattice with only J_{zz} interactions and an applied magnetic field, the excitations are strings connecting pairs of magnetic monopoles, and the field introduces hopping of these strings. In the case of the full pyrochlore lattice and the exchange

Hamiltonian of Eqn. 1.2, under the constraint $J_{zz} \gg g\mu_B B \gg \{J_{z\pm}, J_{\pm}, J_{\pm\pm}\}$, where B is the strength of the magnetic field applied along $[001]$, the excitations are similar to those of the simpler checker board lattice problem. The additional terms in the exchange Hamiltonian serve to generate hopping of the strings. Wan et al. [56] acknowledge that the constraints they apply to the exchange energies are not consistent with those of $\text{Yb}_2\text{Ti}_2\text{O}_7$, but speculate that quantum strings may persist if the phase occupied by $\text{Yb}_2\text{Ti}_2\text{O}_7$ is continuously connected to the magnetised state they use as their ground state, which may be the case. Benton et al. [57] explored the photon excitations expected in a Quantum Spin Ice material, based on a simplified model where $J_{\pm\pm} = 0$, and $J_{zz} \gg J_{\pm}, J_{z\pm}$. They find that the photons evolve out of the pinch points observed in classical spin ice materials, and develop a bandwidth of $\Delta E \sim 4.0c/a_0$, where $c = 0.6 (12J_{\pm}^3/J_{zz}^2) a_0$ is the speed of light (photon velocity) in the material, and a_0 is the cubic unit cell lattice parameter.

Yan et al. [58] explore the full coupling parameter space of Eqn. 1.2 (though using a different convention for the exchange interactions) more fully, not using a perturbative approach, but rather by minimising the exchange energy between classical spins. They do restrict one of the interactions in their recast exchange Hamiltonian to be zero, the Dzyaloshinskii-Moriya [59] interaction. Yan et al. [58] find that in their three exchange parameter phase space, there are four stable $\mathbf{q} = 0$ states that minimise the energy of classical spins, a ferromagnetic phase, and three phases referred to as $\Psi_{1,2,3}$. Ψ_4 is an antiferromagnetic state where the spins lie in the local XY planes at each corner of a single tetrahedron, and also in the same global plane, known as the Palmer-Chalker phase [60]. Ψ_3 is a second coplanar antiferromagnetic phase selected via order-by-disorder from a manifold of classically equivalent states [61]. Ψ_2 is again an antiferromagnetic phase where the magnetic moments lie in the local XY planes, but they no longer share a global plane [62]. Within this classical phase diagram, Yan et al. [58] find that $\text{Yb}_2\text{Ti}_2\text{O}_7$ lies in the ferromagnetic phase very close to the phase boundary between the ferromagnetic phase and the Ψ_3 phase. They also propose an explanation for the dimensional reduction seen in the quasi-elastic

neutron scattering of Refs. [18, 20, 22, 25, 29, 40], that it arises from dimensionally reduced spin wave excitations related to the excitations of the Ψ_3 phase. They also comment that the quasi-elastic neutron scattering seen in samples of $\text{Yb}_2\text{Ti}_2\text{O}_7$ that do not show evidence of long range order appear to be more similar to the excitations in their calculations very close of the FM/ Ψ_3 phase boundary, rather than the scattering calculated using the Hamiltonian of Ross et al. [27]. Finally Yan et al. [58] explore the effects of quantum fluctuations on their phase diagram, but do not find it to especially impact on the case of $\text{Yb}_2\text{Ti}_2\text{O}_7$. Jaubert et al. [42] continue on from the work of Yan et al. [58], investigating the implications of the nearby Ψ_3 phase on the physics of $\text{Yb}_2\text{Ti}_2\text{O}_7$. They find that thermal and quantum fluctuations due to the close proximity of the Ψ_3 phase resist ferromagnetic order, and for certain values of the local Ising exchange term, multiple phase transitions as a function of temperature are possible, with the system first entering the Ψ_2 phase followed by a second phase transition to ferromagnetic order as the temperature is decreased. Jaubert et al. [42] state that this is consistent with the findings of L’Hotel et al. [37], making the questionable interpretation that a small bump in the magnetisation at the temperature of the sharp peak in the heat capacity is a separate transition from the lower temperature transition seen in magnetisation measurements with no coexisting peak in the heat capacity. Even when thermal and quantum fluctuations are taken into account at finite temperatures, the exchange coupling of Ross et al. [27] would seem to be far enough from the phase boundary that this would not occur. They also find that the application of a small magnetic field along an unspecified direction may increase T_c as seen in L’Hotel et al. [37] due to the field stabilising ferromagnetic order and destroying the Ψ_2 phase. Overall, it is likely that the nearby phase boundary in the magnetic exchange phase diagram is important to the physics of $\text{Yb}_2\text{Ti}_2\text{O}_7$, but it is not clear that the multiple phase transition picture is correct.

After the work of Ross et al. [27], other works have estimated the exchange couplings in $\text{Yb}_2\text{Ti}_2\text{O}_7$ from experimental measurements [33, 38, 40]. Chang et al. [33] attempted to extract exchange parameters from quasi-elastic neutron scattering

as was done by Thompson et al. [29], but suffered similar problems, in that they were fitting correlations at a temperature quite close to the mean-field ordering temperature ($T = 0.3$ K, $T_{\text{c,RPA}} = 0.29$ K in their case) of the Hamiltonian. This Hamiltonian is discussed in the supplemental material of Applegate et al. [32], where it is shown that non-local cluster heat capacity calculations based on this model place the high temperature bump in the heat capacity at a much lower temperature than is observed experimentally. It is also discussed in Hayre et al. [31], where they compute the magnetisation based of $\text{Yb}_2\text{Ti}_2\text{O}_7$ based on the Hamiltonian of Chang et al. [33], finding that it reproduces experimental measurements poorly. Pan et al. [38] refine the exchange interactions of Ross et al. [27] based on $\mathbf{Q} = 0$ THz spectroscopy measurements of the excitations in $\text{Yb}_2\text{Ti}_2\text{O}_7$ as a function of applied field. The only information they have to base this refinement on is the energy dependence of the $\mathbf{Q} = 0$ susceptibility (which can be extracted from THz measurements), which is highly unlikely to be completely constraining on the Hamiltonian. Pan et al. [38] search a region in exchange coupling space that spans twice the uncertainty in the exchange interactions of Ross et al. [27], and find $J_{zz} = 0.16$ meV, $J_{\pm} = 0.065$ meV, $J_{\pm\pm} = 0.02$ meV, and $J_{z\pm} = -0.134$ meV, very similar to the exchange interactions of Ross et al. [27], which is to be expected based on constraints used in their research. More interestingly Pan et al. [38] provide the first evidence of two magnon scattering in $\text{Yb}_2\text{Ti}_2\text{O}_7$, observing excitations whose energy increases at twice the rate of single magnon excitations as a function of applied field, though these excitations are sharp modes, and two magnon excitations should form a continuum of absorption, as the energy from the THz field is distributed between a pair of magnons.

Very recently, Robert et al. [40] refit the " B " = 2 T inelastic neutron scattering data of Ross et al. [27], along with their own quasi-elastic neutron scattering measurements, collected at $T = 0.05$ K. They find that the inelastic data of Ross et al. [27] is insufficiently constraining on the exchange Hamiltonian. They find that fits to the inelastic neutron scattering allow for a continuum of exchange Hamiltonians that satisfactorily describe the broad features of the inelastic neutron scattering,

with only small details of the excitation spectrum changing within their continuum of solutions. Within this continuum they find that there is an approximate linear relationship between the four terms in each set of exchange energies, allowing the space of four exchange energies to be reduced to a single parameter. After this reduction to a single variable problem, Robert et al. [40] perform calculations of the quasi-elastic neutron scattering pattern, which they measured, as a function of this single parameter and determine that a set of interactions equivalent to $J_{zz} = 0.07$ meV, $J_{\pm} = 0.085$ meV, $J_{\pm\pm} = 0.04$ meV, and $J_{z\pm} = -0.15$ meV best fit their data. In this parameterisation, J_{zz} is no longer the dominant exchange term, with $J_{z\pm}$ now being the exchange energy with the largest magnitude by a significant margin. Robert et al. [40] discuss the implications of their exchange Hamiltonian in terms of the work of Yan et al. [58] and Jaubert et al. [42] that place the exchange parameters of Ross et al. [27] very near the line of Quantum Phase Transitions in the phase diagram obtained by varying the energies associated with the four symmetry allowed exchange terms. The exchange parameters of Robert et al. [40] appear on this phase diagram even closer to this phase boundary, and they discuss the idea that competition between ferromagnetic order, predicted for the region of the phase diagram occupied by $\text{Yb}_2\text{Ti}_2\text{O}_7$, and the order in the nearby phase may explain the odd low temperature behaviour of $\text{Yb}_2\text{Ti}_2\text{O}_7$, its sample dependence, and even the strange continuum of inelastic scattering observed in their work and Ross et al. [22]. Robert et al. [40] also find that for their set of exchange parameters, the work of Yan et al. [58] and Jaubert et al. [42] indicate that the system should undergo three different phase transitions as a function of temperature in zero applied magnetic field. In addition to extracting and studying an exchange Hamiltonian for $\text{Yb}_2\text{Ti}_2\text{O}_7$ from experimental data Robert et al. [40] find that the anisotropic quasi-elastic neutron scattering studied in Refs. [18, 20, 22, 25, 29], which Refs. [22, 25] find begins to condense around Bragg peak locations, is unchanged at temperatures well below the temperature of the phase transition, suggesting that its source is not dimensional reduction, or if it is, this effect persists in the low temperature phase. Robert et al. [40] also find

that the continuum of scattering in zero magnetic field, which extends in energy up to 1.5 meV, persists at high temperatures ($T = 2$ K, the same temperature where a broad feature is observed in the heat capacity [9]), and is not a unique feature of the low temperature phase. They analyse this continuum based on the idea of splitting it into quasi-elastic and inelastic components, but this interpretation is dubious, as the energy ranges of the two components are similar. They attempt to explain the continuum of scattering using both ferromagnetic spin waves and classical dynamics, but neither provides a good fit to the data. They speculate that the continuum may arise either from fractionalisation of excitations (made more likely by the possibility of dimensional reduction) or from multi-magnon scattering, but do not explore either option [40].

Finally, Onoda et al. [28] performed purely first principles investigation of the exchange parameters in Eqn. 1.2 exploring the relative strengths of the four symmetry allowed interactions on the pyrochlore lattice based on transfer integrals between Oxygen and Ytterbium ion orbitals. As the sign and magnitude of these integrals change, they provide a range of interactions as a function of the ratio of two transfer integrals. This work does not fix the nature of the exchange interactions of $\text{Yb}_2\text{Ti}_2\text{O}_7$, but does provide a range of values for Hamiltonians extracted from experiments to be compared to.

To summarise, a great deal of work has been done to determine the exchange Hamiltonian that governs the interactions of the magnetic moments in $\text{Yb}_2\text{Ti}_2\text{O}_7$. All of this work focuses on pseudo-spin- $1/2$ models of the system due to Yb^{3+} being a Kramers ion with a well isolated ground state doublet. There are currently competing views on the nature of the exchange Hamiltonian, with two works finding very different exchange Hamiltonians, one with a dominant local Ising exchange term with other terms that would drive quantum fluctuations [27], and another that finds a term that drives quantum fluctuations to be dominant [40]. The exchange Hamiltonian with dominant Ising interactions has driven work on the idea of a Quantum Spin Ice being the explanation for the strange behaviour of $\text{Yb}_2\text{Ti}_2\text{O}_7$ [27, 54, 56], while other work

has focused on the proximity to a phase boundary in the exchange space driving the strange physics in $\text{Yb}_2\text{Ti}_2\text{O}_7$ [42, 58].

1.8.4 The Crystal Field of $\text{Yb}_2\text{Ti}_2\text{O}_7$

The magnetism in $\text{Yb}_2\text{Ti}_2\text{O}_7$ is due to the presence of Yb^{3+} ions, neither the Ti^{4+} , nor the O^{2-} ions in the system are magnetic. The valence electrons in the Yb^{3+} ion are in a $4f^{13} {}^2F_{7/2}$ configuration, meaning that the ion has $S = 1/2$, $L = 3$, and $J = 7/2$, with J being the good quantum number for the system. The magnetism of the Yb^{3+} ion will be strongly influenced by the surrounding ions in the crystal. In the case of $\text{Yb}_2\text{Ti}_2\text{O}_7$, the Yb^{3+} ion occupies the A site in the lattice, and is surrounded by a cage of O^{2-} ions in the form of a distorted cube, compressed along the local $[111]$ axis, as shown in Fig. 1.8. This gives the local environment of the Yb^{3+} ion D_{3d} point group symmetry [12]. The electric field resulting from the surrounding O^{2-} ions will split the ${}^2F_{7/2}$ eight-fold degenerate ground state. As Yb^{3+} has an odd number of valence electrons (half integer total spin), it is a Kramers ion. This means that the CEF states of the ion will be split into a series of doublets (there must be eight states in total as J_z spans the range $7/2 \dots -7/2$).

Determining the nature of these doublets with respect to their J_z character, especially the nature of the lowest energy doublet, is important to the low energy magnetism that is of interest in this system. Many works have sought to quantify the crystal field (CEF) of the Yb^{3+} ion in $\text{Yb}_2\text{Ti}_2\text{O}_7$, as the crystal field plays a key part in the magnetism of this material. Early works [11, 63] used an incorrect description of the local environment of the Yb^{3+} ion. The local environment of the Yb^{3+} is D_{3d} , and thus six Stevens Operator Equivalents [64] are required to correctly describe the crystal field Hamiltonian of the Yb^{3+} ion [12], Townsend et al. [63] and Dunlap et al. [11] used too few parameters. The next work on the CEF states, Sengupta et al. [12], used the correct set of crystal field parameters, but concluded based on the available data that the ground state doublet had an Ising-like character and consisted




Figure 1.8: The oxygen environment of a Yb^{3+} ion in $\text{Yb}_2\text{Ti}_2\text{O}_7$. The light blue sphere represents the Yb^{3+} ion, and the red spheres are O^{2-} ions. The blue vector represents the crystallographic $[111]$ axis.

mainly of $m_J = \pm 3/2$, which is inconsistent with a large body of more recent works [5, 6, 8, 13, 15, 17, 23] that find the ground state to be XY-like in character and to consist mainly of $m_J = \pm 1/2$. All of these works find that the lowest energy doublet is separated from the next excited doublet by an energy gap of 600 – 900 K (50 – 80 meV). This means that for the temperature and energy ranges of interest in this work, only the ground state doublet will be significantly populated at low temperatures.

As only the lowest energy doublet is important for the magnetism, the degrees of freedom can be described using an effective spin-1/2 model, where the two states in the doublet are treated as the spin-up and spin-down states of an effective spin-1/2. In order to use this description we must incorporate the magnetic anisotropy present in the ground state doublet via an anisotropic $\mathbf{g}$ tensor, which is defined by $\mathbf{J} \equiv \mathbf{g}\mathbf{S}/g_J$, where the equivalence holds for the projection of $\mathbf{J}$ onto the ground state doublet. $\mathbf{J}$ is the total angular momentum, $\mathbf{S}$ is the effective spin-1/2, and g_J is the Landé g-factor, which is 8/7 in the case of Yb^{3+} . The quantisation axis z of $\mathbf{J}$ in zero field is set as the local $\langle 111 \rangle$ axis at each corner of the tetrahedral sublattice. This direction changes for each Yb^{3+} ion, but if $\mathbf{g}$ is defined in terms of the local coordinates, it remains the same

for each Yb^{3+} ion. In this form, the $\mathbf{g}$ tensor is diagonal, and has only two components, the local z component, $g_{\parallel}$, and the local xy component, $g_{\perp}$. These components are defined as $g_{\parallel} = 2g_J \langle \Psi_0^1 | J_z^a | \Psi_0^1 \rangle$ and $g_{\perp} = 2g_J \langle \Psi_0^1 | J_x^a | \Psi_0^2 \rangle = i2g_J \langle \Psi_0^1 | J_y^a | \Psi_0^2 \rangle$, where $\Psi_0^{1,2}$ are the two states that make up the ground state doublet of the crystal field, and J_{α}^a are the local Cartesian components of the total angular momentum on sublattice site a . Using these two components, in local coordinates, the $\mathbf{g}$ tensor is written as:

$$\mathbf{g} = \begin{pmatrix} g_{\perp} & 0 & 0 \\ 0 & g_{\perp} & 0 \\ 0 & 0 & g_{\parallel} \end{pmatrix}. \quad (1.4)$$

Table 1.1 gives the $\mathbf{g}$ tensor components estimated from the most recent works on the ground state doublet. The average values of Table 1.1 are $g_{\parallel} = 2.1(3)$ and $g_{\perp} = 4.0(2)$, where the uncertainty is the standard deviation of the values. These indicate that the ground state doublet in all of these studies is XY-like in nature. All of the work in this thesis uses the effective spin-1/2 model of $\text{Yb}_2\text{Ti}_2\text{O}_7$. The XY-like nature of the ground state doublet is significant as it is responsible for a great deal of the current interest in $\text{Yb}_2\text{Ti}_2\text{O}_7$, as an XY-like ground state doublet allows for quantum fluctuations between the spin-up and spin-down states of the effective spin-1/2, as opposed to the Ising-like ground state doublet seen in the classical spin ice materials $\text{Ho}_2\text{Ti}_2\text{O}_7$ and $\text{Dy}_2\text{Ti}_2\text{O}_7$ where quantum fluctuations are suppressed. The ability for the system to support quantum fluctuations allows for the possibility of unconventional magnetic behaviour in this material.

1.8.5 Other Works on $\text{Yb}_2\text{Ti}_2\text{O}_7$

In addition to all of the works previously discussed in this chapter, there have been several recent works on $\text{Yb}_2\text{Ti}_2\text{O}_7$, the results of which do not fit in with the previous sections. These include further work by Pan et al. [51] where they attempt to extract monopole inertia from THz and microwave cavity measurements on a single crystal

Table 1.1: Anisotropic $\mathbf{g}$ tensor components ($g_{\parallel}$ and $g_{\perp}$), in the local $[111]$ frame, of the ground state CEF doublet from various works. All of these works find that the the ground state doublet is local planar in nature, with $g_{\perp} > g_{\parallel}$.

Publication	$g_{\parallel}$	$g_{\perp}$
Hodges et al. [5, 15]	1.79	4.27
Yasui et al. [17]	2.6	3.9
Malkin et al. [6]	1.836	4.282
Cao et al. [23]	2.24	3.98
Gaudet et al. [8]	1.85 ± 0.10	3.62 ± 0.15

sample of $\text{Yb}_2\text{Ti}_2\text{O}_7$. The model they fit to their measurements purports to show that monopoles describe the dynamic susceptibility at low temperatures in zero field, but a large peak in the real part, and a crossover in the imaginary part, of the AC magnetic conductivity, the most obvious features of their model, occur in large frequency gaps in their experimental data, making the analysis questionable.

Two works on thermal conductivity in single crystal samples of $\text{Yb}_2\text{Ti}_2\text{O}_7$ by Shimoyama et al. [52] and Li et al. [50] have also been published. Shimoyama et al. [52] find evidence of strong magnon-phonon coupling, and that the thermal conductivity drops suddenly at the same temperature that a phase transition is observed in the heat capacity. They also find that large applied magnetic fields enhance the thermal conductivity, likely due to the field reducing magnetic disorder which reduces phonon scattering. At very low temperatures, in weak applied magnetic fields, the thermal conductivity is observed to decrease before increasing in the high field regime. The authors speculate that this is due to magnetic monopoles carrying heat at low temperatures, and these monopoles being suppressed in applied field, but there is no direct evidence for this being the source of the observed behaviour. Li et al. [50] see similar results to Shimoyama et al. [52], again observing a sharp drop in the thermal conductivity at T_c . Li et al. [50] find that at low temperatures, the thermal conductivity of $\text{Yb}_2\text{Ti}_2\text{O}_7$ increases with temperature as T^2 , not consistent with phonon based thermal conductivity, which is expected to scale as T^3 . They propose that the

reduction of the scaling of the thermal conductivity as a function of temperature is likely due to magnetic effects, and that these effects are stronger in $\text{Yb}_2\text{Ti}_2\text{O}_7$ than in $\text{Dy}_2\text{Ti}_2\text{O}_7$. In large applied magnetic fields Li et al. [50] find that the thermal conductivity scales as $T^{2.5}$, an effect also observed in Shimoyama et al. [52]. An open question is why even at high fields the thermal conductivity does not increase as T^3 . In magnetic fields applied along the $[001]$ direction, Li et al. [50] find no evidence of a phase transition for fields greater than 0.1 T, and in weak fields along this direction, the thermal conductivity behaves differently to other applied field directions where the thermal conductivity decreases before increasing. In weak fields applied along $[001]$, the thermal conductivity immediately increases, likely due to spin polarisation and the gapping out of low lying excitations [50].

Chapter 2

Heat Capacity

This chapter will provide a brief introduction to the concept of heat capacity, and its relation to other thermodynamic quantities and material properties. The usefulness of heat capacity for mapping out phase diagrams in materials and indirectly probing excitations will be discussed. Additionally, the theory behind the AC measurement technique used for the measurements reported in this work, the design and construction of the apparatus used, and the methods used to analyse the data will be discussed. Finally test measurements on a material with a well characterised low temperature magnetic transition, Cs_2CoCl_4 , will be shown, and a correction factor to account for the apparatus configuration will be determined for application to the other heat capacity measurements in this thesis.

2.1 Introduction

This thesis seeks to understand the magnetic phase diagram of the frustrated pyrochlore $\text{Yb}_2\text{Ti}_2\text{O}_7$, as probed by thermodynamic measurements. One key aspect of the study of that phase diagram is ensuring the high quality of the single crystal sample, manifested in a sharp transition into a low temperature phase below ~ 230 mK. As shown in Chapter 1, the heat capacity of single crystal samples is a good indicator of their quality. In addition, heat capacity measurements will be used to

probe the temperature-magnetic field phase diagram of $\text{Yb}_2\text{Ti}_2\text{O}_7$, and the nature of its excitations in very small fields. The ability of this technique to address all of these different questions makes the heat capacity a very versatile technique, and an important one for this work.

2.2 Theory

The heat capacity C is defined as the ratio between the amount of heat dQ absorbed by a system, and the resulting rise in temperature dT ,

$$C = \frac{dQ}{dT}. \quad (2.1)$$

As will be shown, the heat capacity of a material can depend on many external parameters, and is thus very useful for providing insight into the phase diagrams of materials as a function of those parameters.

2.2.1 The First Law of Thermodynamics

The first law of thermodynamics explains how measurement conditions can affect the heat capacity of systems. The first law states that within a system, energy is conserved, and heat and work are forms of energy. The only way to change the internal energy of the system (U) is to either supply (or remove) heat (Q) to (from) the system, or do (or allow the system to do) work (W). The change in internal energy can be written [65] $\Delta U = \Delta Q + \Delta W$, where ΔQ is the heat supplied to the system, and ΔW is the work done on the system. For a differential change in internal energy dU , this can be rewritten as [65]

$$dU = dW + dQ, \quad (2.2)$$

where $d\bar{}$ represents an inexact differential. Equation (2.2) can then be rearranged to show how the heat capacity of a system is dependent on both the internal energy of the system, and any work taking place during the measurement, resulting in

$$\frac{dQ}{dT} = \frac{dU}{dT} - \frac{dW}{dT}. \quad (2.3)$$

In the case of magnetic solids, pressure (P), volume (V), magnetic field (H), and magnetisation (M) are the experimentally accessible quantities that must be included in the work term, resulting in [66]

$$d\bar{Q} = dU + PdV + \mu_0 M dH. \quad (2.4)$$

The third term is an external energy that comes from the magnetic field applied to the system. In order to internalise this energy a new quantity is defined $U' = U + \mu_0 MH$ [66], accounting for the magnetic energy of the system. Using this new internal energy U' ,

$$d\bar{Q} = dU' + PdV - \mu_0 H dM. \quad (2.5)$$

In this form, the second term can be neglected because in all of the experiments in this thesis, $P = 0$ as the experiments are performed in a vacuum. Under these conditions, two heat capacities that can be defined [66]

$$C_M = \left(\frac{\partial Q}{\partial T} \right)_M, \quad (2.6)$$

$$C_H = \left(\frac{\partial Q}{\partial T} \right)_H, \quad (2.7)$$

the heat capacity at constant magnetisation, C_M , and the heat capacity at constant magnetic field, C_H , respectively. As the magnetisation of the sample is not something that can be easily controlled in experiment, the second quantity, C_H , is the experimental quantity measured in this work.

2.2.2 The Second Law of Thermodynamics and Entropy

The second law of thermodynamics leads to the definition of a quantity S known as entropy. The entropy is related to the reversibility of thermodynamic processes: the entropy of a system and its environment remains constant during a reversible process, and otherwise always increases. Assuming a quasi-static process where the system is always in thermal equilibrium with its environment,

$$dQ = TdS, \quad (2.8)$$

where dS is the change in the entropy of the system. This definition allows the heat capacity to be recast in terms of changes to the entropy of the system in question via the first law. If $C = dQ/dT$, then the heat capacity can be expressed in terms of the entropy as

$$C_x = T \left(\frac{\partial S}{\partial T} \right)_x, \quad (2.9)$$

where x is a property being fixed during the measurement of C , e.g. P, V, M, H .

This results in the heat capacity being a very useful tool for probing the phase diagram of materials, due to the behaviour of the entropy at phase boundaries. Phase transitions exist in many forms, primarily classified as first order or continuous [65]. First order transitions are those which have an associated latent heat, while continuous transitions have no latent heat. For a first order transition, the entropy (the first derivative of the free energy of the system) as a function of the temperature is discontinuous as shown in Fig. 2.1(a). As the heat capacity can be expressed as a derivative of the entropy, this results in a large sharp singularity in the heat capacity as the system crosses the phase boundary (see Fig. 2.1(a)). In the case of continuous transitions, the heat capacity can also show signatures of phase transitions. For example, in the Landau (mean-field) picture of a continuous transition, where the entropy is continuous, but its derivative is not (more correctly the second derivative of the free energy is not continuous), the heat capacity will show a discontinuity at the

phase boundary. In reality, critical fluctuations not accounted for in Landau theory will mask the discontinuity and the heat capacity will have a peak at the temperature of the phase transition, as shown in Fig. 2.1(b). These features make the heat capacity a very useful quantity in the study of phase diagrams. In this work, the phase diagram of $\text{Yb}_2\text{Ti}_2\text{O}_7$ will be probed as a function of temperature and magnetic field using temperature sweeps at constant applied field to study the field evolution of such features.




Figure 2.1: Column (a) shows the entropy S and heat capacity C as a function of temperature across a first order phase transition at $T = T_c$. The discontinuity in S at T_c , ΔS , due to the latent heat of transition L , leads to a singularity in C . Column (b) shows the entropy and heat capacity of a typical continuous phase transition as a function of temperature. The entropy shows a kink at the transition temperature, which leads to a peak in the heat capacity as the derivative of the entropy is not smooth. The singularity in the heat capacity associated with the first order phase transition, and the peak associated with the continuous transition are both easily detected experimentally.

2.2.3 Entropy and Statistical Mechanics

Another very useful aspect of the entropy is that it has a basis in statistics, where it can be defined as [65]

$$S = k_B \ln \Omega, \quad (2.10)$$

where k_B is Boltzmann's constant, and Ω is the number of microstates in the macrostate occupied by the system. If we do not know the exact state of the system in question only the probability, then the entropy can be computed from the partition function Z as

$$S = \frac{\partial}{\partial T} (k_B T \ln Z), \quad (2.11)$$

where

$$Z = \sum_i \exp\left(-\frac{E_i}{k_B T}\right), \quad (2.12)$$

where, i indexes all of the available states of the system with energies E_i , and T is the temperature of the system.

This relationship between the heat capacity and the energies of the various states of the system means that excitations in a system have an impact on the heat capacity. The heat capacity is not nearly as direct a probe of excitations as techniques such as neutron scattering, but some properties of the excitations can be extracted. One specific example that will be useful in the analysis of the heat capacity of $\text{Yb}_2\text{Ti}_2\text{O}_7$ is the Schottky anomaly [67], a peak in the heat capacity that arises due to the presence of gapped excitations in the system. The origin of this peak in the heat capacity can be seen by looking at the relationship between the entropy and the partition function of the system.

If we consider a simple two level system, with degeneracies n_0 for the ground state, and n_1 for the excited state at an energy gap Δ , then the partition function is given by

$$Z = n_0 + n_1 \exp\left(-\frac{\Delta}{k_B T}\right). \quad (2.13)$$

Using this form of the partition function it is possible to derive the heat capacity due to the excited states [67]

$$C_{\text{Schottky}} = N_A k_B \left(\frac{\Delta}{k_B T}\right)^2 \frac{n_0}{n_1} \frac{\exp\left(\frac{\Delta}{k_B T}\right)}{\left(1 + \frac{n_0}{n_1} \exp\left(\frac{\Delta}{k_B T}\right)\right)^2}, \quad (2.14)$$

where N_A is Avogadro's number. This leads to a peak in the heat capacity, as shown in Fig. 2.2. Such features are observed in the heat capacity of $\text{Yb}_2\text{Ti}_2\text{O}_7$ in finite applied magnetic field, and will be discussed in Ch. 3.




Figure 2.2: The feature, known as a Schottky anomaly, that arises in the heat capacity as a function of temperature in a two level system with an energy gap of $\Delta = 1.0$ K. For this plot, the degeneracy of the ground state and excited state are one.

2.3 Methods

As discussed in the previous section, measuring the heat capacity of a material is a matter of measuring the amount of heat that must be supplied to a sample to change its temperature by a measurable amount. This typically involves attaching a heater and thermometer to a thermally isolated sample. The heater supplies a controlled amount of heat to the sample, while the thermometer measures the resulting change in temperature. There are various methods in the literature for performing these measurements including the adiabatic method [68], the relaxation time method [69], and the AC method [70]. In this thesis all measurements will be performed using the AC method.

2.3.1 AC Calorimetry

The AC method of measuring the heat capacity of materials presents several advantages when compared to the other available methods. First, it has less need for thermal isolation of the sample, as a reasonable relaxation time for the sample temperature is actually a requirement of the method, as opposed to the adiabatic method, where any thermal connections to the surroundings must be kept to an absolute minimum. This allowance for finite thermal connections between the sample and its surroundings also allows for stronger sample supports, which can be necessary when measuring strongly anisotropic magnetic samples in large magnetic fields, where significant torques can result, to ensure that the sample remained fixed in position. This consideration was not necessary for the experiments in this work, but it is desirable in principle for future experiments. The periodic nature of the measurement also allows the use of modern signal processing methods to maximise the signal to noise ratio of the measure, and also allows for continuous measurements as a function of changing temperature or other variable, unlike other methods, where the heat capacity is probed via discrete heat pulses.

The theory behind AC heat capacity measurements can be understood by considering the simple model shown in Fig. 2.3(a), where a sample, with heat capacity C_s and temperature T_s , is connected to a thermal bath (with temperature T_b), heater (with heat capacity C_h and temperature T_h), and thermometer (with heat capacity C_θ and temperature T_θ), via heat conductances K_b , K_h , and K_θ . If a current through the heater is used to generate an oscillating amount of heat $\dot{Q} = \dot{Q}_0 (\cos(\frac{1}{2}\omega t))^2$, and the variations in T_h , T_s , and T_θ are small, then the steady state solution for T_θ is given by [70]

$$T_\theta = T_b + \frac{1}{2}\dot{Q}_0 \left(\frac{1}{K_b} + \frac{1-\delta}{\omega C} \cos(\omega t - \alpha) \right), \quad (2.15)$$

where $C = C_s + C_h + C_\theta$, and α is a phase offset. Focusing on the third term of Eqn. (2.15) the relationship

$$\tilde{T}_\theta = \frac{\dot{Q}_0}{2\omega C} (1 - \delta) \cos(\omega t - \alpha) \quad (2.16)$$

emerges, which shows that the oscillating component of the temperature of the thermometer is inversely proportional to the heat capacity of the system. By isolating this term in experiments, the heat capacity can be probed.

The quantity $1 - \delta$ is a frequency dependent normalisation term. If we define the relaxation times

$$\begin{aligned}\tau_s &= C/K_b, \\ \tau_\theta &= C_\theta/K_\theta, \\ \tau_h &= C_h/K_h,\end{aligned}\tag{2.17}$$

and place the conditions on the system that i) $C_h, C_\theta \ll C_s$, ii) that the sample, heater and thermometer come to equilibrium with a time constant $\ll 1/\omega$ (more specifically $\omega^2 (\tau_\theta + \tau_h)^2 \ll 1$), and iii) the frequency is much larger than the inverse of the sample to bath relaxation time such that $\omega\tau_s \gg 1$, then $1 - \delta$ can be simplified to [70]

$$1 - \delta = \frac{1}{(1 + (\omega\tau_s)^{-2} + \omega^2 (\tau_\theta^2 + \tau_h^2))^{-\frac{1}{2}}}\tag{2.18}$$

to first order in $(\omega\tau_s)^{-2}$ and $\omega^2 (\tau_\theta^2 + \tau_h^2)$. The general frequency dependence of this normalisation factor is shown in Fig. 2.3(b), where it can be seen that the function has a plateau as a function of frequency where its value is ~ 1 . Therefore, the frequency of the temperature oscillations used to perform an AC heat capacity experiment should be tuned to lie in this plateau. The phase α in this term can be related to the various properties of the system, but is not of particular interest in the determination of the heat capacity [70].

In real experiments, the sample has finite thermal conductivity, meaning that it takes a finite amount of time for the temperature of the sample to equilibrate throughout its volume, and this must be taken into account to properly understand the temperature oscillations of the system. To understand the effect of finite thermal conductivity, we consider a slab sample of thickness L , cross sectional area A , thermal




Figure 2.3: (a) Schematic diagram of a sample with heat capacity C_s and temperature T_s connected to a thermal bath (with temperature T_b), heater (with heat capacity C_h and temperature T_h), and thermometer (with heat capacity C_θ and temperature T_θ) via thermal conductances K_b , K_h , and K_θ respectively. (b) a plot showing (red) the behaviour of the frequency dependent normalisation factor $1 - \delta$, described by Eqn. (2.18), for $\tau_s = 200$ s and $\tau_h^2 + \tau_\theta^2 = 1.0$ s, and (blue) Eqn. (2.21) with $2K_b/3K_s = 0.5$. The purpose of the blue curve is to illustrate the effect of the finite thermal conductivity of a real sample on the correction factor. (c) The same schematic diagram as in (a), but for a sample of finite thickness L , density ρ , thermal conductivity κ , and heat capacity per unit mass c . This figure is inspired by Sullivan and Seidel [70].

conductivity κ , heat capacity per unit mass c , and density ρ , as illustrated in Fig. 2.3(c). The sample is uniformly heated on one side ($x = 0$), and connected to the thermal bath (with temperature T_b) on the opposite side ($x = L$) via a thermal conductance K_b . The same condition as used in the ideal model, $\omega^2 (\tau_\theta + \tau_h)^2 \ll 1$, is imposed, which means that the thermometer and heater are very well coupled to the sample. It is also useful to define the following quantities [70]: the thermal diffusivity $n = \kappa/\rho c$, the characteristic length $l_0 = (2n/\omega)^{1/2}$, and its inverse $k = 1/l_0$.

Under these conditions it is possible to solve for the oscillating component of the temperature at $x = L$ (the location of the thermometer). If, as in the ideal case, the condition that $\omega\tau_s = 2k^2\kappa LA/K_b \gg 1$ is imposed, as well as $kL \ll 1$ (the sample thickness is small compared to the characteristic length), then [70]

$$\tilde{T}(L, t) = \frac{-i\dot{Q}_0 e^{i\omega t}}{2\omega C} \left(1 + \frac{1}{\omega^2 \tau_s^2} + \omega^2 \tau_{\text{int}}^2 + \frac{2K_b}{3K_s} \right)^{-1/2}, \quad (2.19)$$

where τ_s is the same quantity as in Eqn. (2.17), while $\tau_{\text{int}} = L^2/(90)^{1/2} n$ is associated

with the time it takes for the sample to reach internal thermal equilibrium. Comparing the term in brackets in Eqn. (2.19) to Eqn. (2.18) shows that they are functionally equivalent, with the only difference being a frequency independent term $\frac{2K_b}{3K_s}$. Figure 2.3(b) shows that the effect of this term is to lower the peak value of $1 - \delta$ from ~ 1 in the ideal frequency plateau. The value of this new term can be estimated from the relation $\frac{2K_b}{3K_s} = \frac{2L^2}{3n\tau_s}$. Reference [70] gives typical values at low temperature for the relevant quantities of $n = 10^4 \text{ cm}^2/\text{sec}$, $\tau_s = 1 \text{ s}$, and $L = 1 \text{ cm}$, yielding a value for this term of $< 10^{-4}$, which should have no impact on $1 - \delta$.

Finally, it is useful to cast Eqn. (2.19) in terms of time invariant quantities [70]

$$|\tilde{T}(L, t)| = \frac{\dot{Q}_0}{2\omega C} (1 - \delta), \quad (2.20)$$

where $|\tilde{T}(L, t)|$, the temperature at the thermometer, and $\dot{Q}_0$, the amplitude of the AC heat, are experimentally measurable and controllable respectively, allowing direct access to the heat capacity C by inverting Eqn. (2.20). The only unknown in Eqn. (2.20) is the value of

$$1 - \delta = \left(1 + \frac{1}{(\omega\tau_s)^2} + (\omega\tau_{\text{int}})^2 + \frac{2K_b}{3K_s} \right)^{-\frac{1}{2}}, \quad (2.21)$$

but performing a measurement where the frequency of the oscillating heat is varied while all other quantities remain the same, then a plot of $\omega |\tilde{T}(L, t)|$ will allow for the determination of the optimal value of ω for the experiment.

2.4 Experiment Design

As discussed in Section 2.3, measuring the heat capacity of a sample involves supplying a known amount of heat to a thermally isolated sample and measuring the resulting change in its temperature. The AC method allows for a finite thermal link to a heat bath, but the basic requirements of the experiment remain unchanged: a way to thermally isolate the sample from its surroundings save only for a small heat link, a

heater and thermometer thermally connected to the sample of interest, an electrical connection to the thermometer to record its temperature, and an electrical connection to the heater to supply it with energy.

Measuring the heat capacity without changing any parameters of the system, such as temperature or magnetic field is not particularly interesting. In order to vary the temperature of the sample, the experiment required a heater on the thermal bath in order to vary the background temperature of the experiment, and the apparatus was placed in the bore of a 16 T superconducting magnetic, in order to probe the heat capacity in applied magnetic fields. The desire to apply large magnetic fields to the sample in specific orientations also imposed conditions on the design of the apparatus: that it must be capable of maintaining the orientation of the sample, even when subject to external torques. The design of the apparatus used to meet these requirements followed that of Rost [71]. An overview of its design and construction will be given in the following section, along with the configuration of the electronics used for the experiment.

2.4.1 Sample Environment

Before the heat capacity of the sample could be measured, it first had to be cooled to the temperature range of interest. In this thesis, all measurements were performed at millikelvin temperatures by connecting the heat capacity apparatus to an Oxford Instruments Kelvinox 25 dilution refrigerator. A dilution refrigerator uses the process of transporting ^3He across a phase boundary within a mixture of ^3He and ^4He , between a ^3He rich phase and a ^3He poor phase. This process has an associated latent heat of mixing, which provides the cooling power of the system. This is the only continuous cooling method that can access temperatures of < 300 mK. A more thorough explanation of the mechanism behind the operation of a dilution refrigerator can be found in Pobell [72].

There were two aspects of the operation of the dilution refrigerator that were

important to the construction of the heat capacity apparatus that was used in this thesis [71]. The first was that the sample space of the dilution fridge was under high vacuum, due to the requirement that the mixing chamber must be thermally isolated from its surroundings. This was convenient for the construction of the heat capacity apparatus, as it meant any thermal contacts between the sample and its surroundings would not come from the atmosphere, only physical connections.

Second, the cooling power of a dilution fridge occurs in a volume called the mixing chamber. In this thesis the heat capacity of $\text{Yb}_2\text{Ti}_2\text{O}_7$ was measured both in zero field, and in applied fields of up to 1.5T, made possible by the fact that the dilution fridge used was fitted with a 16T superconducting magnet. The heat capacity apparatus must be centred in the homogeneous field region of the magnet. The mixing chamber of the dilution refrigerator was located outside of the bore of the magnet, due to size constraint issues. This meant that the apparatus needed to be thermally linked to the cooling power of the mixing chamber by a device known as a coldfinger, a length of thermally conductive material that supported the apparatus and allowed it to cool.

Following Rost [71], the coldfinger was constructed from a 4 mm diameter Oxygen Free High Conductivity Copper (OFHC) rod surrounded by a 10 mm thin walled stainless steel (SS) tube to provide structural support. The ends of the OFHC rod and SS tube were terminated in OFHC copper plates, one of which was anchored to the mixing chamber, while the other was anchored to the heat capacity apparatus, ensuring thermal contact. This design provided structural rigidity while minimising the electrically conducting cross sectional area, and thus the effects of eddy current heating in a changing magnetic field. The coldfinger connecting the mixing chamber and the heat capacity apparatus is shown in Fig. 2.4. Not shown is a thin copper shell placed around the heat capacity apparatus to eliminate any possibility of radiative heating of the sample from the 4.2 Kelvin surface of the vacuum can that surrounded the experiment.




Figure 2.4: Labelled photograph of the AC capacity setup, showing in addition, the supporting coldfinger, and the mixing chamber of the dilution fridge.

2.4.2 Apparatus Design

Figure 2.5(a) shows the layout of the apparatus used to perform the heat capacity measurements, along with a photograph of the key components of the apparatus, while Fig. 2.5(b) is a simplified view of the experiment showing the fundamental components and thermodynamic links. The desire to maximise (minimise) the thermal conductivity of the heat links shown in black (red), along with minimising the heat capacity of the sample platform and other components connected to the sample (the addenda), determined the choice of materials for various components of the heat capacity apparatus.

The design of the apparatus began with a circular OFHC thermal bath. Within the ring of the thermal bath, a silver sample platform was suspended via Kevlar threads. On the underside of this platform, a thin film strain gauge was mounted to act as a heater. The top side of the platform was reserved for the sample, on top of which was mounted a small resistive thermometer. This configuration is different from that of Rost [71], with the positions of the heater and thermometer being reversed. The impact of this configuration on the relationship between the theory of Sullivan and Seidel [70], and the experiment was not fully comprehended at the time, but through experiments shown later, it will be shown that the effect of this change is a rescaling of the size of the measured temperature oscillations with respect to the amount of

heat supplied to the system. Finally, the sample platform was linked to the thermal bath by a Pt wire which acted as a tunable thermal link. In this section the reasoning behind each of these components and their relevant properties is discussed.

The thermal bath acted as a heat sink for the experiment, and set the background temperature for the measurements. It was strongly thermally linked to the mixing chamber of the dilution fridge via the coldfinger, and four cylindrical OFHC supports shown in Fig. 2.4. The tunable thermal link between the bath and the sample platform meant that the temperature of the bath was the “DC” temperature of the experiment onto which the AC oscillations were superimposed. The temperature of the thermal bath was controlled via two resistive heaters, shown in Fig. 2.5(a). These heaters were used to vary the background temperature of the experiment from its base temperature (~ 100 mK). The heaters used for this purpose were NiCr foil strain gauges.

The other major physical component of the apparatus was a $5\text{ mm} \times 5\text{ mm} \times 150\mu\text{m}$ 99.99% pure Silver platform. This material was chosen as it had relatively high thermal conductivity at low temperatures [66, 72], could be worked using standard machine tools allowing for easy cutting and shaping, and did not deform easily [71]. It also allowed for the easy smoothing of edges, which was important as sharp edges may have cut the threads suspending the platform.

The sample platform was suspended in space by a series of Kevlar threads passing through the sample platform and thermal bath. This design minimised the thermal connection to the bath due to the support structure, so that the thermal link between the platform and the bath could be controlled via a tunable Platinum wire. The arrangement of the Kevlar threads is shown in Fig. 2.5(a). Two threads were used, with each thread forming a loop that made four connections between the platform and the thermal bath, passing through two phosphor bronze springs, the thermal bath, and holes drilled in the sample platform. The tension on the threads fixed the position of the platform in the centre of the thermal bath, and kept the platform from twisting due to any external forces. Resistance to twisting forces was the reason for the use of




Figure 2.5: (a) Diagram and photographs showing the design and layout of the AC heat capacity apparatus with the key components labelled. (b) A simplified experimental layout showing the key components along with the thermal links (zig-zag lines) between them. The thermal links in black are the links for which it is desirable to have the highest possible thermal conductivity. The red thermal links arise due to the wires necessary to supply heat to the system and measure the temperature of the thermometer, their thermal conductivity should be minimised. The green thermal link is the tunable link that controls the optimal frequency of the thermal oscillations in the experiment.

Kevlar for the supporting threads, instead of more traditional materials such as nylon fishing line or dental floss. In some experimental configurations, the sample can exert a significant torque on the sample platform in the presence of an applied magnetic field (e.g. the Ising ferromagnet in transverse magnetic field, which can generate a sufficient torque to bend entire dilution fridge inserts). The resulting twisting action has the potential to stretch or snap the supporting threads, thus Kevlar was chosen for its enhanced strength. The Kevlar thread used (AR305722 Polyaramid Fibre) consisted of ~ 70 strands of $17\text{ }\mu\text{m}$ diameter. In order to fit through the holes in the sample platform and apparatus, and to further minimise its thermal conductivity, the thread was thinned to ~ 35 strands. The arrangement of the Kevlar threads also allowed for the platform to be fixed in place without the use of any adhesives on the sample platform, which would add to its heat capacity. One undesirable aspect of Kevlar is that it has a negative coefficient of thermal expansion [73], meaning that it expands when cooled. This expansion necessitated the use of leaf springs made from bent sheets of Phosphor Bronze. These springs ensure that the Kevlar threads remained tensioned at all temperatures.

With the physical structure of the apparatus in place, all that remained was to attach the sample, heater, and thermometer to the apparatus. As illustrated in Fig. 2.5(a), the sample was attached to the top of the silver platform, with the thermometer placed on top of the sample. In order to obtain good thermal contact between the thermometer, sample, and platform, while still allowing for the sample and thermometer to be removed without undue strain, Apiezon N grease was used for these connections. This material is a viscous fluid at room temperature, but freezes at $T = 289\text{ K}$, giving it sufficient structural integrity to hold the components of the system together. Rost [71] used Apiezon N grease to connect the thermometer to the apparatus, but used silver epoxy to connect the sample to the platform. As expressed in Schink and Lohneysen [74], considerations of ease of use vs. structural integrity were more important than bonding materials possessing relatively higher thermal conductivity, as boundary resistances between the materials being linked dominate the thermal con-

ductivity at very low temperatures. Rost [71] used epoxy to mount samples, as the experiments carried out there involved large forces on the sample in applied magnetic fields. In the present setup, the heater was connected to the bottom of the sample platform using GE 7031 varnish, a common low temperature adhesive. This material was chosen as it has good thermal conductivity at low temperatures, is commonly available, and is easy to work with. It dries at room temperature, but can be easily removed using the correct solvent, allowing for changes to be made to the apparatus if necessary.

In order to measure the resistance of the thermometer, and apply heat using the heater, it was necessary to make electrical connections to these components, allowing them to be connected to the experimental electronics. Each of these components was contacted using two twisted pairs of wires, allowing for the application of current and simultaneous measurement of the voltage across the component. The ends of the twisted pairs not connected to the heater and thermometer were heat sunk to the thermal bath using contact pads. These wires form a thermal link between these components and the thermal bath. In order to minimise this link, 30 μm diameter Manganin wire was used. Manganin is an alloy of Copper, Manganese, and Nickel noted for its relatively low thermal conductivity for a metal [72]. In the case of the heater, these wires were simply soldered to the contacts of the heater, using a very small amount of Lead-Tin solder. The connection to the thermometer required a more complex approach, in order to avoid any unnecessary heating of the thermometer, which would change the resistance to temperature relation of the thermometer. The thermometer was contacted using 10 μm gold wires which were attached to the thermometer using ultrasonic wire bonding. A small amount of Silver epoxy was then applied to the contacts to provide additional structural support. The gold wires connected to the two contacts on the thermometer were then connected to the Manganin twisted pairs via Lead-Tin solder. A pair of tweezers acted as a heat sink for the gold wires, to protect the thermometer from the heat of the soldering iron. From the thermal bath, the twisted pairs were connected to 30 μm wires heat sunk to the

mixing chamber of the dilution fridge, before being connected to the wiring travelling up and out of the fridge.

Thermometer

The thermometer used to measure the temperature of the sample was a Lakeshore RX-102A-BR RuO₂ thermometer. Such thermometers consist of a film of RuO₂ embedded in lead silicate glass on top of a polycrystalline Al₂O₃ substrate. The resistance of the RuO₂ film changes as a function of temperature, so by measuring its resistance, the temperature of the thermometer can be determined. RuO₂ thermometers are the preferred thermometer for mK temperatures because their resistance changes almost exponentially as a function of temperature at low temperatures, as shown in Fig. 2.6. This exponential increase in the resistance with decreasing temperature is due to the conduction mechanism within the glass, variable-range electron hopping in 3D between the RuO₂ embedded in the glass matrix [72]. The large change in resistance with small changes in temperature at low temperatures make these chips very sensitive thermometers below ~ 1 K.




Figure 2.6: Typical resistance curves of Lakeshore RuO₂ thermometers as a function of temperature. The resistor model used in this work is RX-102A. Figure is taken from Ref. [75].

The overall dimensions of the thermometer chip used were $1.45 \text{ mm} \times 1.27 \text{ mm} \times 0.65 \text{ mm}$. The chip used in the experiments is shown in the top right image in

Fig. 2.5(b). One problem with the chip was that Al_2O_3 has a very low thermal conductivity. In order to ensure a short time constant between the temperature of the sample and the temperature of the thermometer, the thermometer was thinned from 650 μm to 140 μm , leaving ~ 100 μm of Al_2O_3 to provide structural integrity to the chip. Rost [71] pointed out that the polishing process can introduce mechanical stress to the chip, which can have an impact on the reproducibility of the resistance temperature curve of the thermometer. Rost [71] and Goodrich et al. [76] showed that thermal cycling of the chip can remove these stresses, and after sufficient (~ 100) cycles between room temperature and 4.2 K, the resistance of the RuO_2 thermometer at 4.2 K does not change between cycles. Based on these findings, the thermometer chips used for this experiment were cycled between 300 K and 4.2 K 100 times, using a Quantum Designs PPMS as a cryostat. Resistance as a function of cycle plots are not available, as the electrical contacts on the thermometers used during this cycle were not stable, and subsequently had to be replaced with the method discussed previously. After this thermal cycling the thermometers were calibrated as a function of temperature and applied magnetic field. This calibration process is described in Appendix A, and the resulting Resistance to Temperature Conversion functions are given in.

The heater attached to the sample platform was a Kyowa KFL-05-120-C1-11 foil strain gauge. This device was made from a NiCr alloy foil, attached to a Polyimide base. The grid of wires formed by the NiCr foil was 0.5 mm $\times$ 1.4 mm in size, and gave the gauge a resistance of ~ 120 Ω . The resistance of the heater was measured during the heating process, via a four wire measurement, ensuring that any variation in the resistance of the heater was accounted for in the calculation of the heat capacity. The Polyimide base of the strain gauge was originally 3.3 mm $\times$ 2.7 mm, but was trimmed to minimise the amount of excess material contributing to the addenda heat capacity.

Tunable Thermal Link

The final physical aspect of the construction of the heat capacity apparatus is the thermal link between the bath and the sample platform, constructed from a length of 25 μm Pt wire. The thermal conductivity of this link sets τ_s , which must, in combination with the frequency of the temperature oscillations, meet the condition $\omega\tau_s \gg 1$. Equation (2.17) shows the relationship between τ_s , the heat capacity of the sample, and the thermal conductivity of the tunable link. This form of τ_s requires some modification, to account the heat capacity of the addenda in the experiment and the other heat links between the platform and the bath due to the Kevlar and Manganin wires, resulting in

$$\tau_s = \frac{C + C_{\text{addenda}}}{K_b + K_{\text{addenda}}}, \quad (2.22)$$

where C is the heat capacity of the sample, C_{addenda} is the heat capacity of the other components in the system, and K_{addenda} is the thermal conductance of the support structure and measurement leads. Rost [71] performed a detailed analysis of the heat capacity of the addenda, accounting for contributions from the Silver sample platform, the RuO_2 thermometer, the Manganin wires, and the various amorphous solids used to connect the various components. The only fundamental difference between the components used in the apparatus and that of Rost [71] was the use of GE 7031 varnish in the place of other adhesives. The low temperature heat capacity of GE 7031 varnish is reported by Pobell [72], and is very similar to that of Apiezon N grease and the Stycast Epoxy used by Rost [71], consisting of terms linear and cubic in T . The use of this component was not expected to change the analysis of Rost significantly. In order to set the thermal conductivity of the Pt link, 1 K was chosen as a reference temperature at which to set the heat capacities and thermal conductivities of the various components. At this temperature, the heat capacity of the 6.25 mg sample platform, computed based on the heat capacity of Silver measured and fitted by Martin et al. [77] was $5.17 \times 10^{-7} \text{ J K}^{-1}$. Rost [71] estimated the remainder of the

addenda including the Manganin wires, adhesives, and thermometer to contribute an additional $1 \times 10^{-7} \text{ J K}^{-1}$ to the addenda. Algra [78] gives the heat capacity of the test material, Cs_2CoCl_4 , as $\sim 2.83 \text{ J K}^{-1} \text{ mol}$. The sample used in the experiment has a mass of 6.2 mg, resulting in a heat capacity of approximately $3.76 \times 10^{-5} \text{ J K}^{-1}$, two orders of magnitude larger than the addenda heat capacity.

Now the thermal link, K_b , necessary to achieve $\omega\tau_s \gg 1$ can be computed. Rost [71] estimated the contributions to this thermal link from the Kevlar support threads and Manganin twisted pairs as 3.1×10^{-8} and $3.6 \times 10^{-8} \text{ W K}^{-1}$ respectively at 1 K. The thermal link used in the apparatus was a 1 cm length of 25 μm Pt wire. Pobell [72] gives the thermal conductivity of Pt as $0.33 \text{ W cm}^{-1} \text{ K}$ at 1 K. For a wire of diameter 25 μm and length 1 cm, the thermal conductance will be $1.62 \times 10^{-6} \text{ W K}^{-1}$. Combining this thermal conductance with those of the Manganin wires and Kevlar threads, and using the previously given sample and addenda heat capacity yields $\tau_s = 23 \text{ s}$. Assuming a typical operating frequency for the experiment of $1 \times 10^{-2} \text{ Hz}$, then $\omega\tau_s = 1.45$ marginally satisfies the condition $\omega\tau_s \gg 1$. This condition can be better satisfied using a higher operating frequency, but it will be shown later that the actual value of τ_s realised in the experiment is much larger than $\sim 23 \text{ s}$, allowing for an operating frequency of $1 \times 10^{-2} \text{ Hz}$. The increase in τ_s is likely due to boundary thermal resistances at the ends of the heat link, which are not accounted for in these calculations, decreasing its thermal conductance.

With the thermal link in place, the design of the heat capacity apparatus was complete. The next section will discuss the electronics involved in the measurement, necessary to apply current to the heater, measure the resulting voltage, and measure the resistance of the RuO_2 thermometer.

2.4.3 Experimental Electronics and Control

Figure 2.7 shows the basic layout of the electronics and equipment used to operate the heat capacity apparatus in this work. The operating controls for the dilution




Figure 2.7: Schematic diagram of the instruments and connections necessary for low temperature AC heat capacity measurements showing the dilution fridge, superconducting magnet, heat capacity apparatus, measurement electronics, and control computer.

fridge (represented by the blue rectangle) and superconducting magnet are omitted from this diagram, as they were controlled independently of the experiment. The dilution fridge was cooled to its base temperature, and all variation in temperature was performed by using heaters attached to the thermal bath of the heat capacity experiment. The upper temperature limit was determined by the temperature at which the dilution fridge lost stability, ~ 700 mK. Figure 2.7 also omits the path of the measurement wires down the dilution fridge insert. In the actual experiment, the fine wires from the experiment (shown as fine black lines forming twisted pairs in Fig. 2.7) ran up the dilution fridge insert, being heat sunk at various points. At the top of the dilution fridge, they were passed through a vacuum tight fitting, and each wire was passed through a pi filter (Tusonix 4209-003LF) which acted as low pass filters to prevent high frequency noise from reaching the bottom of the dilution fridge, as this noise would cause undesirable heating. The box containing the pi filters was grounded to the shield of a lab cable made up of shielded twisted pairs. The twisted pairs in the cable are represented by single heavy black lines in Fig. 2.7. Just before the wires reached the experimental electronics, represented by grey boxes, they were broken out from the single heavy cable and run via coaxial cables for ease of use. From the heat capacity apparatus to the measurement electronics, the wires were run as twisted pairs whenever possible, in order to help minimise any noise in the signal due to magnetic flux passing through the wire loops formed by the leads. As for the superconducting magnet, the magnetic field was only varied between heat capacity measurements, not during measurements, so those electronics are not fundamental to the experiment. For reference purposes, the magnet used was an Oxford Instruments 16 T superconducting magnet, but the maximum field used during these experiments was 1.5 T.

The primary component of the AC heat capacity experiment was the thermometer placed on the sample, used to measure the temperature oscillations in the sample. In this work, the resistance of the thermometer was measured using a four wire measurement performed by a combination of a custom made optically isolated current source

and a Stanford Research SR830 Lock-in Amplifier. This combination of measurement electronics was chosen over the use of an AC resistance bridge as it allowed for more frequent measurements of the resistance, which was important for the sampling of the temperature oscillation waveform.

A Lock-in Amplifier works by multiplying an input voltage signal by a reference signal, with a frequency equal to that of the oscillating signal of interest. The multiplied signal is then time averaged, and due to the orthogonality of sinusoidal functions, all components of the signal which do not share the same frequency as the reference signal will average to zero, leaving only the desired signal. This allows for very accurate measurements of the amplitude of the oscillation. As the AC heat capacity measurement relied on accurate measurements of the amplitude of temperature oscillations, it would seem obvious to set the frequency of the reference signal to that of the temperature oscillation, and use the Lock-in Amplifier to directly and very accurately measure the amplitude of the voltage oscillations across the thermometer that result from the temperature oscillations of the sample. Unfortunately this was made impossible by the very low frequency of the temperature oscillations used, on the order of 10^{-2} Hz, too low for the Lock-in amplifier to operate. Instead the thermometer resistance was measured at a much higher frequency, 13.4 Hz. This frequency was used in order to avoid contamination of the measurement by line noise. The use of a much higher frequency than that of the temperature oscillations allowed for the oscillations to be recorded as a function of time for later analysis.

The purpose of the optically isolated current source was to prevent the existence of ground loops in the measurement of the resistance of the thermometer. This was important as an accurate measurement of the resistance of the thermometer was key to an accurate measurement of the temperature. The optically isolated current source generated a 13.4 Hz triangle wave AC current with a defined amplitude between 3 nA and 100 μ A. The power for the current source came from a 9 V battery, and it sent a timing pulse to the Lock-in Amplifier, which set the frequency of the reference signal, via optical fibre. The combination of battery power supply and optical fibre

timing pulse eliminated any possibility of an electrical loop being formed between the current source electronics and any of the other equipment in the lab via the mains ground. The current from the current source passed down a pair of leads to the RuO₂ chip thermometer. A second pair of leads connected to the thermometer were used to probe the voltage drop across the sample generated by the current. These leads were also connected to the optically isolated current source, where the voltage signal was amplified by a factor of 1000. After this amplification, the AC signal was passed to the Lock-in amplifier. The very low currents generated by the optically isolated current source were necessary to avoid self heating of the thermometer. Passing a current through the sample generates power $P = I^2 R$, which, due to the very high resistance of the RuO₂ thermometer at base temperature ($\sim 25 \text{ k}\Omega$), can have a significant effect on the temperature of the thermometer, even if the sample and other components are at a lower temperature. The current amplitude was tuned to minimise noise in the measured resistance while at the same time minimising self heating of the thermometer.

Oscillations in the temperature measured by the RuO₂ thermometer were generated by passing a very low frequency sinusoidal current through the heater connected to the sample platform. This current was generated by a Keithley 6221 programmable current source. Typical currents used in experiments were of the order of $10 \text{ }\mu\text{A}$. The current supplied had frequency f , while the temperature oscillations will have frequency $2f$, as the power generated by a heater is given by $P = I^2 R$, and the product of two sine waves of frequency f , is a sinusoid with frequency $2f$, along with a constant offset that can be neglected. As previously mentioned, the frequency of the temperature oscillations used was of the order of 10^{-2} Hz .

In order to be able to determine exactly how much heat was added to the sample-heater-thermometer-platform system, a voltmeter was connected across the heater on the sample platform, such that a four wire measurement of the resistance of the heater was performed at the same time the heat was applied to the system. This allowed the heat added to the system to be easily calculated via the relation $P = IV$, where

I is the amplitude of the current sine wave and V is the amplitude of the voltage sine wave. This meant that any variation in the resistance of the heater as a function of temperature or magnetic field did not lead to inaccuracies in the calculated amount of heat entering the system.

In order to vary the temperature of the apparatus, current was passed through the heaters connected to the thermal bath. This was done using a simple programmable DC current source (Keithley 224). This current source was used to generate a smoothly ramping current as a function of time, which lead to a steadily increasing power generated by the thermal bath heater, leading to a smooth increase in the temperature of the experiment.

All of the electronics used in the experiment were connected to a control computer running the Fridge Control package, created by various members of the Coldea Group for the primary purpose of running heat capacity experiments. This Python based software allowed for the simultaneous control and readout of multiple different experimental components. It allowed for the generation of oscillating heat using the AC current source, whilst simultaneously recording the resistance of the RuO_2 thermometer via the optically isolated current source and Lock-in amplifier, and changing the background temperature of the experiment by slowly varying the current supplied by the DC current source. All of the inputs to and outputs from the electronics were then stored to a time indexed text file for later processing.

2.5 Analysis Procedure and Test Measurement on Cs_2CoCl_4

This section will discuss the data collected during a heat capacity experiment, and how it is analysed to determine the heat capacity of a sample via zero field measurements on a single crystal sample of Cs_2CoCl_4 , a quasi-1D XY antiferromagnet with a phase transition near 220 mK. First, the various variables controlled and measured during

the experiment are discussed, and how they are converted into the physical quantities necessary to compute the heat capacity of the sample from the experimental data. Then the calculated heat capacity are compared to literature values, and a scaling factor necessary to account for the experiment construction will be determined.

2.5.1 Analysis Procedure

The variables recorded during a heat capacity experiment were: voltage across the RuO₂ thermometer, voltage across the sample heater, the amplitude of the AC current through the heater, and the frequency of the AC current through the heater. These values were recorded as a function of time by the Fridge Control package for later analysis.

In order to convert these into physically relevant quantities, several steps were taken. Figure 2.8 shows the data collected during a heat capacity experiment, and the various steps involved in converting that data into a value for the sample heat capacity. The first step was to convert the voltage signal measured by the Lock-in Amplifier into resistance, by dividing by the current supplied by the optically isolated current source. The resulting signal as a function of time is shown in Fig. 2.8(a). The inset of Fig. 2.8(a) shows a small section of the signal, revealing oscillations in the resistance due to oscillations in the temperature of the sample. These oscillations are due to the AC current applied to the heater attached to the sample platform. Figure 2.8(b) shows a typical voltage oscillation in the sample heater due to the applied current. The red line in Fig. 2.8(b) is a least squares fit to the voltage, performed in order to determine the amplitude of the voltage oscillations. Using this fitted amplitude, along with the amplitude of the applied current, the power applied to the system was determined without relying on the assumption that the heater has a constant resistance as a function of temperature and applied magnetic field. This power was a key component in computing the heat capacity via Eqn. (2.20).

Once the thermometer voltage was converted into resistance, each resistance value




Figure 2.8: Figure showing the raw data collected during a heat capacity experiment and how it is converted into a value for the heat capacity. (a) shows the resistance of the RuO_2 thermometer resistance as a function of time, obtained by dividing the voltage measured with the Lock-in Amplifier by the applied current. The inset shows oscillations in the resistance superimposed on the steadily decreasing background. (b) shows a typical oscillation in the voltage across the heater, the heat supplied to the system will behave as the square of this signal. (c) shows the resistance in (a) converted into temperature. (d) shows a small portion of the data in (c), corresponding to the same time window as in (b) showing the temperature oscillations, illustrating that their frequency is double that of the current/voltage oscillations in the heater. (e) shows the amplitude of the temperature oscillations plotted against the background temperature, showing a sharp drop in the amplitude of the oscillations at $T_c = 227(2)$ mK, which corresponds to a peak in the heat capacity, shown in (f). The red lines in (b) and (d) represent least squares fits to the data.

can be equated to a temperature via the resistance to temperature conversion functions in Appendix A. This created a temperature signal as a function of time, shown in Fig. 2.8(c). This signal shows a steadily increasing average temperature as a function of time. The increasing temperature was achieved by the application of a smoothly increasing current to the heaters connected to the thermal bath of the heat capacity apparatus. Figure 2.8(d) shows a small portion of the temperature signal from Fig. 2.8(c), showing the small oscillations in temperature superimposed on the increasing background temperature. These oscillations lead to the finite width of the data plot in Fig. 2.8(c). The highlighted “narrow” point in the data, indicates the point in time where the amplitude of the oscillations is minimised. As the heat capacity is inversely related to the heat capacity, this minimum corresponds to a peak in the heat capacity, indicating the presence of a phase transition. A comparison of Figs. 2.8 (b) and (d), showing the voltage and temperature oscillations in the same time window respectively confirms that the frequency of the temperature oscillations is twice that of the applied current oscillations, as expected given the relationship between the applied current, voltage, and the resulting power from the heater.

The red line in Fig. 2.8(d) is a least squares fit to the temperature signal. The fitting function is given by

$$f = A \sin(\omega t + \phi) + Bt + C, \quad (2.23)$$

where A , ϕ , B , and C are free parameters, ω is the frequency of the oscillation, and t is time. A is the amplitude of the temperature oscillations, while $Bt + C$ at the mid-point of the fitted signal determines the background temperature. These two values are plotted, with the amplitude of the oscillations plotted as a function of the background temperature, in Fig. 2.8(e). The minimum in the amplitude of the temperature oscillations is clearly seen, at $T_c = 227(2)$ mK. Using the amplitude of the temperature oscillations, the power applied to the sample by the heater, and the frequency of the temperature oscillations, the heat capacity of the sample was

determined, as shown in Fig. 2.8(f), where the calculated heat capacity has been normalised by the number of moles of Cs_2CoCl_4 .

Optimal Frequency Determination

In order to determine the optimal frequency for the heat capacity measurement, to maximise $1 - \delta$ (from Eqns. (2.20) and (2.21)), a frequency sweep at base temperature was performed before each experiment. To perform such a sweep, a constant amplitude oscillating current was applied to the sample heater, and the frequency of the oscillations was varied. By measuring the amplitude of the temperature oscillations at fixed temperature as a function of frequency, it was possible to compute a value proportional to $1 - \delta$. By inverting Eqn. (2.20) one obtains

$$\omega \cdot |\tilde{T}| = \frac{\dot{Q}_0}{2C} (1 - \delta) \propto 1 - \delta, \quad (2.24)$$

where $\dot{Q}_0/2C$ forms a proportionality constant if the experiment is performed at fixed temperature. Figure 2.9(a) shows the results of such a frequency sweep, indicating that 10^{-2} Hz is an acceptable frequency for the heat capacity experiment.

The frequency dependence of the data in Fig. 2.9(a) provides some insight into the effects of the changes to the heater-thermometer geometry in the experimental apparatus. There does not appear to be any new effect on the frequency dependence of the correction factor, with the frequency dependence matching that of Eqn. (2.21). Based on the frequency dependence, the only effect of the altered apparatus geometry is likely to be a scaling factor applied to the temperature oscillation amplitude. This scaling factor will arise because the heat supplied by the heater can escape through the thermal link before it reaches the sample and the thermometer. By comparing the raw heat capacity data to values in the literature, it should be possible to determine this scaling factor. Another possibility is a rescaling of the internal relaxation time τ_{int} , as this is defined by the thermal links between the components of the system, but as long as the condition $\omega\tau_{\text{int}} \ll 1$ is maintained, the sample, platform, heater,

and thermometer will be in thermal equilibrium.




Figure 2.9: (a) shows the results of frequency sweep measurements using the AC heat capacity apparatus at approximately constant temperature. The plateau as a function of increasing frequency of the temperature oscillations indicates the optimal frequency for the heat capacity measurements. The frequency chosen for the experiments was 10^{-2} Hz. (b) shows the raw results of heat capacity measurements on Cs_2CoCl_4 compared to the literature values from Refs. [78, 79].

2.5.2 Test Measurement: Cs_2CoCl_4

Cs_2CoCl_4 is a spin-1/2 Quasi-1D XY -like antiferromagnet [80]. Its main property of interest for the purposes of this work is the sharp peak in its heat capacity at $T_c = 222 \pm 5$ mK [78], due to the onset of three dimensional long range magnetic order. This is well within the accessible temperature range of the heat capacity apparatus, making it an ideal test material for the apparatus, by ensuring that it would reproduce known results before using it to test new materials. Additionally, high quality samples of this material were readily available.

The sample used for these measurements was grown from solution by Dr. Zbigniew Tylczynski of Adam Mickiewicz University, Pozna, Poland. The sample used had a mass of $6.2(1)$ mg, with dimensions $2.0 \text{ mm} \times 1.75 \text{ mm} \times 533 \mu\text{m}$, and was cut so that the orthorhombic a axis of the crystal was perpendicular to the large face of the sample.

Figure 2.9(b) shows raw heat capacity data for Cs_2CoCl_4 , composed of the data

shown in Fig. 2.8(f), combined with data from additional temperature sweeps. For this plot, the heat capacity of the silver sample platform, taken from de Nobel and du Chatenier [81] was subtracted from the data. Even with the main contribution from the addenda removed from the data, it can clearly be seen that the heat capacity of the sample of Cs_2CoCl_4 was significantly overestimated. As previously discussed, this was due to the geometry of the thermometer-sample-platform-heater arrangement used to perform the experiment. As the thermal link between the sample platform and the thermal bath was between the heater and the sample-thermometer component of the system, some portion of the heat supplied by the heater escaped via the thermal link, rather than affecting the sample temperature. This lead to reduced temperature oscillation amplitudes compared to those expected based on the actual heat capacity of the sample and the amount of heat supplied by the heater.

Fortunately, the two sources in the literature reporting the low temperature heat capacity of Cs_2CoCl_4 , show a temperature range (290 – 500 mK) where the heat capacity of Cs_2CoCl_4 is consistent between the two data sets. By scaling the measured heat capacity of Cs_2CoCl_4 to match the literature values in this range, it was possible to determine the ratio of the heat that enters the sample and affects its temperature to the heat supplied by the heater, which was found to be 0.674(8). In Ch. 3, it will be shown that this scaling factor also results in good agreement between measurements of the high temperature heat capacity of $\text{Yb}_2\text{Ti}_2\text{O}_7$ and results reported in the literature.

Figure 2.10(a) show the results of this scaling on the heat capacity data, achieving good agreement between the experimental and literature values of the heat capacity of Cs_2CoCl_4 in the temperature range 290 – 500 mK. In addition to this improved agreement in the magnitude of the heat capacity of Cs_2CoCl_4 , two other significant features in the heat capacity of Cs_2CoCl_4 can be observed in Fig. 2.10(a). The AC heat capacity measurements show a very sharp peak in the heat capacity of Cs_2CoCl_4 at $T_c = 227(2)$ mK. This feature is consistent with transition temperatures reports in by Algra et al. [78] and Breunig et al. [79] of $T_c = 222(5)$ mK and $T_c = 220(5)$ mK respectively. The peak observed in the AC heat capacity measurements was much

sharper than that observed in either of the two peaks from the literature. Algra et al. [78] assign this phase transition to the family of transitions known as lambda anomalies, continuous infinite order transitions, while Breunig et al. [79] suggest that their heat capacity experiments contain evidence that the phase transition may be first order. For either type of phase transition, a high resolution heat capacity measurement can be expected to show an extremely sharp anomaly, consistent with the anomaly observed here, suggesting that these measurements were more sensitive than the previous measurements of the transition. The literature values of the heat capacity of Cs_2CoCl_4 below the phase transition are not consistent between the two measurements, so no comment can be made on the inconsistency of the new measurements with the previous results. Also observed in the AC heat capacity measurements is the broad maximum in the heat capacity of Cs_2CoCl_4 reported by Algra et al. [78], indicative of the presence of a Schottky anomaly. This anomaly is associated with exchange interactions within the Quasi-1D chains [78].




Figure 2.10: (a) shows the results of frequency sweep measurements using the AC heat capacity with variations in sample temperature fully accounted for, and corrected for loss of heat to the thermal link before reaching the sample. (b) shows the corrected results of heat capacity measurements on Cs_2CoCl_4 compared to the literature values from Refs. [78, 79], showing good agreement.

Figure 2.10(b) shows the effect of the scaling factor on the correction factor in Eqn. (2.20), $1 - \delta$. The quantity $2\omega |\tilde{T}| C / \dot{Q}_0$ is equivalent to $1 - \delta$ via Eqn. (2.20). By computing this quantity using the recorded temperature during the frequency sweep shown in Fig. 2.8(a), and the computed heater capacity, a plateau as a function of

frequency should be recovered, which should peak at an amplitude of 0.674(8) due to the scaling factor. Examining Fig. 2.10 the peak of the correction factor is found to be 0.669(3), indicating that excepting the reversed geometry of the heat capacity apparatus, all of the requisite conditions on the various thermal conductivities within the apparatus have been met. The values for τ_{int} and τ_s used for the fit to the data are 5(1) s and $\sim 220(20)$ s respectively. The high uncertainty in τ_s is due to the fact that the high frequency fall off in $1 - \delta$ is not well characterised. It is notable that the value of τ_s is much greater than that expected based the on calculations in Section 2.4.2. As the value of τ_s is increased, not decreased, this is not a problem, as it serves to ensure that the condition $\omega\tau_s \gg 1$ is more rigorously enforced.

2.6 Conclusions

This chapter has given an overview of the thermodynamics behind heat capacity and its measurement. It has also given a brief overview of the theory behind the experimental design, and a brief description of both the apparatus design, and the electronics used to perform the measurement. Finally, test measurements on Cs_2CoCl_4 have been described, and their success in reproducing previous experimental results discussed. With the ability of the apparatus to measure heat capacities established, it is now possible to move on to the measurements of interest in this thesis, the heat capacity of $\text{Yb}_2\text{Ti}_2\text{O}_7$.

Chapter 3

Low Temperature Heat Capacity of $\text{Yb}_2\text{Ti}_2\text{O}_7$

This chapter presents the results of heat capacity measurements on a single crystal sample of $\text{Yb}_2\text{Ti}_2\text{O}_7$. Measurements in zero applied magnetic field show a single sharp peak in the heat capacity at $T_c = 214(2)$ mK, indicating that this is a high quality sample. Heat capacity measurements performed in applied magnetic fields of up to 1.5 T along the $[001]$ direction show that very small internal magnetic fields ($\mu_0 H_{\text{int}} \sim 0.018$ T) completely suppress the peak in the heat capacity. This is consistent with the low temperature phase of $\text{Yb}_2\text{Ti}_2\text{O}_7$ having a net moment along one of the $\langle 001 \rangle$ directions, so that when the magnetic field polarises the paramagnetic phase along one of those directions, the transition is replaced by a crossover. The in field heat capacity is also used as an indirect probe of the excitations of the system. By fitting a broad hump observed in the data using a simple two level model, a nominal energy gap to the lowest energy excitation is extracted for varying applied magnetic fields.

3.1 Introduction

This chapter studies the heat capacity of $\text{Yb}_2\text{Ti}_2\text{O}_7$ in order to place the sample used in context with respect to the large body of literature on this material. As

discussed in Chapter 1, the low temperature magnetic behaviour of single crystals of $\text{Yb}_2\text{Ti}_2\text{O}_7$ varies between samples, with some showing evidence of magnetic long range order [17, 33, 37, 39, 41], while others show evidence of persistent magnetic dynamics [5, 15, 16, 18, 20–22, 25, 36, 40]. Similarly, the low temperature heat capacity of single crystal samples varies widely, as shown in Fig. 3.1, with different crystals showing very different peaks in their heat capacity, both in terms of the temperature at which the peak occurs, and the shape of the peak. In contrast, recent measurements on powder samples of $\text{Yb}_2\text{Ti}_2\text{O}_7$ consistently show an extremely sharp peak in the heat capacity near $T \sim 265$ mK (an early measurement of the heat capacity of a powder found a single sharp peak at 214 mK). The differences in the low temperature magnetic behaviour and heat capacity of single crystal samples appear to be correlated, as shown in Table 3.1. Samples with small broad features tend to show evidence of persistent magnetic dynamics, while samples with larger sharper anomalies (more like those found in powder samples of $\text{Yb}_2\text{Ti}_2\text{O}_7$), show evidence of long range magnetic order. Outliers to this trend do exist, for example the heat capacity of the sample of Ross et al. [25], which has been shown to display evidence of magnetic correlations but no long range order at low temperatures, shows two features: one very similar to the peak seen in the heat capacity of the powder sample used to synthesise the crystal and the other similar to peaks in the heat capacities of samples that show evidence of persistent magnetic dynamics at low temperatures. The sample of Bhattacharjee et al. [41] shows indirect evidence of ferromagnetic order at low temperature, but while large, the peak in the heat capacity of the sample used is quite broad. It should also be noted that the sample of Bhattacharjee et al. [41] was produced in a manner much closer to that used to produce the samples of Ross et al. [25] and Yaouanc et al. [26], which show evidence of persistent magnetic dynamics.

The sample used in this work was grown using similar methods to those of Yasui et al. [17], described in Appendix B. This method was used because the sample used in the works of Yasui et al. [17], Chang et al. [33, 39], and L’Hotel et al. [37] shows a single sharp peak in its heat capacity, most similar to those of stoichiometric powder




Figure 3.1: The heat capacity of single crystal samples taken from the literature. The samples of Ross et al. [25], Yaouanc et al. [26], D’Ortenzio et al. [36], and Robert et al. [40] show weak features in the heat capacity at low temperatures (the sample of Ross et al. [25] includes a second, sharp peak). The samples of Chang et al. [33] and Bhattacharjee et al. [41] show larger features in their heat capacity, though the feature of Bhattacharjee et al. [41] is quite broad. The heat capacity of Robert et al. [40] is not included in as they do not make any measurements that probe whether or not the sample shows magnetic long range order.

samples when compared to the heat capacity of any other single crystal sample in the literature. The sample of Yasui et al. [17] and Chang et al. [33] also consistently shows evidence of ferromagnetic long range order below that peak. That crystal has also been shown via EXAFS measurements to have fewer defects and deficiencies than a sample studied in the same work [33] produced using similar methods to the samples of Ross et al. [25] and Yaouanc et al. [26], and which was found to have a smaller broader peak in its heat capacity than that of the primary sample of Chang et al. [33], similar to peaks seen in those works. In order to confirm that this new crystal exhibits a single sharp peak in its heat capacity, the heat capacity of the sample has been measured in the temperature range 100 – 700 mK.

Table 3.1: Heat capacity features of single crystal samples of $\text{Yb}_2\text{Ti}_2\text{O}_7$, catalogued from the literature, along with the low temperature magnetic behaviour observed in those crystals.

Reference	Feature Type	T (mK)	Magnetism	Evidence
Chang et al. [33]	Sharp Peak	196	LRO	Neutron scattering [33], magnetisation [37], and μSR [39]
Bhattacharjee et al. [41]	Large Broad Peak	170	LRO	Ultrasound velocity and attenuation measurements
Ross et al. [25]	Sharp Peak Broad Feature	265 190	Persistent Dynamics	Neutron scattering
Yaouanc et al. [26]	Small Broad Peak	160	Persistent Dynamics	Neutron scattering, μSR , and Mössbauer measurements [20]
D’Ortenzio et al. [36]	Broad Feature	185	Persistent Dynamics	μSR measurements

3.2 Sample

The sample used for the heat capacity measurements was cut from the end of the single crystal described in Appendix B before performing the inelastic neutron scattering measurements described in Ch. 5. The resulting sample is a thin square of dimensions $2.09(1) \text{ mm} \times 2.03(1) \text{ mm}$ in cross section, a thickness of 1.01 mm , and a mass of $9.7(1) \text{ mg}$. This geometry was chosen to ensure good thermal coupling of the sample to the sample platform and thermometer used for the heat capacity measurements. Before the sample was cut, it was aligned such that one of the cubic axes of the crystal was perpendicular to the large flat face of the crystal to within 1° . This was done using the back-reflection Laue method, discussed in Appendix C.

3.3 Zero Field Measurements

Figure 3.2(a) shows the results of heat capacity measurements on the single crystal sample described in Section 3.2 in zero field, compared to the zero field heat capacity of

the single crystal sample of Chang et al. [33]. The $C(T)$ measurements were performed after the sample had been cooled to base temperature in zero field. Following this, a magnetic field of 0.5 T was then applied to the sample, after which the magnetic field was then reduced to the desired field value for the measurement. 0.5 T is much larger than the field required to suppress the transition observed in the magnetisation of single crystal samples [37], so this field should be sufficiently large to uniformly polarise any domains in the low temperature phase. Figure 3.2(a) shows that the peak in the heat capacity of the sample used in this work is larger (by a factor of approximately $\times 1.5$) and sharper (no low temperature shoulder is visible in the heat capacity of the new sample) than the peak in the heat capacity observed by Chang et al. [33]. The peak occurs at $T_c = 214(2)$ mK, compared to $T_c = 195$ mK for the sample of Chang et al. [33]. These features make this crystal the most similar to stoichiometric powders in terms of low temperature heat capacity of any single crystal sample produced to date. In fact, the transition temperature of this crystal is the same as that of the powder sample of Blöte et al. [9], but not of more recently grown powder samples [21, 25].

Figure 3.2(b) shows the quantity $2\omega \mid \tilde{T} \mid C/\dot{Q}$ (parameters as defined in Ch. 2.3.1) computed using the heat capacity shown in Fig. 3.2(a) and the frequency sweep data used to determine the optimal frequency for performing the measurements. As discussed in Chapter 2, this quantity should be equivalent to the frequency dependent correction factor $1 - \delta$, which should have a maximum of ~ 1 , but due to the correction factor discussed in Ch. 2, is expected to have a maximum amplitude of 0.674(8). In the case of these measurements, the peak correction factor is 0.666(3), within the uncertainty of the theoretical value, at a frequency of 3.9×10^{-3} Hz. The times constants used to generate the theoretical curve for $2\omega \mid \tilde{T} \mid C/\dot{Q}$ in Fig. 3.2(b) are $\tau_{\text{int}} = 5.4(5)$ s and $\tau_s = 180(20)$ s, which are similar to those found for the frequency sweep performed during the test measurements on Cs_2CoCl_4 . τ_s is slightly smaller, but not significantly. The frequency of the temperature oscillations used for the experiment was 5×10^{-3} Hz based on this frequency sweep. The use of the same

correction factor for the amplitude of the temperature oscillations, to account for the experiment geometry, is confirmed by the excellent agreement between the high temperature ($T > 0.45$ K) heat capacity of this sample, the sample of Chang et al. [33], and the powder sample of Ross et al. [25] shown in Fig. 3.2(a).

These measurements confirm that by following the growth method of Chang et al. [33], a single crystal sample has been produced that is of as good, if not better, quality than the crystal of Chang et al. [33], based on the fact that it has a larger sharper peak in its heat capacity at $T = 214(2)$ mK than that sample, and the trends in heat capacity vs. sample quality reported in that work. This single sharp peak also makes this the single crystal most similar to powder samples in terms of low temperature heat capacity, though the observed peak is still not as sharp as that observed in stoichiometric powders, as shown in the inset of Fig. 3.2(a).




Figure 3.2: (a) Heat capacity measurements on a single crystal of $\text{Yb}_2\text{Ti}_2\text{O}_7$ compared to previously published measurements on a single crystal samples by Chang et al. [33]. Inset compares the heat capacity of this sample to the heat capacity of a powder taken from Ross et al. [25]. (b) Frequency sweep showing the amplitude of $2\omega |\tilde{T}| C/\dot{Q}$ as a function of the frequency of the temperature oscillations induced in the sample via AC heating. The green line is an approximation to the expected behaviour based on Eqn. (2.21) and an additional scaling factor derived from test measurements.

3.4 Measurements in Applied Field

In addition to characterising the sample to determine its quality with respect to other crystals, heat capacity is useful for studying the temperature-magnetic field phase diagram of $\text{Yb}_2\text{Ti}_2\text{O}_7$. This was done by performing heat capacity measurements in various applied magnetic fields of up to 1.5 T along the [001] axis of the sample. All of the measurements were performed in the same manner as the zero field measurement, with a 0.5 T magnetic field being applied to the sample at base temperature, before the field was reduced to the field at which the measurement was performed. The results of these measurements are shown in Fig. 3.3(a). The frequency of the temperature oscillations used for the in field measurements varied slightly from those of the zero field measurements, but the resulting change in the frequency dependent correction factor is small, and has been accounted for using scaling factors determined from the data in Fig. 3.2(b).

Figure 3.3(a) clearly shows that increasing the applied field along [001] results in a rapid suppression of the sharp peak in the heat capacity at $T = 214(2)$ mK (indicated by the dashed red line in Fig. 3.3, with the peak decreasing in amplitude upon increasing the applied field, before disappearing completely in an applied field of 0.125 T. This is equivalent to an average magnetic field within the sample, after correcting for demagnetisation effects (see Appendix D), of $\mu_o H_{\text{int}} = 0.036$ T. The temperature of the peak does not change within error bars in increasing field.

In other works on $\text{Yb}_2\text{Ti}_2\text{O}_7$, this peak has been associated with a first order phase transition to a ferromagnetically ordered phase at low temperatures, based on thermal hysteresis of magnetisation measurements [37] and polarised neutron scattering measurements [33]. The results reported here cannot indicate whether the observed transition is first order or continuous, as the thermal hysteresis of the transition has not been studied, and the fact that the AC heat capacity technique uses oscillations in the temperature of the system would likely lead to any such hysteresis being smeared away. The temperature oscillations will also lead to some suppression of the sharp-




Figure 3.3: (a) Heat capacity measurements on $\text{Yb}_2\text{Ti}_2\text{O}_7$ in applied magnetic fields of 0 – 1.5 T (0 – 1.37 T internal field) along the [001] direction. The results show that very small magnetic fields along this direction completely suppress the observed phase transition, and introduce a Schottky-like anomaly that shifts to higher temperature upon increasing the applied magnetic field. The zero field transition temperature (214(2) mK) is indicated by the dashed red line. (b) A temperature-magnetic field phase diagram created based on the peak in the heat capacity at low fields. The blue points mark the temperature and internal field of measured peaks in the heat capacity. The fields values used are the average internal field of the sample after correcting for demagnetisation effects. The solid red line is a guide to the eye. FM and PM correspond to “spontaneous ferromagnetic order” and paramagnetism. Strictly speaking, the FM label should only apply at $\mu_0 H_{\text{int}} = 0$ T and $T \leq T_c$. represents the phase boundary between the high temperature correlated paramagnetic phase and the low temperature ordered phase for the fields at which a peak in the heat capacity is observed. The dashed red line represents the field range between the largest field at which the a peak in the heat capacity is observed, and the next measured internal field, where the peak is no longer present.

ness of the peak, making differentiation between first order and continuous transitions difficult.

In the proposed scenario for the low temperature phase of $\text{Yb}_2\text{Ti}_2\text{O}_7$ of Chang et al. [33], the ferromagnetic phase has six possible domains, corresponding to a net magnetic moment along each of the $\langle 001 \rangle$ directions. An applied field along [001] will select one domain, removing the necessity for a phase transition as the symmetry of the high temperature paramagnetic phase has already been polarised by the field, breaking the symmetry of the paramagnet, so that the phase transition will be replaced by a crossover. If the phase transition is first order, it can be expected to survive to finite field, but this field is very small. Alternatively, the persistence of the

peak in finite fields may simply be due to the fact that the magnetic field inside the sample is not uniform, and much smaller internal fields may be sufficient to suppress the transition if they are uniform in the sample.

Griffiths [82] studied the case of a ferromagnet in an applied field, finding that the applied field results in the studied phase transition being rapidly suppressed, with the peak in the heat capacity being broadened and reduced in amplitude upon increasing the field, similar to the effect shown in Fig. 3.3(a). The case of $\text{Yb}_2\text{Ti}_2\text{O}_7$ is not identical to that studied by Griffiths [82], as that work studied a ferromagnet that undergoes a continuous phase transition, whereas in the case of $\text{Yb}_2\text{Ti}_2\text{O}_7$ the phase transition has been proposed to be first order [33, 37]. The two situations may still be comparable because for the $C(T)$ scans in Fig. 3.3(a), the field has been applied along one of the directions along which the ordered phase is expected to have a net moment.

Using the peak as an indication of the presence of a phase boundary, a simple temperature-magnetic field phase diagram for $\text{Yb}_2\text{Ti}_2\text{O}_7$ in a magnetic field along [001] can be constructed, as shown in Fig. 3.3(b). The magnetic field values used in the phase diagram are the average internal magnetic field of the sample, after correcting for demagnetisation effects. The phase diagram is empirical, and is not intended to differentiate between points in the phase diagram where the peak in the heat capacity represents a first order phase transition versus a crossover.

The only other work in the literature to present measurements of the heat capacity in applied magnetic fields is that of Robert et al. [40]. They found that applying a magnetic field to a single crystal of $\text{Yb}_2\text{Ti}_2\text{O}_7$ of unknown orientation resulted in the temperature of the transition being increased significantly (an applied field of 100 mT lead to a peak temperature of ~ 300 mK, as opposed to 175 mK in zero field), while the size of the peak is reduced in increasing field. This result is consistent with the magnetisation measurements of L’Hotel et al. [37] in a field applied along the [110] direction of a single crystal sample, who found that a field applied in that direction leads to an increased temperature for the transition observed in the magnetic

susceptibility.

3.5 Schottky-Like Anomaly

In addition to the sharp peak seen in the heat capacity in Fig. 3.3(a), in the presence of even the smallest applied field, a second distinct feature in the heat capacity can be seen, a broad hump, the peak of which increases in temperature with increasing applied magnetic field. As discussed in Ch. 2.2.3, such peaks are usually due to the presence of a gap to excited states in the system, and are referred to as Schottky anomalies. The fact that such a feature does not exist in the zero field heat capacity implies that in zero field the magnetic excitations are either gapless or have a very small gap.

Figure 3.4 shows the measured heat capacity of $\text{Yb}_2\text{Ti}_2\text{O}_7$ for applied fields $\mu_0 H = B \geq 0.125$ T along [001]. Lower fields are not shown as the heat capacity at those fields includes two features (the sharp peak and the broad hump) and it is difficult to separate the contributions of the two. The heat capacity shown is only the magnetic component of the heat capacity, obtained by subtracting the high field ($B = 1.5$ T) heat capacity from the data; this high field data is considered to have a minimal magnetic contribution as the large applied field has shifted the hump completely out of the temperature range measured. Extrapolating based on the change in the peak temperature as a function of field at low fields leads to a peak temperature of ~ 3 K, while estimates based on the spin wave spectrum measured by neutron scattering (See Fig. 5.2(f) in Ch. 5.2.1) give a peak temperature of ~ 2 K, both of which are well outside the temperature window studied. In order to subtract the high field data from the other data sets, a linear fit was used to parametrise the data, and this fit was subtracted from the lower field data sets.

The obvious excitations to expect in a ferromagnetically ordered system are magnons, as discussed in Chapter 6. Using the exchange Hamiltonian of Ch. 6, based on fits to high field magnon dispersions, the heat capacity resulting from the magnon dis-




Figure 3.4: (a-e) Magnetic heat capacity of $\text{Yb}_2\text{Ti}_2\text{O}_7$, along with (solid green lines) calculations of the spin wave heat capacity, based on the results of fits to neutron scattering, and (solid red lines) the result of fits to the data using a two level Schottky anomaly model (See Eqn. (2.14)). The inset in (d) shows $\ln(CT^2)$ vs $1/T$, computed using the heat capacity data from (d). The dashed red line in the inset shows the linear fit used to determine the gap Δ between the ground and excited states of the model. The red lines in (a-e) are based on these fits. (f) The values of the excitation gap Δ extracted from fits to $\ln(CT^2)$ vs $1/T$ as a function of the internal magnetic field of the sample, showing a non-linearly increasing gap in increasing applied field, the dashed line is guide to the eye.

persion relations shown in Fig. 6.10 at low fields (computed by Dr. Paul McClarty) are shown as solid green lines in Fig. 3.4(a-e), compared to the measured magnetic heat capacity. From this comparison it can be seen that the magnon heat capacity is a poor description of the hump in the heat capacity for all of the fields measured up to $\mu_0 H_{\text{int}} = 0.39$ T, though the fit does appear to be improving at $\mu_0 H_{\text{int}} = 0.39$ T.

In an attempt to determine what type of excitations lead to the observed hump in the heat capacity, various excitation density of states were used to attempt to fit the observed hump in $C(T)$, but none of the proposed excitation spectra provided a good description. Instead, a simplified model has been used to parametrise the data and extract a notional physical parameter that characterises the excitations, namely Δ the energy gap to the lowest energy excited state. The simplified model used is a two level system with a non-dispersive excitation. The heat capacity of this model is described by Eqn. (2.14). The established method in the literature [67] for extracting an energy gap for the excitation based on this model from heat capacity data is to fit a straight line to $\ln(CT^2)$ as a function of $1/T$ below the peak of the Schottky anomaly. If $k_B T \ll \Delta$, then Eqn. (2.14) can be approximated as

$$C_{\text{Schottky}} \approx N_A k_B \left(\frac{\Delta}{k_B T} \right)^2 \exp \left(-\frac{\Delta}{k_B T} \right), \quad (3.1)$$

assuming non-degenerate ground and excited states. Taking the logarithm of Eqn. (3.1) multiplied by T^2 results in

$$\ln(CT^2) \approx \ln \left(N_A k_B \left(\frac{\Delta}{k_B} \right)^2 \right) - \frac{\Delta}{k_B T} \quad (3.2)$$

$$\approx a - b \frac{1}{T} \quad (3.3)$$

which shows that for temperatures sufficiently below the excitation gap, there is a simple linear relationship between $\ln(CT^2)$ and $1/T$, and the slope of the resulting fit will be the excitation gap in units of temperature. The inset of Fig. 3.4(d) shows such a fit, performed using the heat capacity data measured in an average internal field of

$\mu_0 H_{\text{int}} = 0.194$ T, shown in the main panel of Fig. 3.4(d) to Eqn. (3.3) (varying a and b parameters).

The excitation energy gap determined from fitting the data shown in Fig. 3.4(a-e) is shown at the top of each panel and plotted as a function of the internal magnetic field of the sample in Fig. 3.4(f). The results of using these fitted energy gaps in Eqn. (2.14) to compute the heat capacity are shown as solid red lines in Fig. 3.4(a-e). It can be seen that for temperatures much less than the peak temperature of the hump in the data, the Schottky anomaly is a reasonable approximation of the data, though the fit breaks down at and above the temperature of the peak of the hump. This is expected due to oversimplification of the model, in that it only considers two non-dispersing energy levels. Examining low field inelastic neutron scattering measurements shown in Fig. 5.2, it is clear that the observed excitations are not only a single non-dispersing excitation, though the lower bound of the scattering along $(h, 0, 0)$ is dominated by an intense non-dispersing mode in small applied magnetic fields. The existence of an intense non-dispersing mode lends some credence to the use of the simple two level model with a non-dispersing excitation to fit the hump in the heat capacity, while the existence of the higher energy continuum scattering may explain the discrepancies at temperatures above the peak of the hump. The notional excitation gaps extracted from the heat capacity data will be compared to energy gaps extracted directly from neutron scattering measurements in Chapter 6.

Examining Fig. 3.4(f), it can be seen that the gap increases non-linearly as a function of the internal magnetic field. The rate of increase at very low fields is particularly sharp, as the zero field data shows no evidence of a hump, implying that in zero field, the magnetic excitations are gapless within error. It will be shown in Ch. 5 that inelastic neutron scattering also finds that the zero field magnetic excitations are essentially gapless (within error).

3.6 Conclusions

In this chapter it has been shown that the heat capacity of the single crystal sample of $\text{Yb}_2\text{Ti}_2\text{O}_7$ used in this work shows a single sharp peak at $T_c = 214(2)$ mK in zero field. This feature is consistent with, even sharper and at occurring at slightly higher temperature, compared to the peak observed by Chang et al. [33] in their single crystal produced using similar methods. This single sharp peak means that the sample used in this work is of very good quality, based on the work of Chang et al. [33], who showed that samples with sharper single peaks in their heat capacity had reduced vacancies of Yb^{3+} , Ti^{4+} , and O^{2-} ions compared to samples with smaller peaks in their heat capacities.

Additionally, it has been shown that applying a very small magnetic field ($\mu_0 H_{\text{int}} \sim 0.018$ T) along the [001] direction rapidly suppresses the peak observed in the heat capacity. This peak is associated with a first order phase transition into a ferromagnetically ordered low temperature phase [37], with a net moment along one of the $\langle 100 \rangle$ crystallographic directions. The applied field has no significant effect on the temperature of the transition, only reducing the size of the peak observed in the heat capacity. This behaviour is consistent with the phase transition being to a phase with a net moment along [001], as the magnetic field polarises the high temperature paramagnetic phase, rapidly turning the phase transition into a crossover as in the presence of the applied field there is no longer any spontaneous symmetry breaking occurring, i.e. the high temperature paramagnet already has a field induced polarisation along the field direction.

Finally, the observed hump in the finite field heat capacity of $\text{Yb}_2\text{Ti}_2\text{O}_7$ has been fit using a simple two level Schottky anomaly model, in order to extract a notional energy gap to the lowest energy excitation in the system. The extracted gap increases in applied magnetic field, and will be compared to the energy gap in the inelastic neutron scattering in applied magnetic field, shown in Chapter 5.

Chapter 4

Neutron Scattering

This chapter provides an overview of why neutron scattering is a useful technique for condensed matter physics, and how neutrons probe structural and magnetic order, and magnetic excitations. The properties of neutrons that make them particularly useful for probing condensed matter systems are discussed, followed by a description of how the quantity measured in experiment, the number of neutrons scattered into a solid angle per unit time is related to how neutrons interact with ionic nuclei and magnetic moments. In particular the theory of how the magnetic moments present in a system give rise to both elastic and inelastic neutron scattering is discussed. After this, a brief description of the technique behind time of flight neutron spectrometers operate is provided, as all of the neutron scattering measurements in this thesis were performed using this technique. The production of neutrons for scattering is also discussed, along with the relatively new technique of performing “Horace” scans, multi-sample-orientation scans that allow for the construction of four dimensional (three reciprocal space dimensions and energy) data sets that have several advantages over traditional single orientation inelastic neutron scattering data sets.

4.1 Introduction

Neutron scattering is a very versatile tool in the investigation of condensed matter systems, and magnetic systems in particular. There are several aspects of neutrons that make them particularly suited for this purpose, which are laid out very well in Squires [83]. The first such property is that the mass of the neutron leads to a de Broglie wavelength for thermal neutrons that is on the order of the inter-atomic distance in solids and liquids, that is, Angstroms. This means that interference can occur between the wave functions describing the neutrons being scattered by the ions in a material, yielding information about the material's structure. The second important property is that neutrons do not possess a charge. This allows the neutrons to penetrate into the bulk of materials, unlike charged particles, which would be scattered by interaction with the charged particles (protons and electrons) in the atoms that make up the material being studied and only penetrate a small distance under the surface. For the same reason, the lack of a charge on the neutron means that it can closely approach the nuclei and interact with the nucleus via nuclear forces rather than electromagnetic forces. This results in the scattering cross section for neutrons of various nuclei being different from that for x-ray scattering, allowing neutrons to probe the structures of materials where x-rays cannot. Another useful property of neutrons is that the energy of thermal neutrons is similar to that of the excitations in materials, meaning that neutrons are ideally suited to probe these excitations. Finally, neutrons possess a magnetic moment, which allows them to weakly interact with unpaired electrons in materials. This makes neutrons useful for obtaining information about the spins and spatial distributions of the electrons in materials, and it means that information about magnetic excitations in materials can be gained through the measurement of inelastic neutron scattering, which is the primary use of neutrons in this thesis. This chapter will provide a basic overview of the key aspects of the theory of neutron scattering by magnetic materials, which are explained in detail in Squires [83] and Lovesey [84].

4.2 Scattering Theory

The quantity which is measured in a neutron scattering experiment is the number of neutrons scattered per second by a sample into a detector. Fig. 4.1 shows the basic layout of such an experiment. In order to derive any understanding about the system being studied, it is necessary to understand how this scattering arises.

To begin, the scattered intensity in neutrons per second for neutrons scattered into a detector with energy detection range $\Delta E'$ and solid angle detection range $\Delta\Omega$ is defined as

$$I = \Phi \eta \frac{d\sigma^2}{dE' d\Omega} \Delta E' \Delta\Omega, \quad (4.1)$$

where Φ is the incident neutron flux (number of neutrons incident on the sample per second), and η is the efficiency of the detector. The quantity $\frac{d\sigma^2}{dE' d\Omega}$ is called the partial differential cross-section, and gives the fraction of incident neutrons with some incident energy E scattered within the solid angle range $d\Omega$ and in the energy range E' to $E' + dE'$ [84].




Figure 4.1: Schematic diagram of the geometry of a basic neutron scattering experiment, showing incident and scattered neutrons, the neutron source, and detector. Also shown is the scattering triangle the reciprocal space vectors $\mathbf{k}$, $\mathbf{k}'$ and $\mathbf{Q}$.

The neutrons incident on the sample are defined by their wave-vector $\mathbf{k}$ and spin state σ , and when they are scattered, leave the system with wave-vector $\mathbf{k}'$ and spin state σ' . Together, these two variables define the state of the neutron, and can be

expressed as the state $|\mathbf{k}\sigma\rangle$ for the incident state, and $|\mathbf{k}'\sigma'\rangle$ for the scattered state. The state of the sample is defined as $|\lambda\rangle$ before the scattering event and $|\lambda'\rangle$ afterwards. If we treat the neutron and sample as one system, they are described by the product states $|\lambda\mathbf{k}\sigma\rangle$ and $|\lambda'\mathbf{k}'\sigma'\rangle$ before and after the scattering of the neutron respectively. If the potential via which the neutron interacts with the sample, V , is assumed to be weak enough to be treated as a perturbation to the system, then the scattering rate can be computed using Fermi's Golden Rule giving [83, 84]

$$W_{\mathbf{k}\sigma\lambda\rightarrow\mathbf{k}'\sigma'\lambda'} = \frac{2\pi}{\hbar} \rho_{\mathbf{k}'} |\langle \mathbf{k}'\sigma'\lambda' | V | \mathbf{k}\sigma\lambda \rangle|^2, \quad (4.2)$$

where $\rho_{\mathbf{k}'}$ is the density of scattered states per unit energy. Using the functional form of $\rho_{\mathbf{k}'}$ and normalising $W_{\mathbf{k}\sigma\lambda\rightarrow\mathbf{k}'\sigma'\lambda'}$ by the flux of incident neutrons gives [84]

$$\frac{d\sigma}{d\Omega_{\mathbf{k}\sigma\lambda\rightarrow\mathbf{k}'\sigma'\lambda'}} = \frac{k'}{k} |\langle \mathbf{k}'\sigma'\lambda' | V | \mathbf{k}\sigma\lambda \rangle|^2. \quad (4.3)$$

If the energy of the scattering system is allowed to change, that is $E_\lambda \neq E_{\lambda'}$, then

$$\frac{d^2\sigma}{d\Omega dE'_{\mathbf{k}\sigma\lambda\rightarrow\mathbf{k}'\sigma'\lambda'}} = \frac{k'}{k} \left(\frac{m}{2\pi\hbar^2} \right)^2 |\langle \mathbf{k}'\sigma'\lambda' | V | \mathbf{k}\sigma\lambda \rangle|^2 \delta(\hbar\omega + E_\lambda - E_{\lambda'}), \quad (4.4)$$

where the delta function ensures that the energy of the total neutron sample system is conserved. $\hbar\omega$ is the energy transferred from the neutron to the sample, and is defined as

$$\hbar\omega = \frac{\hbar^2}{2m} |k'^2 - k^2|, \quad (4.5)$$

where m is the mass of the neutron.

Finally if Eqn. (4.4) is summed over all of the initial and final states of the scattering system, and the initial and final spin states of the neutron, the partial differential cross-section is obtained [84]:

$$\frac{d^2\sigma}{d\Omega dE'} = \frac{k'}{k} \left(\frac{m}{2\pi\hbar^2} \right)^2 \sum_{\lambda,\lambda'} \rho_\lambda \sum_{\sigma,\sigma'} \rho_{\sigma'} |\langle \mathbf{k}'\sigma'\lambda' | V | \mathbf{k}\sigma\lambda \rangle|^2 \delta(\hbar\omega + E_\lambda - E_{\lambda'}). \quad (4.6)$$

Here ρ_λ is the probability of the scattering system being in a state λ when the neutron is scattered, and ρ_σ is the probability that the incident neutron has spin σ . It should be noted that the neutron is spin half, so that the sum over σ is over only two states.

In order to proceed any further in understanding the theory behind the scattering of neutrons by any system, it is necessary to define the scattering potential, V , which varies depending on the type of scattering system. The next two sections will discuss two different types of scattering: nuclear scattering and magnetic scattering. Before introducing a specific form of scattering potential, there is one further step that can be taken. If the scattering potential is not dependent on the momentum of the incident neutron, $\mathbf{k}$, which is true for all of the forms of scattering discussed in this work, then the Fourier transform of the scattering potential can be taken without specifying the form of V . This is useful as it gives the partial differential cross-section as a function of $\mathbf{Q} = \mathbf{k} - \mathbf{k}'$, the quantity which neutron scattering is traditionally plotted as a function of. The Fourier transformed potential is defined as:

$$V(\mathbf{Q}) = \int V(\mathbf{R}) \exp(i\mathbf{Q} \cdot \mathbf{R}) d^3R. \quad (4.7)$$

This gives the final form of the partial differential cross-section:

$$\frac{d^2\sigma}{d\Omega dE'} = \frac{k'}{k} \left(\frac{m}{2\pi\hbar^2} \right)^2 \sum_{\lambda, \lambda'} \rho_\lambda \sum_{\sigma, \sigma'} \rho_\sigma |\langle \sigma' \lambda' | V | \sigma \lambda \rangle|^2 \delta(\hbar\omega + E_\lambda - E_{\lambda'}), \quad (4.8)$$

where the $\mathbf{k}, \mathbf{k}'$ dependence of the expectation value of the potential operator has been removed, as the potential is by definition not dependent on the momentum of the neutron.

4.2.1 Nuclear Scattering

The partial differential cross-section for scattering of neutrons via the interaction with the atomic nuclei in materials is a useful quantity, as it allows neutrons to be used to

gather information about the crystal structure of materials.

The interaction potential for nuclear scattering is given by [83]:

$$V_j(\mathbf{Q}) = \frac{2\pi\hbar^2}{m} b_j. \quad (4.9)$$

Here m is once again the mass of the neutron, and b_j is the scattering length of the j th nucleus. The scattering length is an abstraction of the ability of a nucleus to scatter neutrons. b varies strongly on an element by element basis.

Neutron scattering based on this form of interaction potential takes two forms, referred to as coherent and incoherent scattering [83]:

$$\frac{d^2\sigma}{d\Omega dE'} = \left(\frac{d^2\sigma}{d\Omega dE'} \right)_{coh} + \left(\frac{d^2\sigma}{d\Omega dE'} \right)_{inc} \quad (4.10)$$

By using the potential form in Eqn. (4.9) and changing the delta function in Eqn. (4.8) to its integral form, these two terms can be expressed as [83]:

$$\left(\frac{d^2\sigma}{d\Omega dE'} \right)_{coh} = \frac{\sigma_{coh}}{4\pi} \frac{k'}{k} \frac{1}{2\pi\hbar} \sum_{j,j'} \int_{-\infty}^{\infty} \langle \exp\{-i\mathbf{Q} \cdot \mathbf{R}_{j'}(0)\} \exp\{i\mathbf{Q} \cdot \mathbf{R}_j(t)\} \rangle \exp(i\omega t) dt, \quad (4.11)$$

and

$$\left(\frac{d^2\sigma}{d\Omega dE'} \right)_{inc} = \frac{\sigma_{inc}}{4\pi} \frac{k'}{k} \frac{1}{2\pi\hbar} \sum_j \int_{-\infty}^{\infty} \langle \exp\{-i\mathbf{Q} \cdot \mathbf{R}_j(0)\} \exp\{i\mathbf{Q} \cdot \mathbf{R}_j(t)\} \rangle \exp(i\omega t) dt. \quad (4.12)$$

In these forms, the sum over the states λ , and λ' are contained in the statistical average $\langle \dots \rangle$, and the sum over the neutron polarisation is dropped, as the nuclear scattering potential has no spin dependence. The two scattering cross-sections, σ_{coh} , and σ_{inc} are defined as [83]

$$\sigma_{coh} = 4\pi(\bar{b})^2, \quad (4.13)$$

and

$$\sigma_{inc} = 4\pi\{\bar{b}^2 - (\bar{b})^2\}, \quad (4.14)$$

where the overline indicates an average over all the nuclei in the scattering system. The important features to note are that the coherent scattering in Eqn. (4.11) arises from scattering from nuclei at different positions $\mathbf{R}_j$ and $\mathbf{R}'_j$ at different times in addition to the same nucleus at different times when $j = j'$, leading to interference effects, which can provide information about the physical structure of the systems [83]. The incoherent term comes only from the same nucleus at different times, and cannot generate interference terms, only a flat background [83].

4.2.2 Magnetic Scattering

The partial differential cross-section of most relevance to this thesis is the cross section for magnetic scattering. The origin of this scattering is the interaction of the magnetic dipole moment of the neutron with magnetic fields present in the sample. These magnetic fields have two sources, the spin of unpaired electrons in ionic orbitals, and the angular momentum of the orbitals of those same electrons. The magnetic field arising from the moments of unpaired spins is given by [83]:

$$\mathbf{B}_S = \nabla \times \mathbf{A}, \quad (4.15)$$

where

$$\mathbf{A} = \frac{\mu_0}{4\pi} \frac{\boldsymbol{\mu}_e \times \hat{\mathbf{R}}}{R^2}. \quad (4.16)$$

Here $\boldsymbol{\mu}_e$ is the electron magnetic moment, $\mathbf{R}$ is the vector that describes an electrons position relative to the nucleus, and $\hat{\mathbf{R}}$ is a unit vector along the direction of $\mathbf{R}$. The magnetic field in the sample arising from the orbital angular momentum of the unpaired electrons is given by [83]:

$$\mathbf{B}_L = \frac{-\mu_B \mu_0}{2\pi} \frac{\mathbf{p} \times \hat{\mathbf{R}}}{R^2}, \quad (4.17)$$

where μ_B is the Bohr magneton, and $\mathbf{p}$ is the momentum of the electron in its orbital. If we combine these two fields into the quantity $\mathbf{B} = \mathbf{B}_S + \mathbf{B}_L$, then the potential

describing the interaction of the neutrons magnetic moment with the magnetic fields inside the sample is given by [83]:

$$V = \boldsymbol{\mu}_n \cdot \mathbf{B}, \quad (4.18)$$

where $\boldsymbol{\mu}_n$ is the magnetic moment of the neutron, equal to $-\gamma\mu_N\boldsymbol{\sigma}$. $\gamma = 1.913$ is the gyromagnetic ratio, μ_N is the nuclear magneton, and $\boldsymbol{\sigma}$ is the Pauli spin operator for the neutron [83].

In this work, the system of interest is one where the magnetism has both spin and orbital angular momentum components, with large spin orbit coupling. The electronic configuration of Yb^{3+} (the magnetic ion in $\text{Yb}_2\text{Ti}_2\text{O}_7$) is $4f^{13}$, with a single hole in the valence orbital [12]. The ground state manifold of this state is $^2F_{7/2}$, so $J = 7/2$ is the good quantum number which describes the system. This means both terms in in $\mathbf{B}$ must be kept in the calculation of the magnetic cross-section.

In order to proceed further, certain assumptions about the electrons in the system must be made. In this work, the magnetism of interest is due to localised moments. Based on this, it can be assumed that the electrons are localised in position near their associated atomic nuclei. For simplicity, it is assumed that these nuclei are located in a crystal on a Bravais lattice ($\text{Yb}_2\text{Ti}_2\text{O}_7$ has an additional sublattice structure, but this does not fundamentally change the neutron scattering). Additionally, the presence of spin-orbit coupling is assumed, so that the electron and orbital angular momenta combine to form total spin and orbital angular momenta, L and S , which combine to form the total angular momentum $J = L + S$ [83]. In order to obtain a simple form for the magnetic scattering due to magnetic moments described by J , the dipole approximation [83, 84] is used, which requires that $|\mathbf{Q}|^{-1}$ is large compared to the mean radius of the electron orbitals. Finally, for simplicity, it is assumed that all of the magnetic ions in the system are identical. Using all of these assumptions and approximations, and inserting the magnetic scattering potential given in Eqn. (4.18) into Eqn. (4.8), the resulting partial differential cross-section for unpolarised neutrons

is given by [83]:

$$\frac{d^2\sigma}{d\Omega dE'} = \left(\frac{\gamma\mu_0 e^2}{4\pi m_e} \right)^2 \frac{k'}{k} \exp(-2W) \frac{g_J^2}{4} |F(Q)|^2 \sum_{\alpha,\beta} \left(\delta_{\alpha\beta} - \hat{Q}_\alpha \hat{Q}_\beta \right) \mathcal{S}^{\alpha\beta}(\mathbf{Q}, \omega). \quad (4.19)$$

Here m_e is the mass of the electron, and α and β represent the Cartesian directions x, y, z . $g_J = (3J(J+1) + S(S+1) - L(L+1)) / (2J(J+1))$ is the Landé g-factor, equal to 8/7 for a free Yb^{3+} ion. $F(Q)$ is the magnetic form factor, which is the Fourier transform of the spatial density of the unpaired electrons in the system [83]. This term is responsible for a decrease in the intensity of the magnetic scattering with increasing $|\mathbf{Q}|$. The term $\exp(-2W)$ is referred to as the Debye-Waller term, and accounts for thermal motion of the ions in the system. W is Q dependent, so this term has a similar effect to the magnetic form factor. W is also temperature dependent, such that this term has less and less effect on the intensity of the scattered neutrons as the temperature of the scattering system is lowered. As such it is neglected in this work, as all of the neutron scattering measurements in this work take place at millikelvin temperatures and at small Q (where magnetic scattering is strong). $\hat{Q}_\alpha$ represents the α component of the unit vector $\hat{Q}$ in the direction of $\mathbf{Q}$.

The most important features of Eqn. (4.19) are the term $(\delta_{\alpha\beta} - \hat{Q}_\alpha \hat{Q}_\beta)$, which indicates that the neutron can only couple to magnetic correlations perpendicular to $\mathbf{Q}$, and the quantity $\mathcal{S}^{\alpha\beta}(\mathbf{Q}, \omega)$, which is the magnetic correlation function and contains the information about the magnetic structure and excitations in the system that are accessible in neutron scattering experiments. In a manner analogous to that used for the nuclear scattering, the magnetic correlation function can be expressed in terms of the total angular momentum as [83]:

$$\mathcal{S}^{\alpha\beta}(\mathbf{Q}, \omega) = N \sum_j \exp(i\mathbf{Q} \cdot \mathbf{R}_j) \frac{1}{2\pi\hbar} \int_{-\infty}^{\infty} \langle J_0^\alpha(0) J_j^\beta(t) \rangle \exp(-i\omega t) dt, \quad (4.20)$$

where $J_j(t)$ is the time dependent total angular momentum operator for lattice site j , and $\alpha = x, y, z$ specify the Cartesian components of that operator. $\langle \dots \rangle$ is the average over the occupation of the various states of the scattering system. The term

$\langle J_0^\alpha(0)J_j^\beta(t) \rangle$ can be split into two parts, its value at $t = \infty$ and a time varying term [83], where the $t = \infty$ component gives rise to elastic magnetic scattering, and the time varying term gives rise to inelastic magnetic scattering.

A more useful form for the magnetic correlation function is [84]:

$$\mathcal{S}^{\alpha\beta}(\mathbf{Q}, \omega) = \sum_{\lambda, \lambda'} \rho_\lambda \langle \lambda | J^\alpha(\mathbf{Q}) | \lambda' \rangle \langle \lambda' | J^\beta(\mathbf{Q}) | \lambda \rangle \delta(E_\lambda - E_{\lambda'} + \hbar\omega), \quad (4.21)$$

where ρ_λ is the probability that the scattering system is in the state λ , the delta function ensures the conservation of energy in the neutron + sample system, and $J^\alpha(\mathbf{Q})$ is the lattice Fourier transform of the α component of the magnetic moment $\mathbf{J}$. This quantity is defined as [84]:

$$J^\alpha(\mathbf{Q}) = \sum_j e^{-i\mathbf{Q} \cdot \mathbf{R}_j} J_j^\alpha, \quad (4.22)$$

where the sum is over all j magnetic ions, whose positions in the lattice are given by $\mathbf{R}_j$. The form of $J^{\alpha\beta}(\mathbf{Q}, \omega)$ in Eqn. (4.21) is convenient for the calculation of inelastic neutron scattering arising from linear spin wave theory, which will be discussed later in this work.

A final aspect to discuss about magnetic neutron scattering is the relationship between magnetic long range order and elastic neutron scattering. As mentioned previously, elastic magnetic scattering arises from the $t = \infty$ limit of Eqn. (4.20). Analogous to the discussion in Section 4.2.1, if the magnetic moments are on a regular lattice, and form some kind of order (Ferromagnetic, Anti-ferromagnetic, etc.), it gives rise to a series of magnetic Bragg peaks in the elastic neutron scattering. The partial differential cross-section for elastic magnetic scattering can be expressed as [83, 84]:

$$\left(\frac{d\sigma}{d\Omega} \right)_{\text{el}} = \left(\frac{\gamma\mu_0 e^2}{4\pi m_e 2\mu_B} \right)^2 |\hat{\mathbf{Q}} \times \{\mathbf{M}(\mathbf{Q}) \times \hat{\mathbf{Q}}\}|^2. \quad (4.23)$$

Here, $\mathbf{M}(\mathbf{Q})$ is the Fourier transform of $\mathbf{M}(\mathbf{r})$, the magnetisation density operator, which includes both spin and orbital components. The series of cross products with

the wave vector transfer $\mathbf{Q}$ indicates that, as previously noted, only magnetic correlations perpendicular to $\mathbf{Q}$ are coupled to the neutron. If we consider a single domain Ferromagnetic crystal, where all the magnetic moments are ordered in one direction, then $\mathbf{M}(\mathbf{r})$ is periodic with the unit cell of the crystal. This means that it is possible to write [83]:

$$\mathbf{M}(\mathbf{Q}) = \int \mathbf{M}(\mathbf{r}) \exp(i\mathbf{Q} \cdot \mathbf{r}) d\mathbf{r} = \mathcal{F}(\mathbf{Q}) \sum_j \exp(i\mathbf{Q} \cdot \mathbf{R}_j), \quad (4.24)$$

where

$$\mathcal{F}(\mathbf{Q}) = \int_{\text{cell}} \mathbf{M}(\mathbf{r}) \exp(i\mathbf{Q} \cdot \mathbf{r}) d^3r. \quad (4.25)$$

The $\int_{\text{cell}}$ indicates that the integral takes places over a single unit cell, as this is the volume over which $\mathbf{M}(\mathbf{r})$ is periodic, and after this, the integral over all space is just a sum over the number of unit cells. $\mathcal{F}(\mathbf{Q})$ is the magnetic unit cell structure factor [84], which contains all of the information about the spatial structure of the magnetisation in the system. By inserting Eqn. 4.24 into Eqn. 4.23, it is possible to obtain [83]:

$$\left(\frac{d\sigma}{d\Omega} \right)_{\text{el}} = \left(\frac{\gamma\mu_0 e^2}{4\pi m_e 2\mu_B} \right)^2 N \frac{(2\pi)^3}{v_0} \exp(-2W) \sum_{\boldsymbol{\tau}} \delta(\mathbf{Q} - \boldsymbol{\tau}) |\hat{\boldsymbol{\tau}} \times \{\mathcal{F}(\boldsymbol{\tau}) \times \hat{\boldsymbol{\tau}}\}|^2, \quad (4.26)$$

where v_0 is the volume of the unit cell, N is the number of unit cells in the sample, and $\boldsymbol{\tau}$ are the reciprocal lattice vectors of the crystal lattice of the system being probed. From this it can be seen that the elastic neutron scattering that arises from magnetic order in a Ferromagnetic system results in peaks in the number of neutron being scattered at the reciprocal lattice vector locations, the same locations where the coherent nuclear scattering also has peaks in the intensity of the scattered neutrons. In the case of Antiferromagnetic systems, where one must consider a larger unit cell than the structural unit cell, and $\mathbf{M}(\mathbf{r})$ becomes the staggered magnetisation, peaks occur at reciprocal lattice vectors for the larger unit cell, in positions not occupied by peaks in the coherent nuclear scattering.

4.3 Time of Flight Neutron Spectroscopy

All of the neutron scattering experiments in this work were performed using the direct time of flight inelastic neutron spectrometer LET at the ISIS neutron source [85]. Furrer et al. [86] provides a basic introduction to how time of flight spectrometry works, a brief description will be provided here.

Time of flight spectrometers work with fixed incident (direct geometry), or scattered (indirect geometry) energy neutrons. For direct geometry the incident beam energy is defined by a series of beam choppers (referred to as monochromators) that filter neutrons by energy. On LET, these monochromators take the form of pairs of counter rotating disks coated in neutron absorbing material placed in the beam. These disk each have a small window that allow neutrons to pass through when it crosses the beam. By controlling the relative phase of two pairs of counter rotating disks, the energy of the incident neutrons may be selected.

If the energy of the neutron in some section of the scattering path is fixed (in the direct geometry case, this is the incident energy), then the total time the neutron spends traversing the instrument, along with the angle at which the neutron is scattered is sufficient to fix the vector momentum and energy of the neutron on the side of the experiment where it is not fixed. Using this relationship, the arrival time of scattered neutrons relative to the arrival time of a short pulse of incident neutrons at some specified point in the instrument is used to work out the wave-vector of the scattered neutron ($\mathbf{k}'$) and how much energy it gained from or lost to excitations in the sample. Figure 4.2 shows a basic schematic diagram of the instrument, specific details can be found in Bewley et al. [85].

To work out the arrival time, and scattering angle, the scattered neutrons are detected by position sensitive detectors. These detectors take the form of long steel tubes, filled with a combination ^3He and other gasses, with a long wire running down the centre of the tube. When a neutron passes through the tube, if it impacts a ^3He




Figure 4.2: Schematic diagram of a direct geometry time of flight spectrometer showing the basic geometry of the experiment. Incident neutrons are produced by the neutron source, then a series of choppers select neutrons with a specific energy. The neutrons then interact with a sample placed in the beam, and are scattered into an array of position sensitive detectors.

atom, the following reaction takes place [86]:



The resulting hydrogen nuclei are charged and ionise the surrounding gas. By applying a large voltage between the tube (cathode) and the wire running down the tube (anode), the electrons released by ionisation can be detected as a current pulse in the central wire. The gasses used are chosen so that many electrons (up to 10^5 [86] per event) are produced per ${}^3\text{He}$ ionisation, making the detection of single neutrons possible. If the central wire has a finite resistance, the relative arrival time of this current pulse at the two ends of the wire can be used to determine where in the tube the ionisation occurred, making the detector position sensitive along its length. On LET, these tubes are arranged side by side, with their long axis vertical, into a wide angle array, such that when the instrument is complete the detectors will span a 180° arc. The electrical pulse arrival time provides vertical position sensitivity, while the position of the tube in which the neutron was detected provides the horizontal

position.

Due to the curved nature of the detector banks, being circularly symmetric around the sample, the path traced by $\mathbf{Q} = \mathbf{k} - \mathbf{k}'$ will also be curved, as shown in Fig. 4.2. In addition to this curvature due to the shape the detector banks, the $\mathbf{Q}$ surface traced through reciprocal space by the detector will be curved due to the reduced energy of the scattered neutrons, which is related to the length of the scattering wave vector by $k = \sqrt{2m_n E / \hbar^2}$. This means that for a neutron scattering data set collected with a single sample orientation, plotted in the typical way where the value of $\mathbf{Q}$ is set at zero energy transferred to the system from the neutron, the value of $\mathbf{Q}$ at finite energy transfer will not be the same as for elastic scattering.

4.4 Neutron Sources

The neutrons used for inelastic scattering experiments can be produced by different methods, the two commonly used methods being nuclear fission, and nuclear spallation. Nuclear fission is the process that occurs inside nuclear reactors where an incident neutron impacts a nucleus, typically ^{235}U , causing the nucleus to decay into multiple fission products including multiple neutrons, some of which are allowed to leave the reactor and are channelled into various experiments [86]. In the case of spallation sources (the type of neutron source used at ISIS), the neutrons are produced by bombarding a target with high energy protons whose de Broglie wavelength is less than the size of the target nucleus. The effect of the incident proton on the nucleus is to leave the nucleus in a highly excited state, at which point it begins to evaporate neutrons [86], and it is these neutrons which leave the target and can be used in scattering experiments. The total process from proton impact to neutron escape from the target takes $\sim 10^{-15}\text{s}$, making the time dependence of the emitted neutrons mainly a function of the form of the incident proton pulse. The low energy part of the energy distribution of the emitted neutrons from a spallation source is very similar in form to the spectrum of the neutrons emerging from a fission source, which is given by a

maxwellian distribution [86]:

$$\phi(v) \propto v^3 \exp\left(-\frac{1}{2} \frac{mv^2}{k_B T}\right), \quad (4.28)$$

where $\phi(v)dv$ is the flux of neutrons with velocities in range dv in units of neutrons/area/second.

When the neutrons leave the spallation source, they have a characteristic temperature of ~ 300 K, and are referred to as thermal neutrons. Neutrons leaving a fission source have a much much higher characteristic temperature, and must be moderated to be used as thermal neutrons. The neutrons used to probe the inelastic scattering of materials in this thesis are cold neutrons, whose characteristic temperature is near 25 K. The temperature of the Maxwellian distribution can be varied by passing the neutrons leaving the source through a moderator, a volume of material which strongly scatters neutrons, kept at the temperature of the desired Maxwellian distribution of the neutrons. Cold neutrons are necessary in this work as the magnetic excitations being probed have very low energies ($E < 10$ meV).

4.5 Horace Scans

In addition to data sets collected with the sample in a fixed orientation, it is possible to collect multiple data sets where the orientation of the sample is changed between each data set. These data sets can then be combined into a 4-dimensional volume of data in $(\mathbf{Q}, E)$ (3 momentum components and energy) space. At the ISIS neutron source, such data sets are referred to as Horace scans [87]. The constituent runs of a Horace scan are usually counted for a much shorter time than single angle data sets. By combining the multiple data sets, it is possible to achieve neutron count statistics approximately equivalent to those of single angle data sets counted for a much longer time.

The advantages of these scans are that they avoid the problem of $(\mathbf{Q}, E)$ coupling

discussed in Sec. 4.3 by collecting the inelastic scattering over a volume of $(\mathbf{Q}, E)$ space. The data obtained from such scans also helps to eliminate the issue of the path through reciprocal space by the detectors at zero energy transfer being curved. By integrating over a small interval in the two directions in reciprocal space perpendicular to a given direction of interest, a plot of the neutron scattering as a function of energy for a straight line through reciprocal space can be obtained. Additionally, the new ability to look at the neutron scattering at fixed energy but varying $\mathbf{Q}$ is possible.

Several such Horace scans were collected in this work, and proved very useful as they provided simultaneous access to the $(\mathbf{Q}, E)$ dependence of magnetic excitations along many non-symmetry equivalent reciprocal directions in space.

Chapter 5

Inelastic Neutron Scattering

Measurements of the Magnetic

Excitations in $\text{Yb}_2\text{Ti}_2\text{O}_7$

This chapter presents the results of inelastic neutron scattering measurements on a single crystal sample of $\text{Yb}_2\text{Ti}_2\text{O}_7$ performed using the inelastic neutron spectrometer LET at the ISIS neutron source. The sample mount, alignment, and environment, which allowed the sample to be cooled to $T \sim 145$ mK and magnetic fields of up to 9 T $\parallel$ [001] to be applied are discussed, along with the details of the instrument configuration used to perform the measurement. The inelastic scattering data along $(h, 0, 0)$ collected in high fields ($B \geq 3$ T) shows four sharp magnon modes, and a broad continuum of scattering at higher energies. The bandwidth of the continuum scattering and the behaviour of its low energy boundary upon changing the magnetic field are consistent with the continuum being due to two magnon scattering processes. For fields of $B < 3$ T, the continuum and the highest energy magnon mode overlap in energy, and the magnon mode is significantly broadened, likely due to magnon decay processes. The remaining three sharp modes have their bandwidth suppressed upon decreasing magnetic field, appearing to collapse into a single degenerate mode at the low energy boundary of the scattering continuum in magnetic fields of $B < 0.75$ T.

The zero field scattering spectrum is a continuum with no sharp features and no evidence of a gap down to the smallest measured energies, extending in energy up to ~ 1.5 meV. The existence of this continuum indicates that at zero field magnons cannot propagate coherently, and instead decay into some more fundamental excitation. The large energy range of the zero field continuum (much higher than expected based on indications of the exchange energy) indicates that the continuum must be due to dispersing excitations of some kind, made up of clusters of spins, or more exotic coherently propagating quasi-particles. Properties of the inelastic neutron scattering over the full Brillouin zone at select fields are also presented and discussed.

5.1 Experimental Details

This section will go over the details of the sample and instrument configuration used in the inelastic neutron scattering works discussed in this thesis. First, the sample orientation and mount will be discussed, followed by the configuration of the time-of-flight neutron spectrometer LET, and finally the sample environment.

5.1.1 Sample Alignment and Mount

The sample used in these experiments was a large section of the crystal described in Appendix B. The crystal was cylindrical in shape, with an average diameter of 7.33(1) mm, a length of 21.84(1) mm, and a mass of 6.3 g. The crystal is shown in Figure 5.1(a) attached to the sample mount used in the experiment.

The sample was aligned such that in the experiment, the magnetic field could be applied along the [001] axis of the crystal, as shown in Figure 5.1(b). The sample mount was machined from Oxygen Free High Conductivity (OFHC) Copper to ensure good thermal conductivity at millikelvin temperatures. An OFHC Copper rod was also incorporated into the mount, which ran up the side of the sample, as shown in Fig. 5.1(b), serving as sample support and an additional cooling path. The crystal




Figure 5.1: (a) Image Furnace grown single crystal of $\text{Yb}_2\text{Ti}_2\text{O}_7$ used in the inelastic neutron scattering experiments, attached to an OFHC Copper mount using plasticine. (b) Sample adhered to the copper mount using GE varnish, with a side copper rod in place. The primary axes of the crystal are indicated by the arrows. (c) The sample mount in its final configuration, wrapped in OFHC Copper foil and anchored using OFHC wire, attached to a dilution refrigerator cold finger. Cadmium has been wrapped around the base of the sample mount to mask scattering from the sample mount.

was placed on the mount such that when the sample rotates about the centre of the $[001]$ axis, the centre of mass of the sample is stationary, and the sample remains in the centre of the neutron beam. The sample was attached to the mount using GE 1076 varnish. OFHC Copper sheeting and OFHC copper wire were then wrapped around the sample to provide further cooling pathways. Figure 5.1(c) shows the fully mounted sample, including a cadmium sheet wrapped around the base of the sample to minimise the amount of spurious signal that could come from neutrons scattering off the sample mount.

The final alignment of the sample in the experiment, determined using nuclear Bragg peak positions with respect to the LET instrument reference frame, was such that the $[001]$ crystallographic axis was almost vertical, tilted only by angle of $\theta = 1.03^\circ$ away from the vertical axis of the instrument. The projection of the tilted c^* axis in the XY plane of the instrument defines an angle $\phi = 56.465^\circ$ with the nominal (un-tilted) $[100]$ axis.

5.1.2 Sample Environment

The sample environment used in these inelastic neutron scattering experiments consists of an Oxford Instruments Kelvinox dilution refrigerator mounted inside a 9 T superconducting magnet. The superconducting magnet was used to apply fields along the [001] axis of the crystal, perpendicular to the horizontal principal scattering plane of the instrument. The magnet, made by Oxford Instruments, was capable of generating magnetic fields of up to 9 T, and has a 90° opening in the scattering plane, and a $\pm 15^\circ$ opening out of the scattering plane of the instrument, allowing access to a wide range of scattering angle 2Θ in the horizontal plane and significant out-of-plane scattering. In this experiment the magnet was positioned such that scattering could be measured for $0^\circ \lesssim 2\theta \lesssim 90^\circ$. The configuration of the magnet and the sample allowed scattering to be measured primarily in the $(h, k, 0)$ plane, and for a small range along l (out-of-plane). The dilution refrigerator was capable of a base temperature of ~ 30 mK at the mixing chamber under no load. A thermometer mounted on the OFHC Copper rod that connects the sample to the mixing chamber of the dilution refrigerator read $T \sim 145$ mK during our experiments, the increase being due to the thermal load and beam heating. This temperature is well below the temperature of the sharp peak in the heat capacity of the sample $T = 214(2)$ mK (see Fig. 3.2), so our sample was well within the low temperature phase. To ensure that the sample contained only one domain in the low temperature ordered phase, the sample was cooled in an applied field of 0.21 T.

5.1.3 Instrument Configuration

The basic design and layout of the inelastic neutron time-of-flight spectrometer LET was discussed in Chapter 4.3. This section will describe the specific configuration parameters used during the experiment. The user configurable parameters of LET are the energy of the incident neutrons, and the rotation rates of two sets of choppers (f_1 and f_2) that define the resolution of the instrument.

For certain chopper phases and frequencies, multiple monochromatic neutron pulses can arrive in the same time frame without overlap. This is called multi-repetition mode. For example, with the instrument configured for a “principal” pulse of 2.5 meV neutrons, secondary pulses of monochromatic 1.34 and 6.3 meV neutrons were also present (see Table 5.1) and produced useful data.

Table 5.1: The energy and chopper frequency configurations of LET used in this work. The last column gives the measured energy resolution (Full Width Half Maximum FWHM) of the elastic line.

Principal Pulse Energy (meV)	f_1 (Hz)	f_2 (Hz)	Monochromatic Energy Peaks in Multi-Repetition Mode (meV)	Resolution (FWHM) (meV)
2.5	160	80	1.34	0.0231(1)
			2.5	0.0555(3)
			6.3	0.2198(6)
4.0	200	100	2.0	0.0942(3)

5.2 Neutron Scattering Data

Several data sets were collected using both of the configurations in Table 5.1 where the sample was kept in a fixed orientation with respect to the incident beam direction. For these data sets, the sample was oriented such that the $\mathbf{Q}$ vectors were primarily along the $(h, 0, 0)$ direction at zero energy transfer. These single angle runs were counted for an average time of 2.5 hours (at an average proton current of $40\mu\text{A}$) to ensure that the uncertainties in the intensity were smaller than the amplitude of any features of interest. Single angle data sets were collected in applied magnetic fields of 0.0 – 9.0 T. Representative plots are shown in Fig. 5.2.

In addition to fixed orientation data sets, four Horace scans were performed, for a selection of applied magnetic fields (0, 0.21, 1.5, and 5 T). These scans spanned angular ranges of 106° for 5 T, 121° for 1.5 T, 133° for 0.21 T and 130° for 0.0 T. These scans allow access to one quadrant of the $(h, k, 0)$ plane, which is symmetry

equivalent to the other four quadrants in the $(h, k, 0)$ plane due to the four-fold rotational symmetry about the $[001]$ axis. All of these scans were collected in 1° steps, with each angle counted for approximately 7 minutes. These data sets were then combined into an $S(h, k, l, E)$ object using the Horace software, from which slices and scans were extracted.

5.2.1 Evolution of Magnetic Excitations in $[001]$ Magnetic Field

Figure 5.2 shows all of the 2.5 meV single orientation data collected during our experiment as a function of applied magnetic field. The data has had an estimate of the non-magnetic background subtracted from it. This “background” has been created by assembling the energy ranges of data where no magnetic inelastic signal is expected to be present (See Appendix E). The data has also been corrected for neutron absorption effects, which are significant for low energy neutrons (the linear attenuation coefficient of $\text{Yb}_2\text{Ti}_2\text{O}_7$ at 1.0 meV is 3.38 cm^{-1}) using the method outlined in Appendix F.

The data in Fig. 5.2 shows a dramatic evolution of the magnetic excitations as a function of applied magnetic field. In large fields ($B \geq 3 \text{ T}$) there are four sharp dispersing modes observed, consistent with magnon excitations. In the pyrochlore lattice, the magnetic Yb^{3+} ions form an FCC lattice with with a four site tetrahedral sublattice. Each corner of the tetrahedron is a distinct magnetic site and gives rise to a distinct magnon mode. Starting from 9 T, the energy of these excitations decreases with decreases upon lowering the field, as expected due to a reduction in the Zeeman energy. For applied fields below 3 T, the uppermost sharp mode in the data can be seen to broaden compared to the same mode in the $B = 3 \text{ T}$ data in the range $-2.0 \leq h \leq -1.5$, as shown in Figs. 5.2(g,h). As the field is decreased further, this mode continues to broaden, eventually evolving into a continuum of scattering above the sharp modes. The other three modes are squeezed together in energy, becoming less and less dispersive, appearing to be mostly non-dispersive and almost




Figure 5.2: Overview of the evolution of magnetic excitations as a function of applied magnetic field. Data was collected using monochromatic incident neutrons with energy 2.5 meV and is integrated over $-0.5 \leq l \leq 0.5$.

degenerate by 1.0 T (except for a small region around $\mathbf{Q} = (-2, 0, 0)$) and completely non-dispersive and degenerate by 0.75 T. This suppression of the one-magnon dispersion bandwidth is likely related to the interactions with high-energy two-magnon continuum, which will be discussed later. Finally at zero field the excitations are dominated by a broad continuum of scattering down to the lowest probed energies.

Figure 5.3 shows scans in energy generated by integrating the data in Fig. 5.2(a-h) over the range $-2.1 \leq h \leq 0.0$. In Fig. 5.3(g,h) it can clearly be seen that when the applied magnetic field is lowered from 3 T to 2 T, the uppermost magnon mode becomes significantly broadened. Figure 5.3((g) and (d)) show that the other three sharp modes move closer in energy and overlap into an apparent single sharp mode as the field is lowered from 2 T to 0.75 T. Below this field, the continuum begins to encroach on this remaining sharp mode such that at 0 T it becomes impossible to distinguish a clear sharp mode from the continuum that now extends all the way down to the lowest measured energies.

Figure 5.4(a) shows energy scans at 0, 0.21 T, and 9 T data. The 9 T data points serve to illustrate the quality of the background subtraction, as at this high field, the magnetic scattering has been pushed to much higher energies, as shown in Fig. 5.6(c) and Appendix E. The overplot of the 0 T and 0.21 T data shows several significant features: the continuum of scattering is quite broad, extending up to $E \sim 1.5$ meV; and the low energy part of the continuum has unusual properties in zero field, and undergoes a dramatic change in a small magnetic field. The broad energy range of the continuum is surprising as this is approximately twenty times the highest experimentally observed Curie-Weiss temperature of this material, which is taken as an indication of the scale of the exchange energies in the system. It is also approximately thirty times larger than the resolution of the instrument for this measurement. This implies that the continuum is due to magnetic excitations involving not simply a single spin flip, but a very large number of spin flips.

Fig. 5.4 shows the very unusual property of the of continuum at low energies, that in zero applied field, the continuum extends down to the lowest measured energy




Figure 5.3: Magnetic inelastic scattering energy scans for applied magnetic fields between 0 T and 3 T. Scans are obtained from the colour intensity maps in Fig. 5.2(a-h) by integrating over $-2.1 \leq h \leq 0.0$.




Figure 5.4: (a) Overplots of energy scans (from the data in Fig. 5.3) at 0, 0.21 T, and 9 T, integrated over $-2.1 \leq h \leq 0$ and $-0.5 \leq l \leq 0.5$. The lowest energy shown is 0.07 meV (the elastic line resolution for this incident energy is 0.0555(3) meV), the background subtraction become unreliable below this energy, due to the strong incoherent scattering in the elastic line. (b,c) The same energy scans for 0 (for reference), 0.5, 0.75 and 1 T (b) and 0, 1.5 and 2 T (c). These scans serve to illustrate how the shape of the high energy continuum is not affected by applied magnetic fields of 0 – 1 T, and is only weakly effected for 1.5 and 2 T. The dashed lines in (b,c) are guides to the eye. All data in this figure has been background subtracted and absorption corrected.

transfer, 0.07 meV, with no evidence of a spin gap (below 0.07 meV contamination by the elastic line makes the data unusable). As we shall see in Chapter 6, spin wave calculations based on a ferromagnetically ordered ground state, as found by Chang et al. [33], predict gapped excitations above 0.21 meV. This means that even though we see strong evidence for multi-magnon scattering in $\text{Yb}_2\text{Ti}_2\text{O}_7$ (see Fig. 5.6, and for multi-magnon decay, the continuum at zero field cannot be described by a simple magnon continuum as a weak perturbation on top of well defined sharp spin waves, as the predicted sharp spin waves are absent. Conspicuously, there is a “bump” in intensity at finite energy transfer, $E = 0.149(5)$ meV, but this feature is quite broad and not characteristic of a sharp excitation, moreover it cannot be separated from the continuum. By examining Fig. 5.2(a) it can be seen that this feature exists mainly near $\mathbf{Q} = 0$, away from this position, the broad “bump” fades in intensity, and the continuum is even more uniform. We note that that very recent polarised neutron scattering reported by Robert et al. [40] also showed evidence for a continuum of scattering at ~ 2 K, suggesting that the origin of the dominant excitation continuum is the same as at base temperature (145 mK) as at 2 K.

It is useful to compare the zero field continuum to the excitations in a small but finite magnetic field. An applied field on the sample of 0.21 T is equivalent to an average internal field of $\mu_0 H_{\text{int}} = 0.17$ T (see Table D.2). Fig. 5.4 shows that the application of such a small field (equivalent energy 0.02 meV for a g factor of 2) has a very large impact on the form of the magnetic scattering at low energies. The application of the weak magnetic field opens a clear gap to the lower bound of the magnetic scattering, and gives rise to a single distinct non-dispersing sharp mode with energy $E = 0.229(3)$ meV at the lower boundary of the continuum scattering.

Figure 5.2 shows this gap to be almost non-dispersive along $(h, 0, 0)$, though the intensity of the mode fades near $(-1, 0, 0)$. It is also interesting to note that while the weak applied magnetic field has a significant effect on the magnetic scattering for $E \leq 0.6$ meV, with the sharp peak being separated from the continuum by a dip in intensity. Above this energy, the continuum appears to be unaffected by this small applied field. This lack of effect of applied magnetic on the upper bound of the continuum can be seen up to applied fields of 1 T (equivalent to $\mu_0 H_{\text{int}} = 0.95$ T), as shown by Fig. 5.4(b) though the energy range that is unaffected by the field decreases in increasing field. For applied fields of 1.5 and 2 T, the form of the upper bound of the continuum scattering is affected by the applied field, but only weakly, leading to a slight increase in the high energy intensity, and the upper bound of the continuum (see Fig. 5.4(c)). For $B \geq 3$ T, the form of the continuum changes more fundamentally, as it no longer overlaps with the single magnon modes. As discussed previously, 1 T is also the applied field at which the observed sharp modes appear to become degenerate and non-dispersive along $(h, 0, 0)$ (see Fig. 5.2(e,f)). This would seem to confirm the idea of a change in the nature of the excitations above and below 1 T, especially since a sharp non-dispersing magnon mode would give rise to a sharp flat band of multi-magnon scattering, not the observed broad continuum.

Fig. 5.5 shows higher energy data (6.3 meV neutron energy) for applied magnetic fields between 2 T and 9 T. The use of higher incident energy neutrons allows access to much higher energy transfer and higher flux so weaker features at high energies

become visible. The resolution of this setting is insufficient to separate the four modes seen in the 2.5 meV data, but Fig. 5.5(i) shows the dispersion of the most intense modes. The intensity scale in Fig. 5.5(a-h) has been set so that the single magnon scattering is saturated and appears as an intense band through the data. This has been done in order to enhance a weak flat band of scattering at higher energies. This band is featureless, and its overall energy appears to increase at twice the rate of the energy of the single magnon scattering.




Figure 5.5: (i) 6.3 meV neutron scattering showing the single single magnon dispersion at 9 T, originally shown in Fig. 5.2(n). The higher incident energy neutrons allow access to a larger energy transfer range, making the entire dispersion visible. The lower resolution of the data means that only the most intense modes are visible, appearing as a single dispersion. The intensity threshold of panels (a-h) is lowered to enhance weak continuum scattering at high energies. The intensity of this scattering increases while its energy decreases in decreasing magnetic field. The slices were created by integrated over $-0.5 \leq l \leq 0.5$.

Fig. 5.6(a) shows an energy scan through the data in Fig. 5.5(b-h), generated by integrating over the range $-4.0 \leq h \leq 0.0$. These scans show the evolution of this continuum of scattering as a function of applied magnetic field. Fig. 5.6(b) shows a fit to the continuum in order to determine the upper and lower energy bounds of the

continuum. The basic fit function is shown by the green line, and is a quadrilateral with the lower bound defined by a Heaviside function, the top and upper edge are sloping lines. This shape is an empirical parameterisation. This profile was convolved with an energy profile estimated from the line shape of the sharp single magnon scattering and fitted to the data. The resulting shape is shown as a solid blue line in Fig. 5.6(b). The fitted boundaries of the continuum, 2.42(3) meV and 3.74(4) meV, are shown by the magenta arrows. This fitting was performed on energy scans through the data for applied magnetic fields in the range $2\text{ T} \leq B \leq 9\text{ T}$, though for fields less than 4 T, only the upper bound could be fit, as due to the coarse energy resolution of the 6.3 meV data set, there is no clear boundary between the continuum scattering and the single magnon scattering below 4 T. It can be seen in the 2.5 meV data (Fig. 5.3(h)) that the continuum has not actually overlapped with the single magnon scattering at 3 T, as it is a distinct weak feature near $E = 1.5$ meV. Below 2 T, the upper bound of the continuum scattering was extracted from 2.5 meV incident energy neutron data by fitting to a single linear slope convolved with the instrumental energy resolution extracted from the line-shape of single magnon modes. The results of the fits to the continuum boundaries from 6.3 meV data are shown by the dashed lines in Fig. 5.6(a), and to all of data sets in Fig. 5.6(c). It should be noted that the energies in Fig. 5.6(c) are plotted as a function of the internal field of the sample, which is different from the applied field and takes into account demagnetisation effects (see Appendix D). Figure 5.6(d) shows an energy integrated slice through 5 T 6.3 meV Horace data showing that the continuum intensity is modulated in the $(h, k, 0)$ plane, with a rod of high intensity along the $[\bar{1}10]$ direction, and a region of low intensity around the FCC Brillouin zone centre at $(-2, 0, 0)$. This slice was created by integrating the data in the ranges $-0.2 \leq l \leq 0.2$ and $2.1 \leq E \leq 3.0$. This energy range encompasses the highest intensity portion of the high field continuum scattering.

Fig. 5.6(c) also shows the field dependence of the one magnon dispersion, determined by Gaussian fits to energy scans through the single magnon scattering in the




Figure 5.6: (a) Energy scans generated by integrating the data in Fig. 5.5 over the range $-4.0 \leq h \leq 0.0$. These scans have been offset to show each field clearly, the magnetic inelastic signal is indicated by the shaded area under the points. (b) 6 T scan from (a) along with the fit function used to determine the boundaries of the continuum, both convolved with the energy resolution (blue) and in its un-convolved form (green). The boundaries of the continuum are indicated by vertical magenta arrows. (c) Summary of the field dependence of the boundaries of the continuum, along with the results of fitting the bandwidth of the single magnon scattering. The dashed lines show linear fits to the lower bounds of the single magnon and continuum scattering. (d) Energy integrated slice through 5 T, 6.3 meV Horace data showing the modulation of the two magnon continuum intensity in the $(h, k, 0)$ plane, with a rod of high intensity along $[\bar{1}10]$ and a minimum in the intensity at $(-2, 0, 0)$, the antiferromagnetic Brillouin zone centre. The slice is generated by integrating over $-0.2 \leq l \leq 0.2$ and $2.1 \leq E \leq 3.0$ meV, which encompasses the two magnon continuum at this field. The black lines show the FCC Brillouin zone boundaries in the $(h, k, 0)$ plane. All data has had a non-magnetic background subtracted and has been absorption corrected.

2.5 meV and 4.0 meV inelastic neutron data. By fitting the energy of the maximum and minimum energy modes, the bandwidth of the single magnon scattering can be determined. From Fig. 5.2(g-n) it is apparent the the minimum in the one magnon dispersion is reached near $h \rightarrow 0$ (the theoretical model to be discussed in Ch. 6 also confirms this), however the region very close to $h = 0$ is not accessible in the data at high energies, so the dispersion in this region needs to be extrapolated from the region nearby. Two methods were used to estimate the minimum energy of the mode based on the available data. The first method, used for fields in the range $1 \text{ T} \leq B \leq 6 \text{ T}$ was to assume that the low energy dispersion of the mode would be quadratic in $|\mathbf{Q}|$ near the minimum, so the modes energy was fit to a function linear in $|\mathbf{Q}|^2$ to determine the energy of the mode at $|\mathbf{Q}| = 0$. At higher fields ($6 \text{ T} \leq B \leq 9 \text{ T}$), where the data available was not in a $\mathbf{Q}$ range where the dispersion could be considered to be quadratic, the available dispersion points were fitted to a linear spin wave model (Ch. 6) and the minimum gap energy was then taken to be the value predicted by this model at $\mathbf{Q} = 0$. For applied fields below 1 T, the dispersion of the lowest sharp mode is very small, so the energy of this mode at the closest point to $\mathbf{Q} = 0$ is used as an estimate for the the minimum energy gap.

As can be seen in Fig. 5.2, as the applied magnetic field is increased, the energy of both the single magnon dispersion and high energy continuum scattering increase, though at different rates. By comparing the energy dependence of the low energy boundary of the single magnon and continuum scattering it is possible to gain insight into the nature of the higher energy continuum scattering. The dashed lines in Fig. 5.6(c) show the result of fitting straight lines to the lower bounds of the single magnon scattering and the continuum as a function of the internal field of the sample (see Appendix D). The extent of the dashed lines shows the points that were included in the fit. For the continuum scattering, only points for $B \geq 5 \text{ T}$ were included in the fit, because below this field, an upturn in the upper boundary of the continuum can be seen. This upturn in the upper bound indicates that the lower bound may deviate from linear behaviour below this field as well. The exact cause this upturn in the

upper bound of the continuum scattering is unknown, but it may be an indication of an avoided crossing of the continuum scattering and single magnon scattering.

A linear fit to the rate of change of the lower bound of the single magnon modes as a function of internal field yields 0.1757(6) meV/T, and a linear fit to the rate of change of the minimum energy boundary of the continuum scattering as a function of internal field yields 0.35(1) meV/T. The ratio between these two rates of change is 2.02(8). It is also possible to compute the bandwidth of the single magnon scattering and the continuum scattering for comparison. By computing the bandwidth of the single magnon and continuum scattering based on the fitted points at 5, 6, and 9 T (the only field for which upper and lower bounds exist for both forms of scattering), the bandwidth ratios are found to be 2.0(1), 2.09(8), and 2.13(8), respectively, with an average of 2.08(10). The fact that both the ratio of the bandwidths of the two sets of scattering, and the ratio of their change in minimum energy as a function of internal field as two to within uncertainties provides strong evidence that the continuum of scattering we observed is due to two magnon scattering.

Two magnon scattering is inelastic neutron scattering where the incident neutron excites two magnon excitations at the same time, sharing the energy and momentum transfer, giving rise to a continuum of allowed final states as illustrated for the one dimensional case in Fig. 5.7 The one magnon dispersion relation is shown in red. The arrows show two selected (Q, E) points on the magnon dispersion, (k_1, E_1) and (k_2, E_2) . Combining these two points, leads to a third point (blue arrow), which makes up one point in the continuum. By shifting k_1 to the left while moving k_2 to the right by the same amount, it is possible to vary the energy of the blue point without changing its Q value, thus giving rise to a continuum of scattering as a function of energy and fixed Q . The green points show the results of a random sampling of pairs of points on the one magnon dispersion to construct the two magnon continuum. In linear spin wave theory, two-magnon scattering arises from terms in $S^{zz}(\mathbf{Q}, \omega)$, meaning that the two magnon continuum scattering arises from longitudinal fluctuations in the magnetic moments. This scattering can only occur if longitudinal, zero-point quan-

tum fluctuations are present in the ground state of the system, i.e. the semi-classical ground state is not an exact eigenstate of the Hamiltonian. Such terms can arise when the ground state of the system contains non-collinear quantisation directions, or when the exchange Hamiltonian contains off-diagonal exchange, both of which are known to be present in $\text{Yb}_2\text{Ti}_2\text{O}_7$ [27, 33].




Figure 5.7: Illustration of how the energy and wave vector of two magnon scattering is related to single magnon dispersion. The single magnon dispersion (red line) is given by $E(Q) = \Delta + D(1 + \cos(2\pi Q))$. E_1 and E_2 are the energies of one magnon excitations at the points k_1 and k_2 (black arrows). The combination of these two (Q, E) points will lead to a point in the continuum (blue arrow). The green points are a numerical sampling of the possible two magnon events.

Now to address the broadening of the highest energy single magnon mode for applied magnetic fields of $B \leq 2$ T. The physical picture suggested by the data is that when the one magnon dispersion overlaps with the two magnon continuum, the one magnon line shape broadens as the one magnon dispersion overlaps with the two magnon continuum, then the line shape broadens as the one magnon excitations decay into two magnons. This can happen generically if the Hamiltonian allows matrix elements for decay. As will be discussed in Ch. 6, the sharp magnon modes in the data can be well described by linear spin wave theory. In this framework the decay of one magnon into two magnons is allowed by the presence of terms of the $\alpha_{q'}^{\dagger, m_1} \alpha_{q-q'}^{\dagger, m_2} \alpha_q^{m_3}$, where one magnon at $\mathbf{q}$ decays into two magnons at $\mathbf{q}'$ and $\mathbf{q} - \mathbf{q}'$. $\alpha^\dagger$ and α are

creation and annihilation operators for magnon modes $m_n = 1, 2, 3, 4$. Such terms can arise through terms in the exchange Hamiltonian of the form $S_i^z S_j^\pm$, which lead to three operator terms after the diagonalising Bogolubov transformation is applied to the Hamiltonian. $S_i^z S_j^\pm$ can occur either due to anisotropic exchange, or through non-collinear quantisation directions, which would allow isotropic Heisenberg exchange to support such terms. Both sources of such decay terms are present in $\text{Yb}_2\text{Ti}_2\text{O}_7$: the exchange Hamiltonian of $\text{Yb}_2\text{Ti}_2\text{O}_7$ contains explicit terms of the form $S_i^z S_j^\pm$, where the z and $\pm$ components of the spin S_i are defined locally, and while in large applied magnetic fields, the magnetic moments of $\text{Yb}_2\text{Ti}_2\text{O}_7$ are expected to be largely collinear, in decreasing field, they will deviate away from this leading to the splayed ferromagnetic phase of Chang et al. [33], which has different quantisation directions for each tetrahedral sublattice site. In order for these terms to lead to the decay of one magnon into two magnons, it must be possible for the decay process to conserve energy. This means that $E_{m_3, \mathbf{q}} = E_{m_1, \mathbf{q}'} + E_{m_2, \mathbf{q}-\mathbf{q}'}$ must be satisfied, which can only happen if the lower bound of the two magnon continuum is lower than the energy of the mode m_3 at wave vector $\mathbf{q}$. As can be seen in Fig. 5.6(c), this energy condition is satisfied for $B \leq 2$ T, which, along with the existence of non-collinear quantisation directions, and $S_i^z S_j^\pm$ terms in the exchange Hamiltonian, make a strong case for multi-magnon decay being the source of the observed broadening in this data.

All of these various features of the low and high field data combine to give a picture of changing excitations as a function of applied magnetic fields. At high fields, the excitations are sharp magnon modes, and our data also reveals a two magnon continuum at higher energies, originating from zero point quantum fluctuations in the size of the ordered moment. As the field decreases, the two magnon continuum increases in intensity and decreases in energy. The the one and two magnon scattering overlap, decay occurs, and the low energy magnons have their dispersions suppressed, and eventually disappear all together at 0 T, where the excitations form a broad continuum.

5.2.2 Dynamical Correlations in the Full Brillouin Zone at Representative Magnetic Fields

As previously discussed, in addition to single angle data, Horace scans were collected at 0, 0.21, 1.5 and 5 T. Figure 5.8 shows three high symmetry slices through the 2.5 meV data sets collected at each of these fields. These slices are generated by integrating two out of the four $(\mathbf{Q}, E)$ dimensions (three reciprocal space and one energy) of the Horace data. In the case of Fig. 5.8, the data is integrated over two directions in reciprocal space, l , and the direction perpendicular to the displayed data. This integration generates a slice showing the inelastic neutron scattering as a function of energy along a specific direction in reciprocal space. Figure 5.10(a) shows the directions of these slices in the $(h, k, 0)$ plane. In the high field 5 T data, it can be seen that sharp magnon excitations are present and well defined in all directions in the $(h, k, 0)$ plane. The measured dispersion relations will be used in the next chapter to extract an exchange Hamiltonian for $\text{Yb}_2\text{Ti}_2\text{O}_7$. The 1.5 T data (Fig. 5.8(d-f)) shows non-dispersing excitations below a broad continuum along the $[100]$ direction, but along the diagonal $[\bar{1}10]$ direction there is a sharp mode, which disperses strongly and remains sharp up to much higher energies ($E \sim 0.9$ meV).

The slices through the 0.21 T data, shown in Fig. 5.8(g-i) show that at this field, the magnetic excitations are gapped all over the $(h, k, 0)$ plane, with the sharp low energy mode being non-dispersing along all of the directions shown, and that there are no sharp modes in the continuum at higher energies. The low energy continuum shows some modulation in the $\mathbf{Q}$ dependence of its intensity. In both the 0.21 T and the 0 T data (Fig. 5.8(j-l)), the highest intensity of the continuum is along the diagonal $[\bar{1}10]$ direction. This is shown in Figs. 5.8(h,i,k,l). Figures 5.8(h,k) show that the continuum along the $[\bar{1}10]$ direction is much stronger than along the $[100]$ direction. Figures 5.8(i,l) show scans across this high intensity rod, showing the drop in continuum intensity away from the $(-\xi, \xi, 0)$ direction. It is notable that this direction of highest continuum intensity coincides with the direction along which the

largest magnon dispersion is observed in the 1.5 T data.

Figure 5.9 shows slices through 1.34 meV Horace data along the $[\bar{1}10]$ direction for applied fields of 0, 0.21, and 1.5 T. This data demonstrates that the intensity near $\mathbf{Q} = 0$ in the zero field slices of Fig. 5.8(j,k) is not a sharp feature that may indicate the existence of a sharp excitation mode, but simply a region of high intensity, much broader than the sharp modes in the 0.21 T and 1.5 T data.

In addition to slices generated by integrating Horace data over two directions in reciprocal space, it is also possible to generate slices at constant energy showing the $(h, k, 0)$ dependence of the scattered neutron intensity in a plane in reciprocal space. Figure 5.10 shows three such slices, showing the dependence of the intensity of the zero field continuum. Figure 5.10(a) shows a slice generated by integrating over $-0.2 \leq l \leq 0.2$ and low energies ($0.06 \leq E \leq 0.4$ meV), while Fig. 5.10(b) shows the same slice in l but at intermediate energies ($0.4 \leq E \leq 0.8$ meV), and Fig. 5.10(c) is at high energies ($0.8 \leq E \leq 1.5$ meV). As previously mentioned, the dashed arrows in Fig. 5.10(a) show the directions of the slices in Fig. 5.8. Figure 5.10(a) clearly shows that at low energies, the zero field continuum contains a rod of high intensity along the $[-1, 1, 0]$ direction. This modulation is similar to the modulation seen in the high energy two magnon continuum scattering in Fig. 5.6(d). This similarity in the intensity modulation is notable as there is no indication that the excitations in zero and high field are related.

This intensity modulation is reminiscent of the excitation spectrum observed in Herbertsmithite [88]. In that material, a spin-1/2 Kagome lattice antiferromagnet, the zero field excitation spectrum is a gapless continuum with significant $\mathbf{Q}$ modulation. Theoretical work on this material [89] has shown that the dynamical structure factor of a Z_2 spin liquid can explain some aspects of the scattering quite well. This model explains the modulation of scattering intensity in the Kagome plane quite well, though it does require the existence of a finite gap in the excitation spectrum. Punk et al. [89] suggest that the fact that this gap is not observed in experiment may be due to impurity effects, which are not expected to be significant in the case of the




Figure 5.8: Representative slices through 2.5 meV Horace scans for applied fields of 0, 0.21, 1.5 and 5 T. These slices are generated by integrating over the range over the ranges $-0.2 \leq l \leq 0.2$ for all of the slices and $-0.2 \leq k \leq 0.2$ for the $(h, 0, 0)$ slices, and an equivalent range perpendicular to the slice direction for the others. All of the data in this figure has been background subtracted and absorption corrected. Figure 5.10(a) shows the positions of the three slices in reciprocal space.




Figure 5.9: Slices through 1.34 meV Horace data for applied fields of (a) 0 T, (b) 0.21 T, and (c) 1.5 T showing that the feature near $\mathbf{Q} = 0$ in the 0 T data is not sharp compared to the sharp mode at the bottom edge of the continuum in the 0.21 T data, and the dispersing mode in the 1.5 T data. The weak dispersing feature below 0.2 meV at high $|\mathbf{Q}|$ is a remnant of the scattering from the superconducting magnet tails. The integration ranges used to generate these slices are $-0.2 \leq l \leq 0.2$ and ± 0.2 in the direction perpendicular to the slice direction in the $(h, k, 0)$ plane.

measurements in this work, but it may be possible to construct such a model where the gap is < 0.04 meV, and is simply impossible to resolve. The existence of such a spin liquid state is not consistent with the ordered state in zero field proposed by Chang et al. [33] based on polarised elastic neutron scattering measurements, and would require the existence of a phase transition as a function of applied field along [001]. The heat capacity measurements in Ch. 3 show evidence for the suppression of the phase transition that occurs in zero applied field when a small magnetic field (0.036 T) is applied along the [001] axis, but there is no direct evidence of a phase transition as a function of magnetic field along [001].

In addition to generating two dimensional slices through the Horace data sets by integrating over two of the available four $(\mathbf{Q}, E)$ dimensions, it is possible to generate one dimensional scans by integrating over three of the four free dimensions. Figure 5.11 shows scans in energy generated by integrating over the three reciprocal space directions (h, k, l) . The scans shown in Fig. 5.11(a) show scans through 0 and 0.21 T 2.5 meV Horace scans generated by integrating over all of the surveyed (h, k, l) volume. These scans show that the behaviour of the low field scattering as a function of energy is not significantly different in regions of reciprocal space away from the $(h, 0, 0)$ direction shown in the scans in Figs. 5.3 and 5.4. This scan has much smaller




Figure 5.10: Panels (a), (b), and (c) show energy integrated slices through reciprocal space of the 0 T, 2.5 meV data. These slices were generated by integrating over the range $-0.2 \leq l \leq 0.2$, and the energy ranges (a) $0.06 \leq E \leq 0.4$ meV, (b) $0.4 \leq E \leq 0.8$ meV, and (c) $0.8 \leq E \leq 1.5$ meV. The arrows in (a) indicate the locations of the slices through the Horace data shown in Fig. 5.8. The solid black lines mark FCC Brillouin zone boundaries.

error bars than the scans in Fig. 5.4(a) and shows that the 0 T data shows no evidence of a gap in the excitation spectrum down to 0.0875 meV, while the 0.21 T data shows evidence of a sharp mode at low energies with a finite energy gap of 0.229(3) meV. Figure 5.11(b) shows similar energy scans through the 1.34 meV Horace data sets in applied fields of 0, 0.21, and 5 T data sets generated using the reciprocal space integration range of $-0.4 \leq h \leq 0.4$, $-0.4 \leq k \leq 0.4$, and $-0.4 \leq l \leq 0.4$. The 5 T data shows that there is no non-magnetic scattering, as for this field the magnetic excitations are outside of the energy range shown in Fig. 5.11(b). The 1.34 meV data has much better energy resolution (The FWHM of the elastic line is 0.0231(1) meV for the 1.34 meV data, as opposed to 0.0555(3) for the 2.5 meV data), allowing better determination of the existence or non-existence of a gap in the excitation spectrum at low energies. Figure 5.11(b) clearly shows that down to the lowest energy plotted (0.037 meV), the zero field data shows no evidence of a gap in the excitation spectrum, whereas the 0.21 T data shows that the magnetic scattering intensity goes to zero at the lowest accessible energy.




Figure 5.11: (a) Energy scan through the 0 and 0.21 T 2.5 meV Horace scans generated by integrating over all of the $(h, k, 0)$ plane, and the range $-0.5 \leq l \leq 0.5$. (b) shows a scan in energy through the 0, 0.21, and 5 T 1.34 meV Horace data sets generated by integrating over the ranges $-0.4 \leq h \leq 0.4$, $-0.4 \leq k \leq 0.4$, and $-0.4 \leq l \leq 0.4$.

5.3 Conclusions

In this chapter, inelastic neutron scattering measurements on a single crystal sample of $\text{Yb}_2\text{Ti}_2\text{O}_7$ at very low temperatures ($T \sim 145$ mK, well within the low temperature phase of $\text{Yb}_2\text{Ti}_2\text{O}_7$ in a wide range of applied magnetic fields ($0\text{T} \leq B \leq 9$ T $\parallel [001]$) were presented and discussed. In large applied magnetic fields ($B \geq 3$ T) the magnetic excitation spectrum is found to take the form of four sharp magnon modes, along with a weak band at twice the energy of the magnon modes. Based on the ratio of the bandwidth of the magnon modes and this band of scattering, and the ratio of the rate of change of the lower energy bound of the magnon scattering and the high energy continuum band as a function of the magnetic field, the continuum of scattering was identified as being due to two magnon scattering processes. When the applied magnetic field was lowered below 3 T, the highest energy single magnon dispersion overlaps with the low energy boundary of the two magnon continuum, and the sharp mode then broadens and becomes part of the continuum. The remaining lower-energy three sharp magnon modes are observed to have their dispersions suppressed as the applied magnetic field is decreased, and appear to collapse into a single, almost dispersion-less mode for fields of $B \lesssim 0.75$ T. Finally, in zero field, no sharp modes are observed in the magnetic scattering and at all probed energies the magnetic scatter-

ing takes the form of a very broad continuum (width ~ 1.5 meV) with no observable gap even at the highest resolutions used in the experiments. Applying a very small magnetic field, equivalent to an internal field of $\mu_0 H_{\text{int}} = 0.17$ T is sufficient to open a finite gap to the lowest energy excitation. The absence of sharp modes in the zero field is indicative of the strong zero-point quantum fluctuations in the ground state which prevent coherent propagation of single-spin-flip magnons down to the lowest probed energies (0.03 meV). Instead single spin flip excitations created in a neutron scattering process decay into other types of excitations. The very broad extent of the continuum, almost an order of magnitude larger than the scale of the largest exchange interaction (see Ross et al. [27]), suggests that the corresponding excitations involve multiple spin flips, which could have a non-local character and occur in a correlated multi-spin clusters, or be due to coherently propagating quasi-particles, like the monopole excitations in quantum spin ice states [31, 32, 54].

Chapter 6

Effective Spin Half Linear Spin Wave Description of the One and Two Magnon Excitations in $\text{Yb}_2\text{Ti}_2\text{O}_7$ in Applied Magnetic Field

This chapter describes fits to high field magnon dispersions from Ch. 5, performed to determine a nearest-neighbour effective spin-1/2 Hamiltonian for $\text{Yb}_2\text{Ti}_2\text{O}_7$. The fits are performed against four different data sets: magnon dispersions along many directions in the $(h, k, 0)$ plane extracted from 5 T $\parallel$ [001] Horace data, magnon dispersions extracted from 9 T $\parallel$ [001] single angle data, the high field magnetisation of a single crystal sample of $\text{Yb}_2\text{Ti}_2\text{O}_7$ in a [001] field, and dispersion points extracted from the inelastic neutron scattering in a $[1\bar{1}0]$ field of Ross et al. [27]. In order to perform this fit, a nearest-neighbour exchange Hamiltonian is defined, and linear spin wave (LSW) theory on the pyrochlore lattice is outlined, along with how this theory is used to compute neutron scattering cross sections. The process of extracting one magnon dispersion relations from experimental data is described, and the details of the least

squares minimisation process to determine the optimal Hamiltonian for $\text{Yb}_2\text{Ti}_2\text{O}_7$ are discussed. The resulting exchange parameters are used to compute the one and two magnon spectra for comparison to the experimental data, finding good agreement with the single magnon dispersion, and the energy boundaries of the two magnon scattering at high fields ($B \geq 3$ T), but underestimating the weight of the two magnon continuum. At low fields ($B < 3$ T) LSW theory is found to break down when the one and two magnon scattering overlap, likely due to magnon interactions and magnon decay processes. In addition to comparing the neutron scattering spectra to LSW theory, the neutron excitation gap is compared to the gap parameter extracted from the in-field heat capacity measurements of Ch. 3, showing that the gap extracted from the heat capacity at very low fields is much lower than the gap observed in neutron scattering measurements. The exchange Hamiltonian is also used to compute a mean field phase diagram in zero field, showing the optimal exchange parameters lie very close to a phase boundary in exchange space, which may provide an explanation of the exotic excitation spectrum observed in zero field. Mean field calculations of the ground state in applied field show that upon increasing the field, the effective spins become more and more polarised along the field.

6.1 Linear Spin Wave Theory

The fits to high field magnon dispersions in this chapter are based on comparing linear spin wave (LSW) theory to experimental data. The linear spin wave theory (LSWT) for the pyrochlore $\text{Yb}_2\text{Ti}_2\text{O}_7$ has been developed by Dr. Paul McClarty, who also obtained the final exchange parameters based on fits to the experimental dispersion points extracted from the neutron scattering data. He was also responsible for calculations of the mean-field magnetisation, and the generation of the mean-field phase diagram. The general outline of the principles behind the LSW description of magnetic excitations as relevant to $\text{Yb}_2\text{Ti}_2\text{O}_7$ is discussed below for completeness, and follows the method outlined in Refs. [27, 90].

To start, it is useful to write the general magnetic Hamiltonian in the form:

$$\mathcal{H} = \mathcal{H}_{\text{Exchange}} + \mathcal{H}_{\text{Zeeman}}, \quad (6.1)$$

where the first term describes (the nearest-neighbour) exchange interactions between the Yb^{3+} magnetic moments in $\text{Yb}_2\text{Ti}_2\text{O}_7$, and the second term describes the interaction of the Yb^{3+} moments with an external applied magnetic field.

As discussed in Ch. 1.8.4, the crystal field configuration of the Yb^{3+} ions allows for the treatment of the magnetic moments using an effective spin-1/2 approximation where the excited crystal field doublets are neglected and the two states in the low energy doublet are treated as an effective spin-1/2. In this approximation the exchange portion of the Hamiltonian is obtained to leading order in the exchange by projecting the exchange between the full $\mathbf{J}$ moments on to the low energy crystal field doublet. This means that the Hamiltonian used here inhabits a subspace of all of the possible exchange interactions between the full $\mathbf{J}$ moments. The exchange portion of the Hamiltonian is also restricted to nearest-neighbour interactions between the effective spin-1/2 moments. The long range dipolar interaction between the magnetic moments is neglected because the strength of this interaction in $\text{Yb}_2\text{Ti}_2\text{O}_7$, calculated using the Ewald summation method, is 0.0016 meV for nearest-neighbour spins [91], which is an order of magnitude smaller than any of the interaction energies determined by Ross et al. [27].

The nearest-neighbour exchange Hamiltonian contains all of the symmetry allowed nearest-neighbour bi-linear exchange terms for the space group $Fd\bar{3}7(O_h^7)$ and point group O_h , which describe the symmetry of the pyrochlore lattice. Using group theory it is possible to show that this symmetry allows only 4 linearly independent nearest-neighbour exchange terms [91]. If the spin-1/2 operators for each sublattice site are expressed as column vectors in terms of global Cartesian coordinates, then the matrix that defines the exchange between spins on sublattice sites 1 and 2 (Table 6.1) is given

by [27]:

$$\mathbf{J}_{1,2} = \begin{pmatrix} J_2 & J_4 & J_4 \\ -J_4 & J_1 & J_3 \\ -J_4 & J_3 & J_1 \end{pmatrix}, \quad (6.2)$$

where $J_{n=1,2,3,4}$ are the energies associated with the four symmetry allowed exchange interactions, and the interaction between the two spins is $\mathbf{S}_1^T \mathbf{J}_{1,2} \mathbf{S}_2$, where $\mathbf{S}_{1,2}$ are effective spins in global Cartesian coordinates. These four interaction terms are linearly independent, and invariant under the symmetry operations of O_h , so any linear combination of these interactions will also be invariant under O_h and can be used to express the nearest-neighbour exchange Hamiltonian. Here we choose to write the exchange Hamiltonian in terms of spin components in the local Ising frame [27]:

$$\begin{aligned} \mathcal{H}_{\text{Exchange}} = \sum_{\langle i,j,a,b \rangle} & \left(J_{zz} S_{ia}^z S_{jb}^z - J_{\pm} (S_{ia}^+ S_{jb}^- + S_{ia}^- S_{jb}^+) \right. \\ & + J_{\pm\pm} (\gamma_{ab} S_{ia}^+ S_{jb}^+ + \gamma_{ab}^* S_{ia}^- S_{jb}^-) \\ & \left. + J_{z\pm} \left[S_{ia}^z (\zeta_{ab} S_{jb}^+ + \zeta_{ab}^* S_{ia}^-) + (S_{ia}^z (\zeta_{ab} S_{jb}^+ + \zeta_{ab}^* S_{ia}^-))^{\dagger} \right] \right), \quad (6.3) \end{aligned}$$

where the sum is over all $\langle i, j, a, b \rangle$ nearest-neighbour pairs of Yb^{3+} ions, where i, j specify FCC lattice sites, and a, b specify tetrahedral sublattice sites. S_{ia} are effective spin-1/2 operators that define the degrees of freedom in the crystal field ground state doublet of the Yb^{3+} ions, as explained in Ch. 1.8.4. In this form the interactions are expressed in terms of $S_{ia}^z, S_{ia}^+, S_{ia}^-$, the local z , raising, and lowering spin operators for the effective spin-1/2 respectively, where the local z directions are defined in Tab. 6.1. The terms γ_{ab} and ζ_{ab} are 4×4 matrices that define phases that depend on which tetrahedral sublattice sites the pairs of effective spin-1/2 operators occupy. The precise form of these phase factors depends in part on the choice of coordinate frame assigning the local z_a directions in the $\langle 111 \rangle$ directions given in Tab. 6.1. In

this case the phase factors are given by [27]:

$$\zeta = \begin{pmatrix} 0 & -1 & e^{i\pi/3} & e^{-i\pi/3} \\ -1 & 0 & e^{-i\pi/3} & e^{i\pi/3} \\ e^{i\pi/3} & e^{-i\pi/3} & 0 & -1 \\ e^{-i\pi/3} & e^{i\pi/3} & -1 & 0 \end{pmatrix}, \quad (6.4)$$

and

$$\gamma = -\zeta^*. \quad (6.5)$$

In this form, the indices a, b of ζ and γ in Eqn. (6.3) index the rows and columns of the matrix. Using this form of the magnetic exchange interactions, the exchange Hamiltonian can be completely specified by the four exchange terms $J_e = \{J_{zz}, J_{\pm}, J_{\pm\pm}, J_{z\pm}\}$.

Table 6.1: The local quantisation directions $\hat{\mathbf{z}}_a$ for the four tetrahedral sublattice sites, along the local $\hat{\mathbf{x}}_a$ directions, taken from Ross et al. [27]. The local $\hat{\mathbf{y}}_a$ directions are defined by $\hat{\mathbf{z}}_a \times \hat{\mathbf{x}}_a$.

Lattice Site	Position	Local $\hat{\mathbf{z}}_a$	Local $\hat{\mathbf{x}}_a$
1	$\frac{a}{8}(1, 1, 1)$	$\frac{1}{\sqrt{3}}(1, 1, 1)$	$\frac{1}{\sqrt{6}}(-2, 1, 1)$
2	$\frac{a}{8}(-1, -1, 1)$	$\frac{1}{\sqrt{3}}(1, -1, -1)$	$\frac{1}{\sqrt{6}}(-2, -1, -1)$
3	$\frac{a}{8}(-1, 1, -1)$	$\frac{1}{\sqrt{3}}(-1, 1, -1)$	$\frac{1}{\sqrt{6}}(2, 1, -1)$
4	$\frac{a}{8}(1, -1, -1)$	$\frac{1}{\sqrt{3}}(-1, -1, 1)$	$\frac{1}{\sqrt{6}}(2, -1, 1)$

For simplicity, the exchange Hamiltonian can be written as:

$$\mathcal{H}_{\text{Exchange}} = \sum_{\langle i,j,a,b \rangle} \mathcal{J}_{ia,jb}^{\alpha\beta} \mathbf{S}_{ia}^{\alpha} \mathbf{S}_{jb}^{\beta}, \quad (6.6)$$

where $\mathbf{S}_{ia}$ are the effective spin operators, and $\mathcal{J}$ is a matrix containing all of the nearest-neighbour exchange interactions. In this form, the sum over all of the magnetic lattice sites is explicit, while the sum over the $\alpha, \beta = z, +, -$ components of the spin-1/2 operators are implicit. Implicit summation over components of matrices and vectors sharing the same index will be used throughout this chapter.

The Zeeman exchange term is defined as:

$$\mathcal{H}_{\text{Zeeman}} = -\mu_0\mu_B H^\mu \sum_{i,a} g_a^{\mu\nu} S_{ia}^\nu. \quad (6.7)$$

H^μ are components of the applied magnetic field in global Cartesian coordinates (the lab frame). $g_a^{\mu\nu}$ are the global Cartesian components of the $\mathbf{g}$ tensor described in Ch. 1.8.4, which characterises the ground state doublet of the CEF states of the Yb^{3+} ion. The form of $\mathbf{g}$ when expressed in terms of global Cartesian coordinates changes depends on the sublattice site a . When expressed in the local $\langle 111 \rangle$ frame $\mathbf{g}$ is diagonal and has two independent terms, $g_{\parallel}$ and $g_{\perp}$ (Eqn. (1.4)), that describe the local Ising and XY anisotropy of the full magnetic moment respectively. The diagonal nature of the $\mathbf{g}$ tensor allows the use of a single index to denote the diagonal elements in the local $\langle 111 \rangle$ frame. These components are included as free parameters in the fit to the spin Hamiltonian, along with the J_e exchange terms.

In the limit of infinite field and an isotropic $\mathbf{g}$ tensor, all effective spins (and magnetic moments) in the system are aligned along the applied field direction. For an anisotropic $\mathbf{g}$ tensor, the effective spins tilt away from perfect parallel alignment to reflect the local $\mathbf{g}$ tensor anisotropy and reduce their Zeeman energy, resulting in an overall magnetic arrangement with slightly canted moments with the overall periodicity of the FCC lattice, i.e. $\mathbf{Q} = 0$ magnetic structure defined by four spins on a tetrahedron. The presence of finite exchange terms affects the local tilts of the magnetic moments but leaves the overall periodicity unchanged. Experiment suggests that is also true in zero field [33]. For each applied field, the zero temperature mean-field ground state of the Hamiltonian (Eqn. (6.1)) was found using an iterative numerical algorithm to determine the minimum energy arrangement of the effective spins with the constraint $\mathbf{Q} = 0$. The Hamiltonian is then recast in terms of the local equilibrium effective spin directions by rotating the spin operators from the local $[111]$ frame to a local quantisation frame via matrices $\bar{R}_a^{\mu\alpha}$, where μ is the quantisation frame index, α is the local $[111]$ frame index, and a specifies the sublattice site.

This matrix is constructed from column vectors of the local quantisation Cartesian directions $\hat{\mathbf{u}}_a, \hat{\mathbf{v}}_a, \hat{\mathbf{w}}_a$ expressed in the local [111] frame, where $\hat{\mathbf{w}}_a$ is a unit vector in the direction of the equilibrium effective spin direction, $\hat{\mathbf{u}}_a = (\hat{\mathbf{w}}_a \times (1, 1, 1)) / |\hat{\mathbf{w}}_a \times (1, 1, 1)|$ and $\hat{\mathbf{v}}_a = \hat{\mathbf{w}}_a \times \hat{\mathbf{u}}_a$. All quantities expressed in the quantisation frame are denoted by a tilde.

Within the local quantisation basis, the spin-1/2 operators for each sublattice site, $\tilde{S}_a^w$, $\tilde{S}_a^+$, and $\tilde{S}_a^-$, can be expressed in terms of Holstein-Primakoff bosons, and then expanded in powers of $1/S$, yielding:

$$\tilde{S}_a^w = S - c_a^\dagger c_a, \quad (6.8)$$

$$\tilde{S}_a^+ = \tilde{S}_a^u + i\tilde{S}_a^v = \sqrt{2S - c_a^\dagger c_a} c_a \approx \sqrt{2S} c_a, \quad (6.9)$$

$$\tilde{S}_a^- = \tilde{S}_a^u - i\tilde{S}_a^v = c_a^\dagger \sqrt{2S - c_a^\dagger c_a} \approx \sqrt{2S} c_a^\dagger. \quad (6.10)$$

Here $c_a^\dagger$ and c_a are creation and annihilation operators for Holstein-Primakoff bosons on the sublattice site a . These operators satisfy the commutation relations:

$$[c_a, c_b^\dagger] = \delta_{ab}. \quad (6.11)$$

All other commutators are zero.

By replacing the effective spin-1/2 operators in $\mathcal{H}$ with their Holstein-Primakoff boson expansion, the Hamiltonian becomes:

$$\mathcal{H} = \mathcal{H}^{(0)} + \mathcal{H}^{(1)} + \mathcal{H}^{(2)} + \dots, \quad (6.12)$$

where the superscript (n) denotes the power of boson operators in that term. LSWT keeps only terms up to $\mathcal{H}^{(2)}$. The ellipses represents terms of cubic and higher order in the boson operators. For the equilibrium mean-field ground state $\mathcal{H}^{(1)}$ is zero.

In this expansion, the first term in the Hamiltonian is given by:

$$\mathcal{H}^{(0)} = S(S+1) \sum_{\langle i,j,a,b \rangle} \mathcal{J}_{ia,jb}^{\alpha\beta} w_a^\alpha w_b^\beta, \quad (6.13)$$

where w_a^α are the $\alpha = x, y, z$ components of $\hat{\mathbf{w}}_a$ in the local $[111]$ frame. The sum is the same as in Eqn. (6.3). This term gives the classical energy of the ground state with $S(S+1)$ replaced by S^2 . Eqn. (6.13) also includes a term linear in S that strictly arises from symmetrising $\mathcal{H}^{(2)}$, but has been included in $\mathcal{H}^{(0)}$ for simplicity.

The first step in dealing with the quadratic term $\mathcal{H}^{(2)}$ is to Fourier transform the Holstein-Primakoff bosons, to obtain the operators

$$c_{a\mathbf{k}}^\dagger = \frac{1}{\sqrt{N}} \sum_{\mathbf{R}_i} c_a^\dagger \exp(-i\mathbf{k} \cdot (\mathbf{R}_i + \mathbf{r}_a)), \quad (6.14)$$

and

$$c_{a\mathbf{k}} = \frac{1}{\sqrt{N}} \sum_{\mathbf{R}_i} c_a \exp(i\mathbf{k} \cdot (\mathbf{R}_i + \mathbf{r}_a)), \quad (6.15)$$

where $\mathbf{k}$ is a vector in the first Brillouin zone, $\mathbf{R}_i$ is an FCC lattice position, and $\mathbf{r}_a$ is a tetrahedral sublattice position. Performing this Fourier transform results in the interaction matrix:

$$\mathcal{J}_{ab\mathbf{k}}^{\alpha\beta} = \frac{1}{N} \sum_{\langle i,j \rangle} \mathcal{J}_{ia,jb}^{\alpha\beta} \exp(i\mathbf{k} \cdot (\mathbf{R}_i - \mathbf{R}_j + \mathbf{r}_a - \mathbf{r}_b)), \quad (6.16)$$

where $\mathbf{R}_i$ is the position of the FCC lattice site i , and $\mathbf{r}_a$ is the relative position of the tetrahedral sublattice site a . This interaction matrix is converted into the local quantisation frame as $\tilde{\mathcal{J}}_{ab\mathbf{k}}^{\mu\nu} = \bar{R}_a^{\mu\alpha} \bar{R}_b^{\nu\beta} \mathcal{J}_{ab\mathbf{k}}^{\alpha\beta}$.

The quadratic term $\mathcal{H}^{(2)}$ in the Hamiltonian will describe the harmonic spin wave excitations of the Hamiltonian. It is helpful to define the boson operators:

$$\tilde{\sigma}_{\mathbf{k}}^u \equiv \sqrt{\frac{S}{2}} \left(\mathbf{c}_{\mathbf{k}}^\dagger + \mathbf{c}_{-\mathbf{k}} \right), \quad (6.17)$$

and

$$\tilde{\sigma}_{\mathbf{k}}^v \equiv i\sqrt{\frac{S}{2}} \left(\mathbf{c}_{\mathbf{k}}^\dagger - \mathbf{c}_{-\mathbf{k}} \right), \quad (6.18)$$

where $\mathbf{c}_{\mathbf{k}}$ stands for the operator column vector $(c_{1\mathbf{k}}, c_{2\mathbf{k}}, c_{3\mathbf{k}}, c_{4\mathbf{k}})^T$. By introducing these vectors, it is possible to rewrite the quadratic Hamiltonian in matrix form as:

$$\begin{aligned} \mathcal{H}^{(2)} &= \sum_{\mathbf{k}} \left((\tilde{\sigma}_{\mathbf{k}}^u)^\dagger (\tilde{\sigma}_{\mathbf{k}}^v)^\dagger \right) \begin{pmatrix} P_{\mathbf{k}}^u & Q_{\mathbf{k}} \\ Q_{\mathbf{k}}^* & P_{\mathbf{k}}^v \end{pmatrix} \begin{pmatrix} (\tilde{\sigma}_{\mathbf{k}}^u) \\ (\tilde{\sigma}_{\mathbf{k}}^v) \end{pmatrix} \\ &= \sum_{\mathbf{k}} \left((\tilde{\sigma}_{\mathbf{k}}^u)^\dagger (\tilde{\sigma}_{\mathbf{k}}^v)^\dagger \right) \mathbf{H} \begin{pmatrix} (\tilde{\sigma}_{\mathbf{k}}^u) \\ (\tilde{\sigma}_{\mathbf{k}}^v) \end{pmatrix}, \end{aligned} \quad (6.19)$$

where the sum is over all $\mathbf{k}$ in the first Brillouin zone.

In this form, the components of the 4×4 matrices $P_{\mathbf{k}}^u$, $P_{\mathbf{k}}^v$, and $Q_{\mathbf{k}}$ are given by:

$$P_{ab\mathbf{k}}^u = \mathcal{R}_a \delta_{ab} + \tilde{\mathcal{J}}_{ab\mathbf{k}}^{uu} + \frac{\mu_0 \mu_B}{2} H^\alpha R_a^{\beta\alpha} g^\beta \bar{R}_a^{z\beta} \delta_{ab} \quad (6.20)$$

$$P_{ab\mathbf{k}}^v = \mathcal{R}_a \delta_{ab} + \tilde{\mathcal{J}}_{ab\mathbf{k}}^{vv} + \frac{\mu_0 \mu_B}{2} H^\alpha R_a^{\beta\alpha} g^\beta \bar{R}_a^{z\beta} \delta_{ab} \quad (6.21)$$

$$Q_{ab\mathbf{k}} = \tilde{\mathcal{J}}_{ab\mathbf{k}}^{uv}, \quad (6.22)$$

where:

$$\mathcal{R}_a = - \sum_c \tilde{\mathcal{J}}_{ac, \mathbf{k}=0}^{ww}, \quad (6.23)$$

g^β are diagonal components of the $\mathbf{g}$ tensor in the local [111] coordinate frame, and H^α are components of the applied magnetic field in the global Cartesian frame. $R_a^{\alpha\beta}$ is a rotation from the global Cartesian reference frame to the local [111] reference frame, where β is the global frame index, α is the local [111] frame index, and a specifies the sublattice site. This matrix has components $R_a^{x\beta} = \hat{\mathbf{x}}_a \cdot \hat{\mathbf{e}}_\beta$, $R_a^{y\beta} = \hat{\mathbf{y}}_a \cdot \hat{\mathbf{e}}_\beta$, and $R_a^{z\beta} = \hat{\mathbf{z}}_a \cdot \hat{\mathbf{e}}_\beta$, where $\hat{\mathbf{e}}_\beta$ are the global Cartesian axes, and $\hat{\mathbf{x}}_a, \hat{\mathbf{y}}_a, \hat{\mathbf{z}}_a$ are the local coordinates, which are defined in Tab. 6.1.

In order to determine the excitation spectrum of the Hamiltonian, it is necessary to diagonalise the matrix $\mathbf{H}$. The diagonalisation process must respect the boson

commutation relations:

$$[\tilde{\sigma}_{\mathbf{k}}^{\alpha}, (\tilde{\sigma}_{\mathbf{k}}^{\beta})^{\dagger}] = S\boldsymbol{\eta}^{\alpha\beta}, \quad (6.24)$$

where

$$\boldsymbol{\eta} \equiv \begin{pmatrix} 0 & -i\mathbf{I} \\ -i\mathbf{I} & 0 \end{pmatrix}, \quad (6.25)$$

and $\mathbf{I}$ is a 4×4 identity matrix.

Using this formulation of the commutation relations, the magnon energies for the mode m , $\omega_{m\mathbf{k}}$, are the positive semi-definite eigenvalues of the matrix $2\boldsymbol{\eta}\mathbf{H}$. The right eigenvectors $\mathbf{v}_{m\mathbf{k}}$ of the matrix $\boldsymbol{\eta}\mathbf{H}$ are the magnon wave-functions in the $a, \mathbf{k}$ basis and preserve the commutation relation among the $\boldsymbol{\sigma}$ operators provided $\mathbf{v}^{\dagger}\boldsymbol{\eta}\mathbf{v} = \mathbf{d}$, where $\mathbf{d} = \text{diag}(\mathbf{1}, -\mathbf{1})$. By periodicity, the spectrum at a general wave vector $\mathbf{Q}$ is the same as at the reduced wave vector $\mathbf{K} = \mathbf{Q} - \mathbf{G}$, where $\mathbf{G}$ is the nearest Brillouin zone centre of the FCC lattice.

As explained in Chapter 4, the inelastic neutron scattering intensity is given by:

$$I(\mathbf{Q}, \omega) \propto \left| \frac{1}{2}F(Q) \right|^2 \sum_{\alpha, \beta} \left(\delta_{\alpha\beta} - \hat{Q}_{\alpha}\hat{Q}_{\beta} \right) \mathcal{S}^{\alpha\beta}(\mathbf{Q}, \omega), \quad (6.26)$$

where $\mathbf{Q} = \mathbf{k}_i - \mathbf{k}_f$ is the wave vector transfer defined in Chapter 4. $\mathbf{k}_i$ and $\mathbf{k}_f$ are the incident and scattered wave vectors of the neutron. $F(Q)$ is the magnetic form factor of the Yb^{3+} ion, $\hat{\mathbf{Q}}$ is a unit vector in the direction of $\mathbf{Q}$, and the indices α, β are Cartesian components in the lab frame.

The inelastic component of the dynamical correlation function $\mathcal{S}^{\alpha\beta}(\mathbf{Q}, \omega)$ is given by:

$$\mathcal{S}_{\text{inelastic}}^{\alpha\beta}(\mathbf{Q}, \omega) = \sum_{a,b} \sum_m \langle 0 | J_a^{\alpha}(-\mathbf{Q}) | m \rangle \langle m | J_b^{\beta}(\mathbf{Q}) | 0 \rangle \delta(\omega - \omega_{m\mathbf{Q}}), \quad (6.27)$$

where α and β are components in the global Cartesian reference frame, and m indexes the magnon modes as before. To connect this to the dynamical correlation function

based on the effective spin-1/2 model, we relate the full magnetic moment $\mathbf{J}_a$ to the effective spin-1/2 object by direct projection via the $\mathbf{g}$ tensor, so that the global Cartesian components of $\mathbf{J}_a$ are given by:

$$\mathbf{J}_a^\alpha = R_a^{\beta\alpha} g^\beta \bar{R}_a^{\gamma\beta} \tilde{\mathbf{S}}_a^\gamma. \quad (6.28)$$

Using this relationship, the one-magnon inelastic scattering cross-section becomes:

$$\mathcal{S}_{\text{inelastic, 1-magnon}}^{\alpha\beta}(\mathbf{Q}, \omega) = \frac{S}{2} \sum_{a,b} \sum_{\mu, \nu=u,v} \Gamma_a^{\alpha\mu} \Gamma_b^{\beta\nu} \sum_m \mathbf{v}_{am, -\mathbf{K}}^{\mu\dagger} \mathbf{v}_{bm, \mathbf{K}}^\nu \exp(-i\mathbf{G} \cdot \mathbf{r}_{ab}), \quad (6.29)$$

where $\Gamma_a^{\alpha\gamma} = R_a^{\beta\alpha} g^\beta \bar{R}_a^{\gamma\beta}$, $\mathbf{v}_{am\mathbf{K}}^\mu$ are the transverse $\mu = u, v$ components of the eigenvectors $\mathbf{v}_{m\mathbf{K}}$ for sublattice site a in the local quantisation frame, and $\mathbf{r}_{ab} = \mathbf{r}_b - \mathbf{r}_a$. The scattering wave vectors $\mathbf{Q} = \mathbf{K} + \mathbf{G}$, where $\mathbf{K}$ is a wave vector inside the first Brillouin zone, and $\mathbf{G}$ is an FCC reciprocal lattice vector.

6.2 One Magnon Dispersion Points

Dispersion data points in $(\mathbf{Q}, E)$ space were extracted from the experimental data for comparison to LSW calculations, in order to find optimal values for the exchange energies and $\mathbf{g}$ tensor components. The data points have been extracted from two different data sets. The first data set is the 2.5 meV Horace scan collected at 5 T. Figure 6.1(a) shows a typical scan in energy through the Horace data, used to determine the energy of a single intense mode at $\mathbf{Q} = (-1.5, 0, 0)$. The energy of the fitted mode is 1.305(3) meV where the error includes the uncertainty in the position of the elastic line zero energy. Figure 6.1(b) shows the results of fits to multiple energy scans at fixed $\mathbf{Q}$ along the $(h, 0, 0)$ direction in reciprocal space. Figure 6.1(c) shows a map of all of the directions in the $(h, k, 0)$ plane in reciprocal space along which $(\mathbf{Q}, E)$ points were extracted from the 5 T data. A total of 354 dispersion points were extracted from the Horace data set for use in the fit.




Figure 6.1: (a) Energy scan through the 2.5 meV, 5 T Horace data set at $\mathbf{Q} = (-1.5, 0, 0)$, generated by integrating over the ranges $-1.55 \leq h \leq -1.45$, $-0.1 \leq k \leq 0.1$, $-0.1 \leq l \leq 0.1$. The solid blue line shows a Gaussian fit to the single visible mode in the data, centred at 1.305(3) meV for the one intense mode visible in the data. (b) Slice along the $(h, 0, 0)$ direction in reciprocal space showing the results of picking the energy position of the most intense modes in the data for multiple points along $(h, 0, 0)$ as white dots. The point in (a) is marked by a magenta triangle. (c) Schematic diagram of the $(h, k, 0)$ reciprocal plane showing directions (dashed red lines) along which $(\mathbf{Q}, E)$ points have been extracted from the volume (Horace) data set. Solid black lines show Brillouin zone boundaries of the FCC structural cell.

In addition to the 5 T Horace data, dispersion points were extracted from the 9 T single angle, 4.0 meV data set. Figure 6.2(a) shows a scan in energy through this data, along with a fit to the intense magnon mode in the data, near $h = -1.5$, giving a peak centred at 2.110(3) meV. The exact point in reciprocal space of this point is actually $(h, k, 0) = (-1.5, -0.31, 0)$, because the detector positions in reciprocal space change as a function of energy transfer, as discussed in Ch. 4. Figure 6.2(b) shows all of the single angle data, along with the results of the Gaussian fits to the magnon mode energies. This data contributed 35 points to the least squares fit data set.




Figure 6.2: (a) Energy scan through the 4.0 meV, 9 T single angle data set at $h = -1.5$, generated by integrating over the ranges $-1.55 \leq h \leq -1.45$, $-0.2 \leq l \leq 0.2$. The solid blue line shows a Gaussian fit to the single visible mode in the data, resulting in an energy of 2.110(3) meV for the one intense mode visible in the data. (b) Slice along the $(h,0,0)$ direction in reciprocal space showing the location of picked dispersion points (white circles) and the data point in (a) (magenta triangle). This slice was generated by integrating over $-0.5 \leq l \leq 0.5$.

In order to use all of these dispersion points, each one has to be uniquely assigned to one of the four magnon modes expected to be present at each $\mathbf{Q}$. At each $\mathbf{Q}$ the modes were labelled in order of increasing energy. At reciprocal space positions where all four modes had finite intensity this process was straight forward, but for positions where not all of the modes have sufficient intensity to be detected in the experiment, mode numbers were assigned using continuity arguments based on comparison to neighbouring points in reciprocal space where the modes have finite intensity. Comparison to the theoretical predictions of the model Hamiltonian proposed in Ross et

al. [27] which in general reproduced the correct energy ordering of the intense modes was also used, even though it could not account quantitatively for the actual values of the mode energies in our experiments.

6.3 Determining the Optimised Exchange Hamiltonian

A least squares minimisation was carried out to find the optimal values of the six free parameters in the Hamiltonian in Eqn. (6.1), $J_e = \{J_{zz}, J_{\pm}, J_{\pm\pm}, J_{z\pm}\}$, and the two symmetry allowed free parameters of the $\mathbf{g}$ tensor, $g_{\parallel}$ and $g_{\perp}$. This determination of the optimal Hamiltonian parameters was necessary as it will be shown later (see Figs. 6.5 and 6.6) that the previous exchange Hamiltonian for $\text{Yb}_2\text{Ti}_2\text{O}_7$ in the literature, determined using the same principle of fitting magnon dispersions (in a $[1\bar{1}0]$ field) by Ross et al. [27], does not accurately describe the magnon dispersions observed in $[001]$ fields when used in linear spin wave calculations.

To be sufficiently constraining four different data sets were used: the magnon dispersion points in a field of 5 T $\parallel$ $[001]$ (2.5 meV Horace data set), the magnon dispersion points at 9 T $\parallel$ $[001]$ (4.0 meV single angle data), dispersion points at 5 T $\parallel$ $[1\bar{1}0]$ from Ross et al. [27], and the value of the magnetisation in an applied field of 7 T $\parallel$ $[001]$. The magnetisation data is discussed in Appendix D. The least squares minimisation process was performed by starting with several thousand sets of randomly generated initial parameter sets, in order to fully explore the parameter space and ensure that the minimisation yielded a global, not a local, minimum of χ^2 . The fitting procedure was first tested by reproducing the analysis in Ross et al. [27], which yielded exchange parameters in agreement with those found in that work, up to the uncertainty.

Initial least squares fits to the 5 T $\parallel$ $[001]$ data alone showed that this data was not sufficient to uniquely constrain all six parameters (specifically the two components

of the $\mathbf{g}$ tensor) in the Hamiltonian. In other words, the minimisation procedure using the 5 T $\parallel$ [001] data set without a constraint on the $\mathbf{g}$ tensor components leads to degenerate solutions with different sets of parameters. In order to add more constraints, dispersion points at 9 T $\parallel$ [001], points extracted from the dispersion at 5 T $\parallel$ $[1\bar{1}0]$ of Ross et al. [27], and 7 T magnetisation data were included in the fit. Combining all of these data sets did provide sufficient constraints to the model to find a global minimum in χ^2 , that fits all of the dispersion points well (see Fig. 6.3(a)). The weights of the different data points were assigned based on the energy uncertainty of the dispersion points, whereas the magnetisation data, being only a single point, was weighted more heavily than indicated by its uncertainty in order to give it a more balanced weight in the fit compared to the other data sets (which constituted 354 dispersion points). The uncertainty of the points extracted from the data of Ross et al. [27] was set to ± 0.02 meV (as the points were extracted by mouse clicks over a digitised figure), and the uncertainty in the magnetisation was $0.004 \mu_B/T$ compared to a magnetisation per Yb^{3+} moment of $1.695 \mu_B/T$. In order to obtain reasonable values of χ^2 per degree of freedom in the fit, it was necessary to increase the uncertainty of the dispersion points extracted from the 5 T $\parallel$ [001] and 9 T $\parallel$ [001] dispersion points to ± 0.01 meV. This reduced the minimum χ^2 from ~ 5 to ~ 2.4 .

The resulting Hamiltonian parameters that lead to the global minimum in χ^2 are shown in the first column of Table 6.2. These parameters will be used for all of the calculations based on spin wave theory in this work. It is helpful to visualise the robustness of the fit by a plot of χ^2 in the plane of parameters $(J_{zz}, J_{x\pm})$, as in Fig. 6.3(b), as these were proposed to be the dominant interactions by Ross et al. [27]. The global minimum does have nearby local minima, one of which is better illustrated by the plot in Fig. 6.3(c) showing the χ^2 dependence along the dashed white line in Fig. 6.3(b), which passes through the global minimum (star) and the nearby local minimum, and shows that the global minimum is significantly deeper.

The uncertainties in the optimal parameters given in Table 6.2 have been deter-

Table 6.2: The optimised exchange parameters compared to those of Ross et al. [27] and Robert et al. [40].

Parameterisation	This Work	Ross et al. [27]	Robert et al. [40]
J_{zz} (meV)	0.0339 ± 0.0025	0.17 ± 0.04	0.07
$J_{\pm\pm}$ (meV)	0.0462 ± 0.0014	0.05 ± 0.01	0.04
$J_{\pm}$ (meV)	0.0717 ± 0.0011	0.05 ± 0.01	0.085
$J_{z\pm}$ (meV)	-0.1618 ± 0.0013	-0.14 ± 0.01	-0.15
$g_{\parallel}$	2.11 ± 0.02	1.80	2.06
$g_{\perp}$	4.05 ± 0.01	4.32	4.09

mined by using the uncertainty estimates for the dispersion energies, and magnetisation. These uncertainties were sampled using a Gaussian distribution to generate 10000 virtual sets of data points was generated, where the data points in each set were shifted from the original values by an amount within the uncertainty of each given point. The least squares minimisation process was carried out on each of these virtual data sets, and the distribution of the resulting set of global minima was used to determine the uncertainty of the optimised parameters. The resulting uncertainties are likely underestimated, as they only take into account energy uncertainty in the magnon dispersions, and the small magnetisation uncertainty, and do not account for any potential systematic errors in the measurements, though these errors have been minimised.

The existence of parameter sets that fit some reduced part of the data almost as well as the optimal set is due to correlations between various parameters in the Hamiltonian, as shown in Fig. 6.3(d). In principal, none of these parameters are truly independent, because the shape of the dispersion is due to all of the terms in the Hamiltonian. In particular, the orientation of the of he moments depends on the field, exchange, and $\mathbf{g}$ tensor combined together. Moreover the hopping terms that affect the bandwidth of the excitations also depend on the effective spin orientations and on all of the exchange terms in H_{Exchange} . Strong correlations (~ 1) arise between $g_{\parallel}$ and $g_{\perp}$ due to the strong coupling to the field introduced by including magnetisation




Figure 6.3: (a) Experimental magnon dispersion energies, used in the least squares fit, plotted against LSW dispersion energies, showing the level of agreement achieved by the optimised Hamiltonian parameters. The green line represents perfect agreement. Dispersion points for " B " $\parallel$ $[001]$ (blue dots) and " B " $\parallel$ $[\bar{1}10]$ (red dots) are shown. (b) Map of χ^2 generated by fixing the values of J_{zz} and $J_{z\pm}$ and optimising the other four free parameters using a least squares routine. The red star shows the position of the optimised Hamiltonian, the magenta triangle shows the position of the fitted parameters of Ross et al. [27]. (c) χ^2 along the dashed line in (b) showing that the minimum containing the optimised Hamiltonian parameters (red star) is deeper than the nearby local minimum by a significant amount. (d) Absolute values of the correlation coefficients of the different free parameters in the Hamiltonian. (e) 10000 $(g_{\perp}, g_{\parallel})$ points (red circles) used to estimate the uncertainty in the optimised values of $g_{\parallel}$ and $g_{\perp}$, showing that there is a strong linear relationship between $g_{\parallel}$ and $g_{\perp}$, the source of the strong correlation between $g_{\parallel}$ and $g_{\perp}$ in (d). The dashed blue line indicates $g_{\perp}/g_{\parallel} = 1.8$, the constraint used by Ross et al. [27], which would make it impossible to obtain the optimised values of $g_{\parallel}$ and $g_{\perp}$. The black arrows show the directions of maximum and minimum constraint on $g_{\parallel}$ and $g_{\perp}$. The inset shows a close-up of the $(g_{\parallel}, g_{\perp})$ points, highlighting the linear relationship between $g_{\parallel}$ and $g_{\perp}$, along with an ellipse indicating the FWHM of the distribution of the points around their optimised values. The strong correlation between and constraint on $g_{\parallel}$ and $g_{\perp}$ is due to constraints imposed by the inclusion of the magnetisation in the least squares fit.

measurements in the fit. This is illustrated in Fig. 6.3(e), which shows the 10000 $(g_{\parallel}, g_{\perp})$ used to estimate the uncertainties in those two parameters. These points show that $g_{\perp}$ and $g_{\parallel}$ are strongly linearly correlated, leading to the correlation value of ~ 1 . This correlation is perpendicular to the direction of maximum constraint (a linear combination of $g_{\perp}$ and $g_{\parallel}$, indicated by the black arrow perpendicular to the cluster of points) the slope of which is imposed by the global z component of the full magnetic moment. These points also show that the constraint $g_{\parallel}/g_{\perp} = 1.8$, used by Ross et al. [27], would have made it impossible to obtain the optimised values of $g_{\perp}$ and $g_{\parallel}$. If the magnetisation constraint is not included, the correlation between $g_{\parallel}$ and $g_{\perp}$ is reduced, and J_{zz} and $g_{\parallel}$ are strongly correlated as they have similar effects on the dispersion. When the magnetisation is included, these two terms are decoupled somewhat. These correlations also demonstrate that determining the correct exchange parameters to describe the magnon dispersions in $\text{Yb}_2\text{Ti}_2\text{O}_7$ is not a simple process of fitting certain features of the data to determine each parameter, as the value of all of the free parameters significantly impact all of the other parameters, with none of the fitted parameters being completely independent. This is in contrast to the case of isotropic Heisenberg exchange and an isotropic $\mathbf{g}$ tensor, where the dispersion bandwidth is simply the Fourier transform of the exchange shifted by the Zeeman energy, so the exchange and $\mathbf{g}$ tensor are almost completely decoupled [92].

As discussed in Section 5.1.1, the final orientation of the sample in the neutron scattering instrument was found have the crystallographic [001] axis tilted away from the field axis of the superconducting magnet by an angle of 1.03° . The least squares fit was performed both with and without this tilting of the applied field, and the resulting changes in the optimised parameters were found to be within the quoted uncertainties in Table 6.2.

All of the magnetic fields used in the spin wave and mean-field magnetisation calculations in the least squares fit were the internal field of the sample, taking into account demagnetisation effects, as discussed in Appendix D. Ross et al. [27] did not account for demagnetisation effects, so the magnetic field used to compute the

spin waves for comparison to those data points was allowed to vary in an attempt to account for demagnetisation effects. The internal field along the $[1\bar{1}0]$ direction that lead to the best fit to their data was $\mu_0 H = 4.93$ T, a small reduction compared to the applied field of 5 T.

Figure 6.4 shows the results of calculations of ΔS , the leading order reduction in the ordered spin moment from the maximum available value of $1/2$, computed using LSWT as follows:

$$\Delta S \equiv \frac{1}{4N} \sum_{i,a} \langle c_{ia}^\dagger c_{ia} \rangle = \frac{1}{4N} \sum_{\mathbf{k}} \sum_{a,m} \Lambda_{a\mathbf{m}\mathbf{k}}^* \Lambda_{a\mathbf{m}\mathbf{k}}, \quad (6.30)$$

where i and a index the FCC lattice sites, and tetrahedral sublattice sites respectively. In the second form of ΔS , $\mathbf{k}$ is a vector in the first Brillouin zone, m indexes mode number, and N is the number of FCC lattice sites. $\Lambda_{m\mathbf{k}}$ are rotated eigenvectors, with elements defined by:

$$\Lambda_{a\mathbf{m}\mathbf{q}} \equiv \frac{1}{2} (v_{a\mathbf{m}\mathbf{q}}^u - i v_{a\mathbf{m}\mathbf{q}}^v). \quad (6.31)$$

In an applied field of 7 T, the calculated reduction in the ordered moment due to zero point fluctuations is less than 1%, compared to a reduction in zero field of 17%. This means that 7 T is sufficiently close to the saturated limit for the mean-field calculated value of the magnetisation per Yb^{3+} moment to be fitted to the experimental value.




Figure 6.4: LSW calculation of ΔS , the leading order quantum reduction in the ordered spin moment from its maximum possible value of $1/2$ as a function of applied field, which shows that the quantum fluctuations in an applied field of 7 T are small.

The set of optimised parameters resulting from this least squares fit is quite different from that of the previous fit to high field magnon dispersions of Ross et al. [27]. In that work, it was found that both J_{zz} and $J_{z\pm}$ were significant terms in the exchange Hamiltonian. In this work, we find that J_{zz} is much smaller than $J_{z\pm}$, in fact $J_{z\pm}$ appears to be the dominant exchange term. This would indicate that any theory based on perturbing away from a dominant J_{zz} theory, as done in Refs. [27, 49, 56, 57, 93] will likely not be applicable to the physics of $\text{Yb}_2\text{Ti}_2\text{O}_7$. Our optimal exchange parameters are very similar to those of Robert et al. [40], which recently appeared in the literature. The parameters of Robert et al. [40] were determined by fits to the data of Ross et al. [27] in an applied field of 2 T, and zero field quasi-elastic scattering. The $\mathbf{g}$ tensor parameters resulting from the least squares fit are very similar to those determined experimentally in the work of Cao et al. [23], where the two parameters in the $\mathbf{g}$ tensor were found to be $g_{\parallel} = 2.24$ and $g_{\perp} = 3.98$.

6.4 One Magnon Dispersions: Data and Spin Wave Calculations

Figures 6.5 and 6.6 show comparisons of all of the slices through the 5 T $\parallel$ [001], 2.5 meV data to LSW calculations based on the optimised Hamiltonian in Table 6.2, and the parameters determined by Ross et al. [27]. The white dots in all of the plots show the energies of the empirically extracted dispersion points that were used in the least squares optimisation. It can clearly be seen that the optimised parameters give a much better fit to the observed dispersion than the parameters of Ross et al. [27], which do not describe the dispersion of several intense modes in the data, see Fig. 6.5(i-l,v-x) and Fig. 6.6(i,w). It is this inability of the Hamiltonian determined by Ross et al. [27] to describe the dispersions observed in the 5 T $\parallel$ [001], 2.5 meV Horace data that necessitated the re-determination of the Hamiltonian. The plots of the spin wave calculations were generated by convolving the model dispersions with a Gaussian line shape with a FWHM of 0.047 meV, assumed constant as a

function of energy transfer. The simulations have also been integrated over the same ranges in reciprocal space that were used to generate the experimental data figures. It should be noted that all of the LSW calculations in applied field presented in this work are performed using the demagnetisation corrected internal fields given in Appendix D for the specified experimental applied field.

Figure 6.7 shows a comparison of LSW dispersions computed using the optimised Hamiltonian parameters and the 9 T $\parallel$ [001], 4.0 meV single angle data. Figure 6.7(a) shows the experimental data, along with the extracted dispersion points, while Fig. 6.7(b) shows the LSW dispersions. It can be seen that the spin wave Hamiltonian captures very well the dispersion of all the modes. The finite width of the magnon modes in Fig. 6.7(b) is a result of convolving the LSW calculations with a resolution profile matched to the line shape of the magnon modes in the 2.5 meV 5 T $\parallel$ [001] Horace data, and an estimated resolution line shape of the 4.0 meV elastic line. This line shape has an effective FWHM of 0.227 meV. The convolution is performed by Monte Carlo sampling of the inelastic line shape, and applying this to the inelastic modes. The simulation is integrated over the same range in reciprocal space as the data.

Figure 6.8 shows the 5 T $\parallel$ $[1\bar{1}0]$ magnon dispersion data of Ross et al. [27], along with calculations of the same data based on the optimised Hamiltonian of this work, and the Hamiltonian of Ross et al. [27]. This figure shows that the optimised Hamiltonian of this work can reproduce the significant features of the dispersion observed in Ross et al. [27] at least as well as the set of Hamiltonian parameters proposed in that work. The simulations in Fig. 6.8 do not included any integration of the scattering data in reciprocal space, as the integration ranges used by Ross et al. [27] are not given, and integrating the dispersion simulations over small ranges perpendicular to the slice direction was not found to have a significant impact on the appearance of the dispersions. The LSW dispersions in Fig. 6.8 have been convolved with a Gaussian with a FWHM of 0.089 meV.

Figure 6.9 shows a comparison of a scan through the 5 T $\parallel$ [001], 2.5 meV single




Figure 6.5: Slices through the 2.5 meV 5 T Horace data. Panels (a-d,m-p) show experimental data, while panels (e-h,q-t) show spin wave calculations for the optimised parameter set, and panels (i-l,u-x) show spin wave calculations for the model of Ross et al. [27]. The white dots represent the $(\mathbf{Q}, E)$ points used to constrain the fitted Hamiltonian. The directions of these slices in the $(h, k, 0)$ plane are shown in the diagram in Fig. 6.1(c). The slices have been generated by integrating over the range $-0.1 \leq l \leq 0.1$, and $-0.1 \leq h \leq 0.1$ (r.l.u.) for panel (a,e,i) and equivalent ranges perpendicular to the slice directions of the other panels.




Figure 6.6: Slices through the 2.5 meV 5 T Horace data. Panels (a-d,m-p) show experimental data, while panels (e-h,q-t) show spin wave calculations for the optimised parameter set, and panels (i-l,u-x) show spin wave calculations for the model of Ross et al. [27]. The white dots represent the $(\mathbf{Q}, E)$ points used to constrain the fitted Hamiltonian. The directions of these slices in the $(h, k, 0)$ plane are shown in the diagram in Fig. 6.1(c). The slices have been generated by integrating over the range $-0.1 \leq l \leq 0.1$, and $-0.1 \leq h \leq 0.1$ (r.l.u.) for panels (m,q,u) and equivalent ranges perpendicular to the slice directions of the other panels.




Figure 6.7: (a) 9 T $\parallel$ [001], 4.0 meV inelastic neutron data, and the dispersion points (white circles) used in the least squares fit. (b) LSW simulations of the data set, computed using the optimised Hamiltonian parameters, compared to the dispersion points used in the least squares fit. The integration ranges used to generate these slices are the same as those used in Fig. 6.2(b).

angle data to show the quantitative agreement of the optimised LSW calculations with the 5 T $\parallel$ [001], 2.5 meV data. The 2.5 meV single angle data set is used instead of the Horace data, as the statistics of single angle data are better than the Horace data. Figure 6.9(a) shows the 5 T $\parallel$ [001], 2.5 meV data set, while Fig. 6.9(b) shows the LSW simulation of the same data set. The spin wave calculations have been convolved with a line shape based on fitting the magnon line shape of the 5 T $\parallel$ [001], 2.5 meV Horace data, consisting of multiple Gaussians with an effective FWHM of 0.214 meV. The overall intensity of the LSW scattering is a free parameter that has been fixed to the intensity of the most intense mode. Figure 6.9(c) shows the 5 T $\parallel$ [001] data, along with the spin wave calculation of the same data, and 9 T data, showing the non-magnetic background. The comparison between the model and the data shows that the model fits the two most intense modes very well, but underestimates somewhat the intensity of the weakest modes. The energy of the highest energy magnon mode is also slightly underestimated. These are regarded as very small differences between the data and model, and the fit could probably be improved further by allowing more terms in the exchange Hamiltonian beyond nearest-neighbour, but since the effect is




Figure 6.8: Spin wave spectra in $[1\bar{1}0]$ field: left column shows data from Ross et al. [27], middle column shows spin wave calculations for the optimised Hamiltonian parameters, right column shows spin wave calculations for the parameters of Ross et al. [27].

very small (1% difference), it is beyond the scope of the present work.




Figure 6.9: (a) shows 2. meV, 5 T single angle inelastic neutron data. This slice was generated using an integration range of $-0.5 \leq l \leq 0.5$ (r.l.u.). (b) shows spin wave calculations of the same data, generated by convolving LSW calculations with an elastic line shape with a FWHM of 0.214 meV. The dashed lines in panels (a) and (b) show the h integration range ($-1.8 \leq h \leq -1.4$) used to generate the scans in (c). (c) shows scans in energy through the experimental data and simulations, along with a 9 T non-magnetic background.

Figure 6.10 shows the 2.5 meV single angle data of Fig. 5.2 compared to spin wave calculations performed in the same manner as for the 5 T data. Examining the data, it can be seen that for applied fields of $B \geq 3$ T, the LSW model and the experimental data are very similar. For fields below 3 T, LSWT does not describe the observed excitations, this will be discussed in Section 6.6.

Figure 6.11 shows a direct comparison of the one and two magnon bandwidth as determined by experiment and calculated by LSWT. This figure shows that for high fields ($B > 3$ T), the spin wave model agrees with the one and two magnon bandwidths quite well. At 3 T, the experimental data indicates that the one and two magnon scattering are very close to cross over (See Fig. 5.3(h)), and it can be seen in Fig. 6.11 that at this field the behaviour of the experimental data and LSW calculations deviate. The behaviour at this and lower fields will be discussed in Section 6.6.




Figure 6.10: Overview of the magnetic excitations as a function of applied field (top rows) compared with LSWT predictions (bottom rows). The spin wave plots are generated using the same method as Fig. 6.9.




Figure 6.11: Comparison of experimental (filled circles) and LSW model bandwidths (lines) as a function of applied field. The calculated one magnon lower and upper bound of the dispersion correspond to the lowest and highest spin wave energies at $\mathbf{Q} = 0$. The two magnon bounds correspond to twice the lower and upper bounds of the one magnon dispersion.

6.5 High Energy Two-Magnon Continuum Scattering: Data and Calculations

As discussed in Chapter 5.2.1, evidence of two magnon scattering is also found in the high field data. It is possible to compute two-magnon scattering using spin wave theory. As discussed, the lowest order term in spin wave theory that contains two magnon inelastic terms arises in the longitudinal components (rather than the transverse components, where the single magnon scattering arises) of the effective spin-1/2 objects. The longitudinal component of the dynamical correlation function, in local quantisation coordinates, written in terms of effective spin-1/2 operators, is given by [84]:

$$\tilde{\mathcal{S}}_{\text{eff},ab}^{ww}(\mathbf{Q}, \omega) \equiv \sum_e \langle 0 | \tilde{\mathcal{S}}_a^w(-\mathbf{Q}) | e \rangle \langle e | \tilde{\mathcal{S}}_b^w(\mathbf{Q}) | 0 \rangle \delta(\omega - \omega_e), \quad (6.32)$$

where e labels states of the system, including excited states that contain pairs of magnon excitations, which will have finite matrix elements for $\tilde{\mathcal{S}}_b^w(\mathbf{Q})$; and $\omega_e = (E_e - E_0)/\hbar$, where E_e is the energy of the state e and E_0 is the energy of the ground

state of the system. The Fourier transformed w (local quantisation frame z) spin-1/2 operator written in terms of Holstein-Primakoff boson operators (Eqn. (6.8)) is:

$$\begin{aligned} \tilde{S}_a^w(\mathbf{Q}) = & \sqrt{N} S \delta_{\mathbf{K},0} \exp(-i\mathbf{G} \cdot \mathbf{r}_a) \\ & - \frac{1}{\sqrt{N}} \sum_{\mathbf{q},\mathbf{k}} \delta(\mathbf{K} - (\mathbf{q} - \mathbf{k})) c_{\mathbf{q},a}^\dagger c_{\mathbf{k},a} \exp(-i\mathbf{G} \cdot \mathbf{r}_a), \end{aligned} \quad (6.33)$$

where $\mathbf{Q} = \mathbf{K} + \mathbf{G}$. The delta functions are evaluated with respect to the shifted $\mathbf{Q}$. $\mathbf{q}$ and $\mathbf{k}$ are wave vectors inside the first Brillouin zone of the FCC lattice, and N is the number of FCC lattice sites. Terms in Eqn. (6.32) when expanded using Eqn. (6.33) containing S and S^2 contribute to magnetic Bragg scattering associated with long range magnetic order. Terms in the expansion containing S^0 are inelastic, and ultimately describe the creation of pairs of magnons, giving rise to finite matrix elements for excited states containing pairs of magnons. The inelastic component of the longitudinal dynamical correlation function is given by:

$$\begin{aligned} \tilde{\mathcal{S}}_{\text{eff } ab, 2\text{-magnon}}^{ww}(\mathbf{Q}, \omega) = & \frac{1}{N} \sum_{\mathbf{q},\mathbf{k}} \sum_{m \geq n} \delta(\mathbf{K} - \mathbf{q} - \mathbf{k}) \delta(\omega - \omega_{m\mathbf{q}} - \omega_{n\mathbf{k}}) \\ & \times \exp(-i\mathbf{G} \cdot \mathbf{r}_{ab}) \Omega_{\mathbf{q},\mathbf{k},amn}^* \Omega_{\mathbf{q},\mathbf{k},bmn} \end{aligned} \quad (6.34)$$

where m, n index single magnon modes. The quantity $\Omega_{\mathbf{q},\mathbf{k},amn}$ is defined as:

$$\Omega_{\mathbf{q},\mathbf{k},amn} \equiv \Lambda_{am\mathbf{q}} \bar{\Lambda}_{an\mathbf{k}} + \bar{\Lambda}_{am\mathbf{q}} \Lambda_{an\mathbf{k}} \left(1 - \left(2 - \sqrt{2} \right) \delta_{mn} \delta(\mathbf{k} - \mathbf{q}) \right). \quad (6.35)$$

where $\Lambda_{am\mathbf{q}}$ are elements of the rotated eigenvectors defined in Eqn. (6.31), and $\bar{\Lambda}_{am\mathbf{q}}$ are elements of the rotated eigenvectors $\bar{\Lambda}_{m\mathbf{k}}$, with elements

$$\bar{\Lambda}_{am\mathbf{q}} \equiv \frac{1}{2} (v_{am\mathbf{k}}^u + i v_{am\mathbf{k}}^v). \quad (6.36)$$

The effective spin-1/2 two magnon scattering, in global Cartesian components,

can be used to compute the full moment (J) two magnon scattering via the relation:

$$\mathcal{S}_{2\text{-magnon}}^{\alpha\beta}(\mathbf{Q}, \omega) = \sum_{a,b} \Gamma_a^{\alpha,z} \Gamma_b^{\beta,z} \tilde{\mathcal{S}}_{\text{eff } ab, 2\text{-magnon}}^{ww}(\mathbf{Q}, \omega). \quad (6.37)$$

The sum in Eqn. (6.35) over $\mathbf{q}, \mathbf{k}$ is performed using Monte-Carlo methods.

Figure 6.12 shows a comparison of LSW calculations of both the one and two magnon scattering in an applied field of 7 T (equivalent to an internal field in the sample of $\mu_0 H_{\text{int}} = 6.93$ T) to experimental data. The spin wave calculations have been convolved with an inelastic line shape based on an estimate of the magnon line shape of the Horace data. In the 7 T data (red circles), the high energy band of scattering (2.7 - 4.0 meV) is a genuine magnetic scattering signal, not present in at low fields, as indicated by the 0.21 T data (black circles). Looking at the spin wave calculations, it can be seen that the single magnon scattering is fit quite well by the model, with only one free parameter, that being the overall intensity of the scattering. The two magnon scattering based on LSWT appears to fit the bounds of the experimental scattering well, but underestimates the intensity of the two magnon scattering. The dashed line shows the two magnon calculation scaled up by a factor $Z = 3.9$, to match the total intensity of the continuum, showing that the line-shape of the two magnon continuum scattering is reproduced reasonably well by the spin wave calculations.

Figure 6.13(a-g) shows the 6.3 meV single angle continuum scattering data, along with LSW calculations for applied fields between 3 T and 9 T. There it can be seen that non-interacting spin wave theory underestimates the weight of the two magnon continuum scattering for all fields where the two magnon scattering can be differentiated from the one magnon scattering (for 3 T, the one and two magnon scattering can be differentiated in the 2.5 meV data, but not in the 6.3 meV data, due to the coarser resolution in the latter case). It is interesting to note that the energy dependence of the two magnon line-shape is similar in shape to the spin wave calculations at high fields, but as the field is lowered below 6 T, more weight builds up at the lower bound-




Figure 6.12: Energy scan at 7 T: the low energy set of unresolved peaks are due to one magnon scattering and the high energy continuum band is due to two magnon scattering. For ease of visualisation, the data above 2.3 meV (dashed red line) has been scaled up by $\times 15$ to make the continuum visible on the scale of the strong one magnon peaks. The solid blue line shows LSW calculations of the one and two magnon scattering. The dashed green line shows the two magnon scattering calculation scaled up by an empirical factor of $Z = 3.9$ to match the total intensity of the data. These scans were generated by integrating 6.3 meV single angle data over $-4.0 \leq h \leq 0.0$ and $-0.5 \leq 0.5$. The LSW calculations were generated by convolving calculated spectra with an experimentally derived line shape. 0.21 T data is provided to show the non-magnetic background in the energy range where the two magnon scattering is present.

ary of the continuum. The origin of the change in the two magnon continuum might be related to magnon-magnon interactions, which are generically expected to become more relevant as the magnetic field is decreased, and may have more effect on the low energy part of the continuum scattering. Figure 6.13(h) shows a plot of the total two magnon scattered intensity in the energy scans of Fig. 6.13 as a function of applied field, for all fields where the two magnon scattering can be clearly differentiated from the single magnon scattering in the experimental data. Also shown is the two magnon intensity from LSWT, multiplied by a scaling factor of $Z = 3.9$. The total intensity is based on summing the scattered intensity as a function of energy over the energy range of the continuum. This data shows that LSWT does correctly compute the field dependence of the intensity, only being off by an overall scale factor. The scaling factor Z has been applied to the spin wave calculations in Fig. 6.13(a-g) (green dashed

lines). After the application of this scaling factor, the spin wave calculations reproduce the two magnon continuum scattering line-shape and intensity well for fields of $B \geq 7$ T, as was seen in Fig. 6.12. This would suggest that the ordered moments of the Yb^{3+} ions may be smaller than LSWT would indicate, even in high fields, leading to the discrepancy in the intensities, as the two magnon scattering arises from fluctuations in the ordered moment, and a reduced ordered moment would allow for greater longitudinal fluctuations.

One possible way to estimate the increase in the longitudinal fluctuations is using the neutron scattering sum rules. The overall neutron scattering sum rule is given by:

$$\frac{1}{N} \sum_{\mathbf{q}, \alpha} \int d\omega \mathcal{S}_{\text{eff}}^{\alpha\alpha}(\mathbf{q}, \omega) = S(S+1), \quad (6.38)$$

where $\mathcal{S}_{\text{eff}}$ is the scattering cross section for the effective spin-1/2 moments, whose components are in global coordinates, N is the number of lattice sites, and $\mathbf{q}$ is a vector inside the first Brillouin zone. This sum rule means that the total of all of the neutron scattering at all energies and wave vectors inside the first Brillouin zone by the effective spins is $S(S+1)$, similarly, if the Hamiltonian were expressed in terms of total angular momentum $\mathbf{J}$, then the total neutron scattering would be $J(J+1)$. For spin-1/2 moments, there are also neutron scattering sum rules for each of the components of the spin in the local quantisation directions:

$$\frac{1}{N} \sum_{\mathbf{q}} \int d\omega \tilde{\mathcal{S}}_{\text{eff}}^{\alpha\alpha}(\mathbf{q}, \omega) = \frac{1}{3} S(S+1), \quad (6.39)$$

where $\alpha = u, v, w$, the quantisation frame components of the spin-1/2. Table 6.3 shows that all of these sum rules are approximately satisfied by LSW calculations based on the optimised Hamiltonian, despite the fact that LSWT is not required by construction to satisfy these rules. The satisfaction of the total scattering sum rule is shown in the sum of all the contributions to the neutron scattering (labelled Total). For scattering from a spin-1/2 system, this sum should be equal to 0.75, with the actual total overestimating the scattering by $< 1\%$. The local sum rules are shown in




Figure 6.13: Comparison of energy scans through the two magnon continuum to spin wave calculations for fields of $3 \text{ T} \leq B \leq 9 \text{ T}$. Spin wave calculations (solid blue lines) underestimate the intensity of the continuum at all fields, but the same calculations multiplied by a factor $Z = 3.9$ (dashed green lines) reproduce the intensity and energy dependence of the continuum well at high fields $B \geq 7 \text{ T}$. Z is the ratio of the total experimental two magnon continuum scattering intensity to the total two magnon spin wave calculation intensity for fields of $5 - 9 \text{ T}$. These two quantities are plotted in (h), showing the good agreement achieved using the scaling factor $Z = 3.9$. The scans are generated by integrating over the reciprocal space ranges $-4.0 \leq h \leq 0.0$ and $-0.5 \leq k \leq 0.5$. The spin wave calculations have been resolution convolved with a 6.3 meV magnon line shape.

Tab. 6.3 via the one magnon scattering intensity, and the combination of the elastic and two magnon scattering intensity. The component sum rules are only true for the local quantisation directions u, v, w , not the global Cartesian coordinates used for the calculations in Tab. 6.3. It is still possible to extract the local sum values from the calculations as in LSWT, the scattering from $\tilde{\mathcal{S}}_{\text{eff}}^{ww}$ contributes only to the elastic scattering and two magnon scattering, as the first arises from the ordered moments, and the two magnon scattering arises from longitudinal fluctuations of the moments, while the local transverse scattering $\tilde{\mathcal{S}}_{\text{eff}}^{uu}$ and $\tilde{\mathcal{S}}_{\text{eff}}^{vv}$ contribute only to the one magnon scattering. This means that the sum of the elastic and two magnon scattering should be equal to 0.25 and the total one magnon scattering should be 0.5. Table 6.3 shows that both the $\tilde{\mathcal{S}}_{\text{eff}}^{ww}$ sum rule and the $\tilde{\mathcal{S}}_{\text{eff}}^{uu} + \tilde{\mathcal{S}}_{\text{eff}}^{vv}$ are overestimated by $< 1\%$.

Table 6.3: Components of the total moment sum rule at 9 T. The elastic contributions is $(S - \Delta S)^2$ with $\Delta S = 0.0034$ computed at 9 T. The one magnon and two magnon contributions were computed using the integral in Eqn. (6.38), replacing $\mathcal{S}_{\text{eff}}$ with the effective spin scattering cross section for the two forms of scattering. The final two entries are the value of the total sum rule (Eqn. (6.38)), and the local z (defined as w) sum rule, which is the sum of the elastic and two magnon scattering intensities.

Elastic	One Magnon	Two Magnon	Total	Longitudinal
0.2467	0.5025	0.0056	0.7548	0.2523

Using these sum rules, it is possible to estimate what value of ΔS would lead to the correct total intensity for the two magnon scattering. This can be done using the relation:

$$\frac{2 \text{ magnon intensity (experiment)}}{1 \text{ magnon intensity (experiment)}} \approx \frac{0.0056 \times 3.9}{0.5025} = \frac{\Delta S (1 - \Delta S)}{0.5}, \quad (6.40)$$

where the relationship between the total two magnon intensity is derived from the relationship $S(S + 1)/3 - (S - \Delta S)^2 = \Delta S(1 - \Delta S)$. This assumes that all of the longitudinal fluctuations are accounted for by the two magnon scattering, which is expected to be approximately true, as at 9 T the system should be approaching the fully polarised phase where LSWT is exact, so higher order terms should be small.

Solving Eqn. (6.40) yields $\Delta S = 0.0222$, an increase of $\sim 6.5\times$, with the new ΔS being 4.4% of S . This result is much larger than the $< 1\%$ of S found for ΔS directly from the LSW fit to the data, but does show that it would require only a small reduction in the ordered moment to lead to a large increase in the two magnon scattering.

This increased ΔS does not invalidate the fit performed using LSWT and zero temperature mean-field theory, as the two magnon scattering intensity is a higher order effect in spin wave theory, that is not expected to be captured perfectly by the model. It has been shown previously that magnon interactions, accounted for in higher order spin wave theory, are important in this material and have a significant impact on the two magnon scattering at lower fields. At infinite order in the Holstein-Primakoff boson operators, where these effects are fully accounted for, spin wave theory is exact and must reproduce the relative intensities of the various scattering terms. Magnon interactions will redistribute intensity between the elastic, one magnon, and two magnon scattering. At high fields, the number of extra terms necessary to improve the agreement between the experimental and calculated two magnon intensity should not be excessive as the experimental conditions approach the infinite field limit where LSWT is exact. Additionally, when the magnetisation constraint, the constraint most likely to be affected by a change in ΔS is not included in the least squares fit, there is still a global minimum in χ^2 near the optimised parameter set in Tab. 6.2, the magnetisation constraint serves mainly to better define the values of $g_{\parallel}$ and $g_{\perp}$.

It is possible that the ordered moment of the Yb^{3+} moments in high field is smaller than is predicted by this model compared to the free ion ordered moment, allowing for more substantial fluctuations of the longitudinal part of the moment, and thus more intense two magnon scattering. In order to determine if this is the case, the field dependence of the elastic magnetic scattering would need to be determined, which is not possible using the data in this work. Comparing the high field magnetisation to the mean-field magnetisation computed using the optimised Hamiltonian will not reveal any information about a reduction in the ordered moment compared to the

model, as the magnetisation is a constraint in the fit, though it can be seen in Fig. D.2 that in the low field regime, the mean-field approach does not correctly predict the ordered moment, overestimating the zero field moment by $\sim 50\%$.

One aspect of the high field two magnon continuum that LSWT appears to capture, other than the energy bounds, is the intensity modulation in $(h, k, 0)$ plane. Figure 6.14 shows a comparison of the LSW calculation of the total intensity of the two magnon continuum in the $(h, k, 0)$ plane in an applied field of 5 T $\parallel$ [001], compared to the constant energy slice originally shown in Fig. 5.6(d). The intensity of the LSW two magnon scattering has been multiplied by a factor of $Z = 3.9$ to put it on the same scale as the experimental data. Both panels of Fig. 6.14 clearly show a minimum in the two magnon intensity around $\mathbf{Q} = (-2, 0, 0)$.

6.6 Magnon Interaction and Decay at Intermediate Applied Fields

As discussed in Chapters 5 and 6.4, for applied magnetic fields below 3 T, the single magnon modes and two magnon continuum intersect, with the uppermost magnon mode having the same energy as the lower boundary of the two magnon continuum. In this field range LSW calculations fail to capture significant aspects of the experimental data. Starting at 2 T, the most significant aspect of the experimental data that LSW calculations fail to capture is the broadening of the highest energy magnon mode when it overlaps with the continuum, as discussed in Ch. 5.2.1. This feature can be attributed to decay of the single magnon mode into two magnon excitations. The other disagreement between the data and LSW calculations at 2 T is near $h = -1.5$ for an applied field of 2 T (see Fig. 6.10(k,o)), the second mode, in order of increasing energy, is concave down in the LSW calculations, whereas it is flat or even concave up in the experimental data. Below 2 T (Fig. 6.10(a-j,m,n)), the experimental data shows three sharp excitations below a continuum of scattering, which is a combination of the




Figure 6.14: Energy integrated slice through the 5 T Horace Scan, 6.3 meV Horace data in the $(h, k, 0)$ plane (a), along with the same slice generated using LSWT (b). The slices were generated by integrating over the range $-0.2 \leq l \leq 0.2$, and the energy range $2.1 \leq E \leq 3.0$ meV, which encompasses the continuum scattering at this field. The LSW calculated intensities have been scaled up by a constant factor of $Z = 3.9$.

uppermost magnon mode and two magnon scattering. As the field is decreased, the three sharp modes flatten (become less dispersive) and appear to become degenerate, with only a single sharp mode visible below the continuum at very low fields. None of this behaviour is captured by the LSW calculations.

As previously discussed, Fig. 5.3 shows that at 3 T, the upper bound of the single magnon modes and the lower bound of the two magnon continuum scattering are very close together, with the bottom boundary of the continuum scattering being unresolvable in 6.3 meV data (see Fig. 6.13). Figure 6.11 shows that for LSW calculations, there is also a very small gap between the single magnon upper bound and the lower boundary of the two magnon continuum. There is some disagreement between LSW calculations and experiment at this field: the upper bound of the two magnon continuum is at higher energy than predicted by LSWT, possibly due to avoided crossing of the sharp modes and continuum, as previously discussed in Section 5.2.1. Below 3 T, the inability of the LSW calculations to describe the upper bound of the magnetic continuum continues down to zero field. At these fields, it is not possible to determine the boundary between one and two magnon scattering, only the upper and lower bounds of the magnetic scattering can be seen in experiment. As discussed, below 2 T, LSW calculations do not reproduce the sharp modes in the magnetic scattering. If the continuum scattering at these fields can still be thought of as a two magnon continuum arising from combinations of these sharp modes, then the inability of LSW calculations to reproduce the single magnon modes means that LSW calculations of the two magnon continuum will not describe the energy boundaries of the continuum scattering well.

In zero applied field, there are two major differences between the LSW calculations and the experimental neutron scattering spectrum: the most obvious is the fact that the experimental zero field spectrum is a continuum, while the zero field LSW spectrum shows distinct modes (Fig. 6.10(a,e)); the other is that the predicted LSW excitations are gapped all along $(h, 0, 0)$, while multiple data sets (Figs. 5.3, 5.4, and 5.11) confirm that there is no gap to within the resolution of the experiments, down

to 0.04 meV. The existence of a gap in the LSW calculations in zero applied field is a predetermined feature of the model, as it is based on excitations from a discrete set of symmetry related ordered ground states, which must have an energy gap associated with its excitations. Figure 6.15 shows a scan in energy through the 0 T data, along with the same scan generated using LSWT, clearly showing the energy gap to the lowest energy excitation in the LSW spectrum, and the lack of any gap in the experimental data.




Figure 6.15: scans in energy generated by integrating the experimental 2.5 meV single angle data over the ranges $-2.1 \leq h \leq 0$ and $-0.5 \leq l \leq 0.5$. The model plot shows the result of convolving the 0 T LSW spectrum with a line shape based on fitting the inelastic mode shape of the 2.5 meV Horace data and integrating over the same ranges in reciprocal space. The 9 T points illustrate the quality of the background subtraction.

All of this evidence combines to show that the LSW method of calculating the excitation spectrum of the Hamiltonian breaks down for applied fields of < 3 T. So long as magnons are indeed the excitations associated with the ground state of the magnetic system, exact spin wave theory (i.e. not restricted to order $1/S$) should be able to describe the observed excitation spectrum. These higher order terms represent magnon interactions, which will be significant in low fields, where the spins are not polarised by the field. The mean-field ground state for all fields is an ordered ferromagnetic phase, so magnons should be the correct description of the excitations in the system, and it should be possible to improve the agreement between the exper-

imental magnon dispersions and calculations at these low fields by including terms to higher order in $1/S$ in the calculations. The zero field mean-field order is similar to that proposed by Chang et al. [33] to explain their observed magnetic Bragg peaks. Any applied magnetic field along [001] reinforces the broken symmetry of the mean-field ground state, making it even more likely that the mean-field ground state is the correct description of the system in an applied field. This means that magnons are likely the correct description of the excitations in all fields, so higher order terms in spin wave theory are likely to account for the differences between the experimental scattering and LSW calculations.

As discussed, when the energy ranges of the one and two magnon scattering overlap, the highest energy single magnon mode is significantly broadened. The likely mechanism for this broadening is the spontaneous decay of the single magnon excited by the incident neutron into two magnons, as this is energetically allowed once the one and two magnon phase space overlap. This is the leading contribution, and is allowed at order $1/S$, without including higher order terms. For the optimised exchange parameters in Table 6.2, the semi-classical magnetic order is non-collinear over the entire field range explored in this work. This non-collinearity means that, in the local quantisation frame, an exchange Hamiltonian with isotropic (Heisenberg) exchange will contain terms that couple $\tilde{S}^u$ at one site to $\tilde{S}^w$ of another site. These exchange terms will then lead to cubic terms when expanded in Holstein-Primakoff bosons, which annihilate one magnon, and produce another two. In the case of $\text{Yb}_2\text{Ti}_2\text{O}_7$, these terms arise due to non-collinear local quantisation frames, but also due to explicit anisotropic exchange terms like $J_{z\pm}\mathbf{S}_{ia}^z\mathbf{S}_{jb}^\pm$.

The decay rate of one magnon into two magnons via three boson terms is given by Fermi's golden rule [94]:

$$\Gamma_{\mathbf{k}} \propto \sum_{\mathbf{q}} |V_3(\mathbf{q}, \mathbf{k})|^2 \delta(E_{\mathbf{k}} - E_{\mathbf{q}} - E_{\mathbf{k}-\mathbf{q}}), \quad (6.41)$$

where $\Gamma_{\mathbf{k}}$ is the decay rate for one magnon at wave vector $\mathbf{k}$ inside the first brillouin

zone into two magnons of mode at wave vectors $\mathbf{q}$ and $\mathbf{k} - \mathbf{q}$. The sum over $\mathbf{q}$ is over the first Brillouin zone. The original magnon has energy $E_{\mathbf{k}}$, while the resulting two magnons have energies $E_{\mathbf{q}}$ and $E_{\mathbf{k}-\mathbf{q}}$. $V_3(\mathbf{q}, \mathbf{k})$ is the potential energy of the three boson interaction where two bosons are created and one is destroyed [94].

Initial calculations based on Eqn. (6.41) show that in LSWT, 3 T $\parallel$ [001] is a threshold applied field in terms of magnon decay. For applied fields of 3 T and above along [001], there is no decay, while below this field, there is a finite density of states for decay, which increases with decreasing field. This is consistent with observations in the experimental data, which shows that for applied fields below 3 T, the uppermost magnon mode is significantly broadened, eventually evolving into a broad continuum of scattering. The other three sharp magnon modes appear to be significantly re-normalised, becoming less dispersive and appearing to merge into a single flat mode at very low applied fields. Investigations of magnon decay and other effects are ongoing.

One aspect of Fig. 6.11 that should be discussed with regards to magnon interactions is that, for all finite applied fields, LSW calculations describe the lower bound of the magnetic scattering quite well. This is confusing, given the inability of LSW calculations to describe other properties of the magnon dispersions for fields < 3 T. The addition of magnon interactions is expected to lower the energies of the magnon modes [83]. This means that magnon interactions that will help to improve agreement between LSW calculations and the experimental dispersions can be expected at lowest order to reduce the agreement between LSW calculation and the experimental lower boundary of the magnetic scattering.

When one carries out higher order spin wave theory, one expects to obtain broadened, re-normalised excitations modes [94], not the continuum of scattering observed in zero field, with no obvious modes. In Section 5.2.2 the $\mathbf{Q}$ modulation of the intensity of the continuum scattering in zero field was compared to the scattering observed in the Kagome antiferromagnet Herbertsmithite. In this material, an apparently gapless continuum of inelastic scattering is observed, with strong intensity modulations [88].

A proposed theory to explain this scattering is that the low temperature magnetism in Herbertsmithite is described by a Z_2 spin liquid [89], though this model has some distinct differences from experiment, namely the continuum of excitations for such a magnetic phase is gapped, though it is possible that the gap is quite small. For such an explanation to apply to the zero field continuum observed in our measurements, a phase transition out of the ordered ferromagnetic phase would be required at low temperature ($T < 215$ mK) and very low fields along [001] ($\mu_0 H_{\text{int}} < 0.036$ T). All of the inelastic neutron scattering measurements were collected in applied fields equivalent to an internal field greater than this, so it is possible that the zero field data may require a fundamentally different method of computing the excitation spectrum from the finite field measurements. This makes some sense, as only at zero field do all of the sharp excitation modes in the inelastic spectrum disappear.

6.7 Comparison of Heat Capacity and Neutron Scattering Energy Gaps

The heat capacity measurements of Ch. 3 in small applied magnetic fields can be used to shed further light on the low field excitation spectrum. Figure 6.16 shows the notional energy gap extracted from fitting the Schottky-like anomaly that arises in the heat capacity of $\text{Yb}_2\text{Ti}_2\text{O}_7$ in small applied magnetic fields along [001], compared to the lower bound of the magnetic inelastic neutron scattering spectrum, and the energy of the lowest energy spin wave mode in linear spin wave calculations. In Fig. 6.16 it can be seen that the gap extracted from the heat capacity is much lower than lower bound of the inelastic neutron scattering spectrum. This is consistent with the scattering at low fields being a continuum of scattering resulting from the incident neutron creating an excitation that immediately decays into other lower energy excitations. In inelastic neutron scattering measurements, the incident neutron excites the system by flipping a spin. If the fundamental excitations are not associated with a spin flip, the spin flip can decay into more fundamental excitations resulting in a

continuum of scattering, similar to the case of a neutron exciting two magnons, discussed in Ch. 5.2.1, and illustrated in Fig. 5.7. If this is the case, then the energy gap obtained from the heat capacity may represent the fundamental minimum excitation energy, while the lower bound of the neutron scattering spectrum may represent some multiple of this fundamental energy. It is also notable that as the internal magnetic field of the sample is increased, the agreement between the gap extracted from the heat capacity, and the gap extracted from the inelastic neutron scattering measurements approach each other. This is consistent with the low field excitations evolving into magnons that are well described by LSW theory.

As discussed in the previous section, magnon interactions can lead to a reduction of the energy of the magnon modes. It may be that higher order spin wave theory for the model proposed for $\text{Yb}_2\text{Ti}_2\text{O}_7$ may lead to an energy gap consistent with the gap extracted from heat capacity measurements. If this is the case, it is inconsistent that inelastic neutron scattering shows no evidence of the lower energy magnons. Alternatively, the fundamental excitations at low field may be some more exotic excitations, such as those associated with a Z_2 spin liquid, though if that were the case, it should only apply for fields below those used in the heat capacity measurements as no phase transition is observed in any of the measured fields. If such a phase transition as a function of field did occur, the cross over behaviour exhibited by the energy gap, where the heat capacity excitation energy agrees better and better with the neutron and LSW energy gap in increasing field would not be expected, as the excitation in finite field phase should already reflect the properties of the high field phase.

6.8 Curie-Weiss Temperature

The Curie-Weiss Temperature (θ_{CW}) of a magnetic material is a reflection of the overall scale of all of the magnetic interactions present between the moments in the system. Chapter 1 discussed some of the early works on $\text{Yb}_2\text{Ti}_2\text{O}_7$ [5, 9, 14] where this quantity was used in an attempt to quantify the magnetic interactions in $\text{Yb}_2\text{Ti}_2\text{O}_7$,




Figure 6.16: A comparison of the gap to the lowest energy excitation extracted from inelastic neutron scattering for fields of ≤ 3 T $\parallel$ [001] to the energy gap extracted from heat capacity measurements for fields of ≤ 0.5 T $\parallel$ [001] and linear spin wave calculations of the minimum excitation gap. The gap extracted from the heat capacity measurements is similar to the neutron scattering gap around 0.5 T, but trends to zero in decreasing field, a very different behaviour from the neutron scattering and spin wave gaps.

finding them to be net ferromagnetic in nature. Many different works [7, 14, 15, 17, 18, 33, 36, 40] have extracted values for this parameter, all of which fall within the range $\theta_{CW} = 0.4 - 0.8$ K. Ross et al. [27] showed that θ_{CW} can be computed for the effective nearest-neighbour exchange Hamiltonian as

$$\theta_{CW} = \frac{-1}{2k_B (2g_{\perp}^2 + g_{\parallel}^2)} \left(2g_{\parallel}^2 J_{zz} - 4g_{\perp}^2 (J_{\pm} + 2J_{\pm\pm}) - 8\sqrt{2}g_{\perp}g_{\parallel}J_{z\pm} \right) \quad (6.42)$$

Using Eqn. (6.42) to compute θ_{CW} for the optimised exchange parameters determined in this work yields $\theta_{CW} = 780$ mK (equivalent to 0.067 meV), within the range of values determined experimentally from susceptibility data.

6.9 Mean-Field Phase Diagram

The semi-classical ground state associated with the optimised exchange parameters in Table 6.2 is ferromagnetic in nature, with an ordering wave vector of $\mathbf{Q} = 0$. The ordered magnetic moments are not collinear, being splayed towards the local $[111]$ directions.




Figure 6.17: Phase diagram in the space of exchange energies $J_e = \{J_{zz}, J_{\pm}, J_{\pm\pm}, J_{z\pm}\}$ computed using zero temperature mean-field theory plus leading order corrections to $1/S$, which are necessary to lift classical degeneracies in the Ψ_2 and Ψ_3 regions of the phase diagram. The origin of the phase diagram is $J_e = (0.458, 0.229, 0.114, -0.1618)$, a Coulomb phase at the classical level. Four long-ranged ordered phases appear in the vicinity of the Coulomb phase: a splayed ferromagnet with a net ferromagnetic moment spontaneously chosen along one of the cubic directions, the Ψ_4 or Palmer-Chalker antiferromagnetic state, and two related ground states Ψ_2 and Ψ_3 . These states are order by disorder states, arising from a degenerate Γ_5 manifold when quantum fluctuations are taken into account. All these phases have ordering wave-vector $\mathbf{Q} = 0$, and can be described by the spin orientations on a single tetrahedron, as shown to the right of the the phase diagram. The two points marked on the phase diagram correspond to the exchange parameters found by Ross *et al.* [27] (blue circle) and the optimised parameters found in this chapter (red square). The axes $X_1 = \{-0.914, -0.394, -0.092, 0\}$ and $X_2 = \{0.354, -0.702, -0.516, 0.340\}$, defined in terms of J_e , are chosen so that the phase diagram includes both of these points.

Figure 6.17 shows a semi-classical phase diagram in the vicinity of the optimised exchange parameters determined in this chapter, and those of Ross *et al.* [27]. This phase diagram was generated by choosing a highly symmetric point in exchange

coupling space $J_e = \{J_{zz} = 1, J_{\pm} = 0.5, J_{\pm\pm} = 0.25, J_{z\pm} = -1/\sqrt{8}\}$ that supports a Coulomb phase at the classical level [95]. This Coulomb phase is similar to the Coulomb phase in the spin ice materials, where the ground state has extensive degeneracy, but where each configuration of the system must satisfy a local constraint: in the absence of excitations, the spins on a single tetrahedron of the pyrochlore lattice must have a “2-in-2-out” configuration [96], which gives rise to “pinch point” features in the quasi-elastic neutron scattering [97]. The local constraint on the Coulomb phase of Benton et al. [95] is not as easily expressed as that of the spin ice materials, but it gives rise to “pinch lines” in the quasi-elastic neutron scattering. These couplings are then rescaled so that $J_{z\pm}$ matches the value in the optimised exchange parameters. Using this rescaled point, the optimised exchange energies, and the exchange energies of Ross et al. [27], a plane in exchange energy space is defined, and the phase diagram was computed in this plane, as shown in Fig. 6.17.

Within the plane of Fig. 6.17 it can be seen that the optimised exchange energies lie in the splayed ferromagnetic phase, along with the exchange energies of Ross et al. The principal difference between these two sets of the exchange parameters is the strength of the J_{zz} interaction, which is much smaller in the optimised exchange parameters compared to the exchange parameters of Ross et al. [27]. The effect of this reduced J_{zz} in the phase diagram is that the optimised exchange energies lie closer to the phase boundary between the ferromagnetically ordered state and the Γ_5 manifold of degenerate states. The much reduced J_{zz} means that Quantum Spin Ice models based on perturbing away from a Hamiltonian where this term is dominant [54, 56], will not be applicable for a model based on the optimised exchange parameters.

The close proximity of the phase boundary between the splayed FM state and the Γ_5 manifold in the phase diagram may provide an explanation for the unusual excitation spectrum observed in zero magnetic field. Refs. [98, 99] found for the Antiferromagnetic Kagome model with further neighbour interactions, degeneracy of two competing ground states along a phase boundary provides the necessary conditions for a chiral spin liquid to emerge as the new ground state of the system in the region

of the phase boundary. A similar phenomenon at the phase boundary between the splayed FM state and the Γ_5 manifold is not completely without foundation, as both the Kagome lattice and the pyrochlore lattice are made up of corner sharing triangles, and the nearest-neighbour exchange Hamiltonian contains chiral terms.




Figure 6.18: Magnetic field dependence of the effective spin-1/2 (top row) and projected full moment (bottom row) directions on a single tetrahedron in the splayed ferromagnetic phase for a magnetic field of 0 – 9 T along the [001] direction. The red arrow indicates the direction of the applied magnetic field. In zero field, the spins point between the local [111] direction for each sublattice site (see Tab. 6.1) and the global [001] direction, such that the tetrahedron has a net moment along the [001] direction. As the field increases, the spins tilt toward the [001] direction, eventually tilting past the [001] direction towards the local XY plane, due to the presence of the anisotropic $\mathbf{g}$ tensor in the Zeeman term of the Hamiltonian. Throughout this process the spins remain in the plane defined by the local [111] and [001] directions. The full magnetic moments include the $\mathbf{g}$ tensor anisotropy, so even in zero field, they are tilted in the local [111]-[001] plane towards the local XY direction. As the applied magnetic field increases, the moments tilt further toward the local XY plane. No new symmetry is broken in increasing field, with the net moment of the single tetrahedron always pointing along the [001] direction, and the effective spins and magnetic moments rotating continuously.

Figure 6.18 shows the evolution of the splayed ferromagnetic phase as a function of applied magnetic field from 0 – 9 T $\parallel$ [001]. In zero field, the effective spins are aligned between the local [111] direction for each sublattice site (see Tab. 6.1) and the [001] direction, such that the tetrahedron has a net moment along the [001] direction. Upon applying an increasing magnetic field along the [001] the effective spins rotate continuously in a plane defined by the local [111] direction for each sublattice site and the [001] direction. In all applied fields the net moment on the single tetrahedron

is along the $[001]$ direction. For high enough applied field, all of the spins point along the $[001]$ direction, after which they rotate “past” the $[001]$ direction, towards the local XY plane due to the effect of the anisotropic $\mathbf{g}$ tensor. As the spins rotate continuously as a function of the applied field, and the net moment of the tetrahedron does not change direction, the symmetry of the system does not change and there is no phase transition. The projected full magnetic moments are simply the effective spins multiplied by the anisotropic $\mathbf{g}$ tensor. They start out in the local $[111]$ - $[001]$ plane, aligned between the $[001]$ direction and the local XY plane, forming an angle with the $[001]$ direction of 1.4° away from the local $[111]$ direction, again so that the net moment of the single tetrahedron is along $[001]$. In increasing applied magnetic field the magnetic moments rotate further towards the local XY plane.

6.10 Conclusions

In this chapter, mean-field methods and linear spin wave theory were used to compute the spin wave spectrum and magnetisation of $\text{Yb}_2\text{Ti}_2\text{O}_7$ respectively, based on a most general symmetry allowed effective spin- $1/2$ nearest-neighbour exchange model with six parameters. These calculations were compared to magnon dispersion relations extracted from inelastic neutron data collected in applied magnetic fields of 5 and 9 T along the $[001]$ direction, as well as single crystal magnetisation in an applied field of 7 T $\parallel [001]$, and magnon dispersion points for an applied field of 5 T $\parallel [1\bar{1}0]$, extracted from the work of Ross et al. [27]. The fitted Hamiltonian parameters included the four symmetry allowed exchange terms $J_e = \{J_{zz}, J_{\pm}, J_{\pm\pm}, J_{z\pm}\}$ and the two terms that make up the anisotropic $\mathbf{g}$ tensor, $g_{\parallel}$ and $g_{\perp}$. By including all of these distinct data sets, a global minimum in χ^2 was found, associated with an optimal set of exchange parameters and $\mathbf{g}$ tensor components. These optimal parameters differ from the earlier parameterisation based on B $\parallel [1\bar{1}0]$ data and fixed $g_{\parallel}/g_{\perp}$ ratio [27] in that the J_{zz} exchange is now found to be a sub-leading term compared to $J_{z\pm}$. In Ross et al. [27] both J_{zz} and $J_{z\pm}$ were found to be the dominant exchanges,

giving rise to a picture of $\text{Yb}_2\text{Ti}_2\text{O}_7$ as a “Quantum Spin Ice” due to significant local Ising exchange. The optimised parameters determined here find a much smaller value for J_{zz} , leaving $J_{z\pm}$ as the dominant exchange term, meaning that Quantum Spin Ice models [54, 56] based on dominant J_{zz} will not be applicable to this Hamiltonian. The optimised exchange Hamiltonian places the model very close to a phase boundary in the semi-classical phase diagram in exchange energy space. Spin wave theory calculations using the optimised exchange Hamiltonian are shown to fit the intensities and dispersions of the single magnon modes at high fields; and the energy bounds and intensity modulation in the $(h, k, 0)$ plane of the two magnon continuum in large applied magnetic fields quite well, but underestimate the weight of the continuum scattering even at the highest fields measured. Calculations of the neutron scattering sum rule and the mean-field magnetisation suggests that these methods over estimate the ordered magnetic moment in zero field. As the first order terms giving rise to the two magnon continuum are longitudinal fluctuations in the effective moment, this overestimation of the ordered moment may be responsible for the underestimation of the overall weight of the two magnon continuum. In small applied fields ($B \leq 3$ T) it is found that linear spin wave theory breaks down as the single magnon modes overlap in energy with the two magnon continuum, and the high-energy magnon peaks broaden and become part of the continuum.

Chapter 7

Conclusions and Future Work

7.1 Conclusions

The objective of this thesis has been to shed further light on the nature of the low temperature magnetic phase in the rare earth pyrochlore material $\text{Yb}_2\text{Ti}_2\text{O}_7$. This material has been of interest to the condensed matter physics community for almost 50 years, ever since a phase transition was observed in heat capacity measurements on a powders sample at $T_c \approx 214$ mK. $\text{Yb}_2\text{Ti}_2\text{O}_7$ has recently attracted much more attention, due to the potential that the unusual physics of the low temperature phase in this material ($T \lesssim 250$ mK) may be related to proximity to a quantum spin ice phase, a quantum paramagnetic phase with dispersing magnetic monopole excitations, photon-like excitations, and dual electric monopoles [3], or another strongly fluctuating cooperative magnetic state.

Because of previously reported variation in magnetic properties between different samples, one of the key aspects of this work in this thesis has been to confirm high quality of the single crystal samples used in the experiments via heat capacity measurements. The zero field heat capacity of the single crystal sample of $\text{Yb}_2\text{Ti}_2\text{O}_7$ used in this work shows a single sharp peak at $T_c = 214(2)$ mK in zero field. This peak is even sharper and occurs at even higher temperatures than observed by Chang et al. [33] for their best single crystal with $T_c = 195$ mK. The sharpness of the transition

and higher value of T_c is generally accepted as a clear indication of high sample quality with a fully coherent, and representative of the limit of minimal structural disorder. Our specific heat characterisation shows that the single crystal of $\text{Yb}_2\text{Ti}_2\text{O}_7$ used for the heat capacity and inelastic neutron scattering measurements in this work is of very high quality.

The first set of new results reported in this thesis concern the phase diagram in magnetic field along a cubic axis, mapped using heat capacity measurement. The data showed that applying a very small magnetic field, equivalent to an internal field in the sample of $\mu_0 H_{\text{int}} \sim 0.018$ T rapidly suppresses the phase transition into the low temperature phase. The magnetic field has no significant effect on the temperature of the transition, only reducing the size of the peak observed in the heat capacity. This behaviour is consistent with the transition being to a phase with a net moment along [001], as the magnetic field polarises the high temperature paramagnetic phase, turning the phase transition into a crossover as the polarisation of the paramagnetic phase removes the need for spontaneous polarisation via symmetry breaking. Magnetisation [37] and polarised neutron diffraction [33] measurements on other high quality samples with evidence of a sharp peak in the specific heat have provided evidence that this phase transition is of first order, to a phase with ferromagnetic Bragg peaks. As the transition to the low temperature phase with spontaneous ferromagnetic polarisation is first order with a finite jump in the magnetization at T_c , in a small applied field there can still be a transition with a small magnetization jump, but a crossover is expected for sufficiently large fields. The in field heat capacity measurements suggest that the transition does not survive above an internal field of $\mu_0 H_{\text{int}} = 0.018$ T, a very small field indeed.

In addition to heat capacity measurements, the primary result of this work is the characterisation of the magnetic excitations in applied magnetic fields of up to 9 T along [001], at $T \sim 145$ mK, well within the low temperature phase of the crystal ($T_c = 214(2)$ mK). Previous studies [27] of the excitations in this phase have found sharp magnon modes in large applied magnetic fields ($B = 5$ T $\parallel$ $[\bar{1}\bar{1}0]$), while in

zero field the inelastic neutron scattering showed a continuum of scattering, but it was left unclear how the continuum scattering evolves into the sharp magnon modes upon increasing applied field.

In relatively large applied magnetic fields ($B \geq 3$ T) the magnetic excitation spectrum was found to take the form of four sharp dispersive magnon modes, along with a weak band of continuum scattering at twice the energy of the magnon modes. Based on the bandwidth ratio between the magnon dispersions and the high energy continuum scattering, and the rate of change of the energy bounds in varying magnetic field, the continuum scattering has been identified as being due to neutron scattering via creating two magnons. This is the first direct observation of two magnon processes in $\text{Yb}_2\text{Ti}_2\text{O}_7$ via inelastic neutron scattering, though high energy excitations associated with two magnon states have been suggested in THz measurements [38].

For applied magnetic fields of $B < 3$ T, the highest energy magnon dispersion overlaps with the high energy continuum, and the sharp mode then broadens and becomes part of the continuum, likely due to magnon decay processes. The three remaining lower-energy magnon modes are observed to have their dispersions suppressed as the applied magnetic field is decreased, and appear to collapse into a single, almost dispersion-less mode for fields of $0.21 \leq B \leq 0.75$ T. In zero field, no sharp modes are observed in the inelastic magnetic scattering, only a very broad continuum extending up to ~ 1.5 meV, with no observable gap even at the highest resolutions used in the experiments. Applying a very small magnetic field, equivalent to an internal field of $\mu_0 H_{\text{int}} = 0.17$ T is sufficient to open a finite gap to the lowest energy excitation. The absence of sharp modes in zero field is indicative of strong zero-point quantum fluctuations in the ground state which prevent coherent propagation of single-spin-flip magnons down to the lowest probed energies (0.03 meV). Instead spin flip excitations created in a neutron scattering process decay into other types of excitations. The very broad extent of the continuum, almost an order of magnitude larger than the scale of the largest exchange interaction suggesting that the corresponding excitations involve multiple spin flips, which could occur in a correlated multi-spin clusters, or be due to

coherently propagating quasi-particles, such as the monopole excitations in quantum spin ice states or other related states [31, 32, 54]. The lack of any measurable gap in the zero field excitation spectrum is surprising, given that the mean field ground state has cubic broken symmetry, with a net magnetic moment along one of the $\langle 001 \rangle$ cubic directions of the crystal and linear spin wave theory predicts a gap of 0.2 meV. Clearly the data shows excitations well below this threshold, showing that spin flips do not propagate coherently at these energies, but decay into other more elementary quasi-particles.

At high fields, where the dominant excitations are well defined sharp magnons, mean-field and linear spin wave theory have been used to compute the spin wave spectrum based on the most general symmetry allowed nearest-neighbour exchange model for $\text{Yb}_2\text{Ti}_2\text{O}_7$, with six free parameters (4 exchange and 2 $\mathbf{g}$ tensor components). These calculations were compared to magnon dispersion relations extracted from inelastic neutron data collected in applied magnetic fields of 5 and 9 T along the $[001]$ direction, as well as single crystal magnetisation in an applied field of 7 T $\parallel [001]$, and already published magnon dispersion points for an applied field of 5 T $\parallel [1\bar{1}0]$, extracted from the work of Ross et al. [27]. By performing a comprehensive search in parameter space with no constraints and comparing the computed spin wave spectrum to the experimental results, the strengths of the four symmetry allowed exchange terms $J_e = \{J_{zz}, J_{\pm}, J_{\pm\pm}, J_{z\pm}\}$, and the two terms that make up the anisotropic $\mathbf{g}$ tensor, $g_{\parallel}$ and $g_{\perp}$, which describes the anisotropic nature of the ground state crystal field doublet, were extracted. This re-evaluation of the Hamiltonian was prompted by the fact that the previous proposals for the magnetic interactions in $\text{Yb}_2\text{Ti}_2\text{O}_7$, by Ross et al. [27] (with $J_{zz} = 0.17$ meV one of the largest exchange terms, along with $J_{z\pm} = -0.14$ meV), could not reproduce features observed in the experimental magnon dispersions when for fields along $[001]$. The optimal Hamiltonian found in the extensive fitting, which explains all available experimental data so far has $J_{zz} = 0.034(2)$ meV which is approximately one quarter of the $J_{z\pm} = -0.162(1)$ meV value. This puts into question treatments of $\text{Yb}_2\text{Ti}_2\text{O}_7$ as a quantum spin ice as

this term is expected to drive quantum fluctuations, and the local Ising exchange expected to enforce the two-in two-out spin ice rule is no longer dominant. Instead, the re-evaluated exchange Hamiltonian places the model very close to a phase boundary in the semi-classical phase diagram in exchange energy space between two distinct magnetic phases [40]. It is interesting to speculate if the proximity to this phase boundary may be responsible for the strongly fluctuating nature of the the ground state, and the complete failure of mean-field theory to capture the spin dynamics in zero field.

Linear spin wave theory calculations of the high field magnon modes using the optimised exchange Hamiltonian fit the intensities and dispersions of the single magnon modes at high fields very well. The energy bounds and intensity modulation in the $(h, k, 0)$ plane of the two magnon continuum in large applied magnetic field are also well captured, but the weight of the continuum scattering even at the highest fields measured is underestimated by the linear spin wave theory approach. Calculations of the neutron scattering sum rule and the mean-field magnetisation suggests that these methods overestimate the ordered magnetic moment in zero field. As the first order terms giving rise to the two magnon continuum are longitudinal fluctuations in the effective moment, this overestimation of the ordered moment may be responsible for the underestimation of the overall weight of the two magnon continuum.

In terms of describing the one magnon states, it was found that linear spin wave theory breaks down at intermediate fields ($B = 3$ T) as the highest energy single magnon mode overlaps in energy with the high energy continuum scattering, and the high energy magnon mode broadens and becomes part of the continuum. Linear spin wave theory also predicts the existence of a finite energy gap to the lowest energy magnon at zero field, in stark disagreement with the inelastic neutron scattering data. In fact, linear spin wave theory breaks down completely in zero field, unable to capture the key observation that magnons do not propagate at any of the energy scales probed.

In small applied magnetic fields, $B \leq 0.5$ T $\parallel$ [001], the heat capacity measure-

ments provided some additional insight into the nature of the excitations that give rise to the observed continuum in the inelastic neutron scattering. These measurements show a broad hump, whose peak in temperature increases upon increasing applied magnetic field. This hump was fit using a simplified two-level Schottky anomaly, in order to extract a notional energy gap to the lowest energy excitation. The extracted gap increases in applied magnetic field, but is significantly smaller than the gap observed in the inelastic neutron scattering at the same fields. This smaller gap extracted from the heat capacity data may represent the gap to the lowest energy fundamental excitation that gives rise to the multi-particle continuum of scattering observed in zero field.

All of these results combine to give a picture of the low temperature magnetism in $\text{Yb}_2\text{Ti}_2\text{O}_7$ where in high magnetic fields, the physics is well captured by a mean-field approach where magnetic moments are largely polarised by the applied field and the excitations are coherently propagating magnons, but where longitudinal fluctuations in the magnetic moments still give rise to two magnon scattering processes, observable at higher energies.

At finite but intermediate fields ($B \approx 0.75$ T) magnon interactions and decay processes destroy the ability of the magnons to coherently propagate, leaving only a continuum of scattering above a sharp dispersionless mode. Already in this intermediate field range linear spin wave theory fails to explain the nature of the excitations, which may be strongly interacting magnons, or may be some more exotic excitation. In zero field, the excitations are essentially gapless and give rise to only a broad continuum of scattering in the inelastic neutron scattering, indicating quantum fluctuations dominate the system, and spin flips decay into some more fundamental excitations.

Specific heat measurements showed a rapid suppression of the transition between the high temperature paramagnetic phase and the low temperature phase with ferromagnetic order, and there was no evidence of a transition between zero field and the high field phase, suggesting that the change in the ground state between these regimes is consistent with a crossover (at least above $\mu_0 H_{\text{int}} = 0.018$ T).

7.2 Future Work

The disagreement between linear spin wave theory calculations of the overall weight of the two magnon scattering in $\text{Yb}_2\text{Ti}_2\text{O}_7$ at all probed fields suggests that longitudinal fluctuations of the magnetic moments may be more significant than can be captured by linear spin wave theory. At low field, fluctuations of the magnetic moments are expected to be the dominant physics, as linear spin wave theory completely fails to explain the inelastic neutron scattering spectrum. The failure of linear spin wave theory could be made more quantitative by a direct measurement of the full ordered moment $\langle J \rangle$ in near zero field (high enough field to put the sample into a single ferromagnetic domain). This would require polarised neutron diffraction in applied field to measure both ferromagnetic and antiferromagnetic peaks (e.g. $\mathbf{Q} = (2, 0, 0)$) to extract the full ordered moment value and see how close the ground state is to the mean-field prediction.

Additionally, further insight into the change in the nature of the magnetic excitations in applied fields of less than 3 T could be gained from performing polarised inelastic neutron scattering measurements. These experiments could provide insight into what portion of the low field continuum is made up of longitudinal fluctuations (as is primarily the case for the two magnon continuum) versus transverse fluctuations (the fluctuations that give rise to magnons in linear spin wave theory). Experiments of this type have been reported by Robert et al. [40] in zero field at paramagnetic temperatures ($T \gtrsim 2.0$ K), showing some differentiation between the polarisation channels. Measurements at base temperature and small fields may provide additional insight, but will be complicated by the non-collinear quantisation directions of the four sublattice sites in the material which will lead to mixing of the longitudinal and transverse fluctuations in the polarisation channels.

The in field behaviour of the heat capacity of $\text{Yb}_2\text{Ti}_2\text{O}_7$ for fields along [001] also suggests measurements that would be complementary to this work. The dilution fridge used to perform the heat capacity measurements in this work limited the easily

accessible temperature range for the measurements to $T < 700$ mK. This limited temperature range meant that for $B > 0.5$ T, the hump in the heat capacity arising from the excitations in the system was pushed to temperature well above 700 mK, so that heat capacity measurements in greater fields did not yield any information about the excitations. If a different cooling method, that allowed access to higher temperatures, were to be used, then the hump in the heat capacity could be measured for applied fields greater than 0.5 T. This measurement would be of interest as it was shown in Ch. 6 that linear spin wave calculations come close to capturing the low temperature part of the hump in the heat capacity at $\mu_0 H_{\text{int}} = 0.39$ T, and the expectation is that at high enough fields ($B > 3$ T), when magnon dispersions are well captured by linear spin wave theory, the spin wave description may be adequate to describe the heat capacity as well. High field high temperature measurements of this feature would allow for the confirmation of this trend.

These three different sets of measurements could help to better understand the evolution of the magnetic excitations in $\text{Yb}_2\text{Ti}_2\text{O}_7$ upon increasing the applied magnetic field, and provide insight into the relationship between the zero field magnetic phase and the high field polarised phase, which appear to have very different associated excitations, despite there being no evidence of a phase transition as a function of applied magnetic field.

APPENDICES

Appendix A

RuO₂ Thermometer Calibration

This appendix describes the methods used to measure and parametrise the relationship between the resistance of the RuO₂ thermometer used to perform the heat capacity experiments in this work and its temperature. This calibration is performed in both zero and finite applied magnetic field.

A.1 Zero Field Calibration

In order to convert the thermometer resistances measured during heat capacity experiments into temperatures, it is necessary to measure and parametrise the relationship between the temperature of the thermometer and its resistance. In zero applied magnetic field, this was accomplished by thermally anchoring a calibrated Lakeshore RX-102B-CB RuO₂ thermometer, and an uncalibrated RX-102A-BR RuO₂ thermometer directly to the mixing chamber of the dilution fridge. The fridge was then cooled to base temperature and allowed to warm slowly to high temperatures multiple times. The temperature of the calibrated RX-102B-CB was measured using a Lakeshore 370 AC resistance bridge, while the resistance of the uncalibrated thermometer was measured using a Stanford Research SR830 lock-in amplifier connected to an optically isolated AC current source and amplifier. The resulting temperature-resistance curve is shown in the main plot of Fig. A.1(a).

The standard method of parametrising such a curve is to use a series of Chebyshev polynomials [100] to fit the temperature as a function of resistance. Chebyshev polynomials are defined as

$$t_n = \cos(n \arccos(X)), \quad (\text{A.1})$$

where n is an integer. Using these functions, the temperature as a function of the parameter X is given by [100]

$$T(X) = \sum_n a_n t_n(X), \quad (\text{A.2})$$

where a_n (starting at $n = 0$) are scalar coefficients that define the fit. The parameter X is related to the measured resistance by [100]

$$X = \frac{(k - k_L) - (k_U - k)}{(k_U - k_L)}, \quad (\text{A.3})$$

where $k = \log(R)$, and k_L and k_U are the lower and upper bounds of range of k to be fit. It is necessary to construct this variable from the resistance, as Chebyshev polynomials are only normalised to have magnitude ≤ 1 in the interval $[-1, 1]$. Figure A.1(a) shows that the resistance vs. temperature curve from very low to very high temperature has been fit by four different series of Chebyshev polynomials (panels i-iv). This was necessary to ensure a good fit to the data without requiring an excessive number of terms in Eqn. (A.2). These ranges overlap, which was necessary as the polynomial fits to the data are poor near the ends of the intervals. Figure A.1(b) shows the fit to the four temperature ranges as a function of the variable X . These fits are performed using the Numpy Chebyshev fitting package, which returns a series of coefficients a_n , with $n_{\text{max}} = 7$. The uncertainty of the fitted parameters will define the uncertainty in the temperature extracted from resistance measurements in the experiment via bootstrap calculations.




Figure A.1: Temperature as a function of resistance calibration in zero field of the Ruthenium Oxide thermometer used for the heat capacity measurements. (a) shows the resistance temperature data (blue points) used to parametrise the relationship between the resistance of the thermometer and its temperature, along with the four part Chebyshev polynomial fit to the data. The insets (i-iv) show close-ups of the four ranges fit with Chebyshev polynomials. (b) shows the same parameterisations as a function of the variable X , defined in Eqn. (A.3).

A.2 Finite Field Calibration

In addition to measurements in zero field, several experiments were performed in finite magnetic field, which has an impact on the resistance temperature relationship of the RuO_2 thermometer [101]. In large applied magnetic fields, the magnetic field dependence of the resistance can be well described as proportional to $B^{1/2}$ [101], but this approximation breaks down at low fields. In order to construct a function that can determine the temperature of the thermometer as a function of resistance in an applied magnetic field, a function of two variables must be constructed.

Other approximations are available, but for ease of use, and the ability to fit measured (T, R, B) curves, a series of two dimensional Chebyshev polynomials was used:

$$T(X, Z) = \sum_{m,n} a_{m,n} t_m(X) t_n(Z), \quad (\text{A.4})$$

where $t_m(X)$ and $t_n(Z)$ are one dimensional Chebyshev polynomials, and $a_{m,n}$ are a series of coefficients defined by fitting a series of (R, B) curves at fixed T . The

experimental (T, R, B) curves, shown in Fig. A.2, were obtained by stabilising the heat capacity apparatus with no attached sample at a fixed temperature and slowly ramping the applied magnetic field to minimise eddy current heating. These sweeps were combined with the temperature sweep data in Fig. A.1(a). The definition of X is the same as in Eqn. (A.3), while Z is defined as

$$Z = \frac{(B - B_{\min}) - (B_{\max} - B)}{(B_{\max} - B_{\min})}, \quad (\text{A.5})$$

where B is the applied magnetic field amplitude and $B_{\min}$ and $B_{\max}$ are the minimum and maximum magnetic fields used (0 and ~ 1.5 T respectively).

The results of fitting Eqn. (A.4) with $m_{\max} = 7$ and $n_{\max} = 3$, leading to a 32 parameter fit, are shown in Fig. A.2(a). The gaps in the (T, R, B) curves used in the fit are due to heat pulses that arise in the sample environment when the superconducting magnet was energised. Figure A.2(b) shows the same fit to the calibration data as a function of X and Z . Unlike the zero temperature data, only a single series of Chebyshev polynomials was necessary to parametrise this data, as the temperature range fitted was much smaller ($T < 700$ mK). Once again, the uncertainty of the fitted parameters will define the uncertainty in the temperature extracted from resistance measurements in the experiment via bootstrap calculations. Figure A.2(c) shows a constant field ($B = 0.75$ T) cross section, showing the degree of fit to the calibration data achieved by the 2D Chebyshev polynomial for a temperature sweep in fixed applied field, the type of sweep performed in this thesis. Figure A.2(d) shows a constant temperature cross section ($T = 85$ mK) through the calibration data and the 2D chebyshev polynomial for the same temperature. The fit to the data is not perfect, due to the use of only Chebyshev polynomials of up to $n = 3$ for the field dependence. This was necessary due to the previously discussed gaps in the data; adding higher order terms lead to unphysical oscillations in these empty regions. The level of fit achieved decreases the field induced error in the thermometer from 12.6% (~ 11 mK) to 1.1% (< 1 mK) at $B = 1.3$ T, a significant improvement.




Figure A.2: Temperature as a function of resistance and applied magnetic field of the Ruthenium Oxide thermometer. (a) The experimental data points describing the relationship between thermometer resistance, temperature, and magnetic field, along with a two dimensional Chebyshev polynomial fit to the data. (b) The same data points and fit in terms of the reduced variables X and Z . (c) Constant field cross section through (a) showing the quality of the fit to the resistance as a function of temperature in finite applied magnetic field. (d) Constant temperature cross section through the data in (a) showing the fit to the resistance as a function of applied magnetic field. The green line illustrates how poor the fit to the data would be if the resistance of the the thermometer was constant in applied magnetic field.

These two sets of calibrations were used to generate all of the temperature values used in the heat capacity measurements in this work.

Appendix B

Sample Growth

The single crystal samples of $\text{Yb}_2\text{Ti}_2\text{O}_7$ used in this work were grown using a combinations of the methods of Chang et al. [33] and Prabhakaran and Boothroyd [102] by Dr. Dharmalingam Prabhakaran at the University of Oxford. At the time of the growth of the crystals used in this work, Chang et al. [33] was the only work to present evidence of a single crystal sample of $\text{Yb}_2\text{Ti}_2\text{O}_7$ with a single sharp peak in the heat capacity, evidence interpreted to mean that it was the best single crystal grown to date. The method used to grow single crystals by Chang et al. [33] was to first combine stoichiometric quantities of Yb_2O_3 and TiO_2 , which were mixed, pressed into rods and sintered at a temperature of 1150 °C. A rod produced this way was then placed in an optical floating zone furnace in air, and grown into a single crystal at a rate of 1.5 mm/hr.

In the case of the sample used in this work, high purity ($> 99.99\%$) Yb_2O_3 and TiO_2 were combined using the standard solid state reaction technique in air. A mixed powder of these two chemicals was sintered at temperatures between 1200°C and 1300°C for 48h, with intermediate grinding, following the method Prabhakaran and Boothroyd [102]. This single phase powder was then compressed in the form of cylindrical rod of diameter 10mm and length 140mm and sintered at 1250C for 6h.

Using this rod, a single crystal was grown in a four-mirror optical floating-zone furnace (Crystal System Inc.) in argon rich atmosphere with a growth rate of 1-

2mm/h. Due to the argon atmosphere, the as-grown crystal was oxygen deficient and dark in colour. In order to improve the oxygen stoichiometry, the as-grown crystal was annealed at 1200°C for 72h under oxygen flow atmosphere and the crystal become transparent as shown in Fig. 5.1.

The crystal produced in the growth had an average diameter of 7.33(1) mm, and was transparent, with a yellow tint. The portion of the crystal used for inelastic neutron scattering measurements had a mass of 6.3 g. The crystal was confirmed to be a single crystal by neutron diffraction. Fig. B.1(a) shows the $(4, 0, 0)$ structural Bragg peak in the $(4, k, l)$ plane of reciprocal space. The fact that there is only a single peak indicates that the sample is a single crystal. Figure B.1(b) shows a scan along the vertical $(4, 0, l)$ direction, showing a cross section of the peak perpendicular to the $(h, 0, 0)$ direction. The vertical direction was used to study the cross section of the peak due to the increased instrumental resolution in that direction. From the fitted width of the scan, an effective upper bound on the mosaicity of the crystal is estimated to be 0.6° (FWHM). This value incorporates the actual mosaicity of the sample, along with the neutron beam divergence of the instrument, and other instrumental effects, so that actual mosaicity of the sample will be lower than the combined effect. The lattice parameters used in Fig. B.1, and in the rest of the thesis are $a = b = c = 10.023\text{\AA}$.




Figure B.1: (a) Slice showing the $(4,0,0)$ Bragg peak in the $(4,k,l)$ diffraction plane. This slice was obtained by integrating white-beam diffraction data collected on the neutron instrument LET at $T \sim 300$ K over the range $3.8 < h < 4.2$. (b) A scan along the $(4,0,l)$ direction through the Bragg peak, showing a fitted Gaussian cross section of the peak. The scan was generated by integrating over $3.8 < h < 4.2$ and $-0.1 < k < 0.1$ in reciprocal space.

Appendix C

Laue Alignment

All of the samples used in this work were aligned using a back-scattering Laue diffractometer. In this instrument the crystal is placed in a white beam of x-rays (x-rays with a continuum of wavelengths), and an x-ray camera is placed between the sample and the x-ray source (with a path through the camera for the x-rays).

The beam of x-rays is diffracted by the crystal, leading to a series of Bragg peaks on the camera when the Laue equations [103]:

$$\mathbf{a} \cdot \mathbf{Q} = 2\pi h, \tag{C.1}$$

$$\mathbf{b} \cdot \mathbf{Q} = 2\pi k, \tag{C.2}$$

$$\mathbf{c} \cdot \mathbf{Q} = 2\pi l, \tag{C.3}$$

are satisfied. $\mathbf{Q} = \mathbf{k}_i - \mathbf{k}_f$ is the scattering wave vector, $\mathbf{a}, \mathbf{b}, \mathbf{c}$ are primitive lattice vectors, and h, k, l are reciprocal lattice indices. The use of a white beam of x-rays allows Eqns. (C.1)-(C.3) to be satisfied simultaneously for multiple values of h, k, l because $\mathbf{k}_i$, does not have a fixed length in reciprocal space, due to the presence of multiple wavelengths of x-rays. The angle between $\mathbf{k}_i$ and $\mathbf{k}_f$ can vary, allowing $\mathbf{Q}$ to change in length and direction, sampling a large volume of reciprocal space and producing many structural Bragg peaks in the x-ray camera.

Figure C.1 is a sample Laue camera image of the Bragg peaks generated when

$\mathbf{k}_i$ is parallel to the $[001]$ axis of the crystal and the $[100]$ axis is horizontal in the figure. By matching the pattern of spots in this and similar images to calculations, the orientation of the sample can be determined and refined. The software used in this work to simulate the pattern of Bragg peaks is QLaue [104].




Figure C.1: Laue snapshot of a crystal of $\text{Yb}_2\text{Ti}_2\text{O}_7$ with the $[001]$ axis is pointing out of the camera plane. The large dark spot in the middle of the image is the shadow where the x-ray guide tube passes through the x-ray camera panel. Note the 4-fold symmetry around the out-of-plane direction, as expected for cubic symmetry.

Appendix D

Magnetisation Measurements and Demagnetisation Corrections

D.1 Magnetisation Measurements

D.1.1 Introduction

To provide additional constraints on the least squares fit in Ch. 6, the magnetisation of a small sample of $\text{Yb}_2\text{Ti}_2\text{O}_7$ was measured in a Quantum Designs Magnetic Properties Measurement System (MPMS). This Appendix will provide a very brief summary of how a Superconducting Quantum Interference Device (SQUID) Magnetometer functions, and give results of the magnetisation measurements and mean-field magnetisation calculations. Demagnetisation effects for the samples used in all of the experiments in this work will also be calculated.

D.1.2 The SQUID Magnetometer

The magnetometer in the Quantum Designs MPMS is based on the use of a Radio Frequency (RF) SQUID, a loop of superconducting material containing a single Josephson junction. A Josephson junction is small insulating barrier through which

the Cooper pairs of the superconductor can tunnel. To use this device as a magnetometer, it is inductively coupled to an RF tank circuit, which is driven with an RF current. The resulting RF voltage of the tank circuit is modulated by any magnetic flux passing through the SQUID [105]. This modulation has periodicity of one quantum of flux $\phi_0 = \frac{h}{2e} = 2.067 \times 10^{-15}$ Wb. Using amplifiers, it is possible to resolve this modulation to better than one flux quantum, allowing sensitivities to external fields of 10^{-15} T [106].

In the MPMS, the sample is passed through a series of superconducting pickup coils inside a superconducting magnet. The moment of the sample passing through the coils generates a time varying magnetic field, which generates a current in the pickup coils. This current is then passed through a second set of coils inductively coupled to the SQUID, which is isolated from the sample and applied magnetic field [106]. The response of the SQUID to the inductive field is proportional to the sample magnetisation.

D.1.3 Sample

The magnetisation measurements were performed on a single crystal sample of $\text{Yb}_2\text{Ti}_2\text{O}_7$. The sample mass was 28.2(1) mg and it was cut from the same sample of $\text{Yb}_2\text{Ti}_2\text{O}_7$ that was used the neutron scattering and heat capacity experiments. The geometry of the sample was a flat plate with one end tapering towards the edge, measuring 2.09(1) mm by 1.81(1) by 1.01(1) mm thick. The magnetisation was measured with the large face perpendicular to the applied field. This is not an ideal geometry for demagnetisation corrections, but was necessary because as the flat face is known to be perpendicular to the [100] direction of the crystal, ensuring the the magnetic field geometry is the same as that used in all of the other experiments in this work. The sample was mounted in the MPMS by fixing it between two pieces of Teflon, which were placed inside a plastic pill capsule inside a plastic straw, the standard method for mounting samples in the MPMS.

D.1.4 Demagnetisation Correction

It is important to take into account demagnetisation effects when measuring the magnetisation of a sample. These effects arise due to the requirement that $\nabla \cdot \mathbf{B} = 0$ [1]. The magnetic field $\mathbf{B}$ can be made up of two quantities, $\mathbf{H}$, the Auxiliary Field, and $\mathbf{M}$, the Magnetisation, via the relationship $\mathbf{B} = \mu_0 (\mathbf{H} + \mathbf{M})$, μ_0 is the permeability of free space. In a magnetised sample, $\mathbf{M}$ is finite, while outside of the sample it is zero, leading to a divergence in $\mathbf{M}$. In order to ensure the condition that $\nabla \cdot \mathbf{B} = 0$, then it must be true that $\nabla \cdot \mathbf{H} = -\nabla \cdot \mathbf{M}$, leading to a finite divergence of the $\mathbf{H}$. This finite divergence of $\mathbf{H}$ at the surface of the sample acts like a source of counteracting field $\mathbf{H}_d$, the demagnetising field. It acts to reduce the size the field in the sample ($\mathbf{H}_{\text{int}}$) via the relationship:

$$\mathbf{H}_{\text{int}} = \mathbf{H}_{\text{app}} + \mathbf{H}_d, \quad (\text{D.1})$$

where $\mathbf{H}_{\text{app}} = \mathbf{B}/\mu_0$ is the applied field outside the sample. When reporting magnetisation results, it is correct to report the magnetisation as a function of this internal field, $\mathbf{H}_{\text{int}}$, as this is the field the magnetic moments in the sample actually experience.

The size of the demagnetising field is related to the geometry of the sample and its magnetisation. If the sample has finite susceptibility, the field can also vary as a function of position throughout the sample. The simplest case is that of a ferromagnetic ellipsoid, where the demagnetising field is given by [1]

$$\mathbf{H}_d = -\mathbf{N}\mathbf{M}, \quad (\text{D.2})$$

where $\mathbf{N}$ is the demagnetisation tensor. If the external magnetic field is applied along one of the principal axes ($\hat{\mathbf{x}}, \hat{\mathbf{y}}, \hat{\mathbf{z}}$), then this tensor is diagonal, and the resulting demagnetising field becomes:

$$\mathbf{H}_d = -(N_x M_x \hat{\mathbf{x}} + N_y M_y \hat{\mathbf{y}} + N_z M_z \hat{\mathbf{z}}). \quad (\text{D.3})$$

In experiments samples are not elliptical in cross-section. In these cases the demagnetisation can be computed numerically, which can be quite computationally intensive, or certain assumptions can be made to simplify the problem. The assumption made here is that the susceptibility of the sample is zero, which means that the magnetisation of the sample is uniform. We also deal only with the volume average (magnetometric) demagnetisation factor, as all of the quantities of interest in this thesis are measured over the whole volume of the sample.

For the two sample geometries encountered in this work (rectangular prism and cylinder), the volume average demagnetisation factor for $\chi = 0$ has been worked out in the literature. Chen et al. [107] present an analytic form for the volume average demagnetisation factor of a square prism of size $2a \times 2b \times 2c$ with the field applied along the c direction [107]:

$$N_m = \frac{2}{\pi} \arctan \left(\frac{ab}{c\sqrt{a^2 + b^2 + c^2}} \right) + \frac{1}{3\pi abc} (F_2 + F_3) + \frac{1}{2\pi abc} (F_m(a, b, c) + F_m(b, a, c) - F_m(c, a, b) - F_m(c, b, a)). \quad (\text{D.4})$$

Here, the terms F_2 , F_3 , and F_m are defined as [107]:

$$F_2 = a^3 + b^3 - 2c^3 + (a^3 + b^3 - 2c^3)\sqrt{a^2 + b^2 + c^2}, \quad (\text{D.5})$$

$$F_3 = (2c^2 - a^2)\sqrt{a^2 + c^2} + (2c^2 - b^2)\sqrt{b^2 + c^2} - (a^2 + b^2)^{3/2}, \quad (\text{D.6})$$

and

$$F_m(u, v, w) = u^2 v \ln \left(\frac{(u^2 + w^2)(u^2 + 2v^2 + 2v\sqrt{u^2 + v^2})}{u^2(u^2 + 2v^2 + w^2 + 2v\sqrt{u^2 + v^2 + w^2})} \right). \quad (\text{D.7})$$

This demagnetisation factor is relevant to both the magnetisation and heat capacity measurements.

For the magnetisation measurements, the geometry of the sample has been approximated as a rectangular prism for the purposes of demagnetisation correction. Using Eqns. (D.4)-(D.7), the demagnetisation factor for the sample is found to be

$N_m \sim 0.486$. Any effect of the uncertainty of the sample dimensions on this value will be far outweighed by the error introduced by treating the sample as a rectangular prism. This value is substituted into the equation

$$H_d^z = -N_m \frac{16}{(10.023\text{\AA})^3} \mu_B \mu_0 M_s \quad (\text{D.8})$$

to obtain the average demagnetising field. There, N_m is the demagnetisation factor, and M_s is the magnetisation of the sample in units of Bohr magnetons per magnetic ion. The prefactor $16/(10.023\text{\AA})^3$ represents the number of Yb^{3+} magnetic moments in the unit cell of $\text{Yb}_2\text{Ti}_2\text{O}_7$, divided by the volume of the unit cell. The resulting values of H_d^z are used in Eqn. (D.1) to determine the average magnetic field inside the sample.

D.1.5 Results

Figure D.1 shows the results of a magnetisation measurement performed at $T = 2.0$ K, corrected for demagnetisation effects. This plot shows the positive field portion of a hysteresis loop measurement. No hysteresis was observed, which is expected as 2.0 K is well above the transition temperature of $\text{Yb}_2\text{Ti}_2\text{O}_7$. The maximum applied field in the experiment was 7 T, which gave a maximum internal field of 6.86 T, and a magnetic moment at this field of $1.695(4)\mu_B/\text{T}$. This value was used as a constraint in our determination of the exchange energies in Ch. 5.

D.2 Mean-Field Magnetisation for Least Squares Fit

The $\mu_0 H_{\text{int}} = 6.86$ T single crystal magnetisation at $T = 2$ K was used in the least squares minimisation to determine the free parameters in the nearest neighbour exchange Hamiltonian. Figure D.2 shows zero temperature mean-field magnetisation




Figure D.1: Magnetisation measurements on a single crystal sample of $\text{Yb}_2\text{Ti}_2\text{O}_7$ collected at $T = 2.0$ K. The data has been corrected for demagnetisation effects, and the field was applied along the $[100]$ axis of the crystal.

computed using the optimised exchange parameters determined by the least squares fit, along with high (2 K, present data) and low (90 mK, extracted from L’Hotel et al. [37]) temperature single crystal magnetisation. For the highest magnetic fields, $\mu_0 H_{\text{int}} \geq 6$ T, the mean-field calculations reproduces the experimental magnetisation reasonably well, but for lower fields, the mean-field calculation overestimates the sample magnetisation. Some of this overestimation may be due to the neglect of thermal fluctuations in the mean-field calculations, but measurements performed at $T = 90$ mK should be well within the ordered phase of $\text{Yb}_2\text{Ti}_2\text{O}_7$, where these effects should be small, which would imply that the reduced magnetisation seen in experiment is due to zero point quantum fluctuations not accounted for in mean-field theory.

D.3 Demagnetisation Corrections for Neutron Scattering and Heat Capacity Measurements

D.3.1 Introduction

It is also necessary to perform demagnetisation corrections for the heat capacity and neutron scattering data to ensure that quantities extracted from both data sets as




Figure D.2: Zero temperature mean-field magnetisation calculations compared to single crystal magnetisation at 90 mK and 2 K with the field applied along [001]. The 90 mK data comes from L’Hotel et al. [37].

a function of magnetic field are correctly compared. As both the heat capacity and inelastic neutron scattering experiments were carried out at very low temperatures, it is necessary to use magnetisation measurements collected at very low temperatures to perform the demagnetisation correction. Since we did not have the necessary equipment available to perform the necessary measurements, the data reported by L’Hotel et al. [37], as shown in Fig. D.3, was used for this analysis.

D.3.2 Heat Capacity Demagnetisation Correction

The largest applied field used in the heat capacity experiments was “ B ” = 1.5 T, so the data of L’Hotel et al. [37] is sufficient to determine the average internal field experienced by the sample. This was done using the same method as in Section D.1.4. The sample geometry used in the heat capacity experiment was a thin flat square plate with dimensions 2.09(1) mm by 2.03(1) mm by 1.01(1) mm thick. The measurements were performed with the large face of the sample perpendicular to applied field. Again, this is not an ideal geometry for demagnetisation effects, but was necessary to have the [001] axis of the crystal parallel to the applied field, and to ensure good thermal coupling to the sample. The volume average demagnetisation factor for this sample is $N_m \sim 0.5$. This value was used to determine the average internal field experienced

by the sample for each external field used in the experiment. The largest effect on the field inside the sample occurs in very low fields, where the applied field can be reduced by a factor of ~ 1.75 .

One subtle point that should be mentioned is that in order to compute the average internal magnetic field experienced by the heat capacity sample, Eqn. (D.1) must be inverted. Using the demagnetisation factor of our sample, 0.5, the equivalent external field can be computed from the values of M and H_{int} provided by L’Hotel et al. [37]. By performing a linear interpolation to the average internal field as a function of the computed external field it is possible to convert applied fields into the average internal field experienced by the magnetic moments in the sample. The values of the internal field for the applied fields used in the heat capacity experiments in Ch. 3 are given in Table D.1.

B (T)	0.05	0.075	0.1	0.125	0.15	0.21	0.30	0.5	1.5
$\mu_0 H_{\text{int}}$ (T)	0.004	0.006	0.018	0.036	0.057	0.11	0.194	0.385	1.37

Table D.1: Values of H_{int} for corresponding values of the applied magnetic field for heat capacity measurements performed in this work.

D.3.3 Inelastic Neutron Scattering Demagnetisation Correction

Correcting for demagnetisation effects in the sample used in neutron scattering experiment has two different complications: some of the measurements were performed in fields well above the field range available in the data of L’Hotel et al. [37], and the sample used for the inelastic neutron scattering experiments is a tilted cylinder. Simple formulae exist for computing the demagnetisation factor of a cylinder with the long axis parallel and perpendicular to an applied magnetic fields, but if the cylinder is tilted away from either of these two cases, computing the effect of demagnetisation becomes more complex.

In order to address the problem that the single crystal data in L’Hotel et al. [37] extends up to only 4 T, and neutron data was collected in applied fields of up to $B = 9$ T, a series of data combinations and extrapolations were necessary. In order to extend the range of magnetisation data that can be used to compute the internal field, the low temperature data was combined with the data collected at $T = 2$ K for $\mu_0 H_{\text{int}} > 5$ T. L’Hotel et al. [37] find that for $B \gtrsim 5$ T, the powder magnetisation measured at $T = 2$ K and $T = 90$ mK are equal, making this combination justified. Based on this, when computing the internal field for neutron scattering measurements, the magnetisation used is a combination of 90 mK data from L’Hotel et al. [37] below $\mu_0 H_{\text{int}} = 4$ T, a linear interpolation between the 90 mK and 2 K magnetisation between $\mu_0 H_{\text{int}} = 4$ and 5 T, and $T = 2$ K between $\mu_0 H_{\text{int}} = 5$ and 6.86 T.

For fields greater than $\mu_0 H_{\text{int}} = 6.86$ T an extrapolation must be made to estimate the sample magnetisation at higher fields so that the demagnetisation correction can be computed. This has been done by numerically computing the susceptibility (dM/dH_{int}) for the $T = 2$ K data and performing a linear fit to the high field ($5.5 < \mu_0 H_{\text{int}} < 6.86$ T) susceptibility. Using this fit to the susceptibility, the magnetisation can be extrapolated iteratively using the formula:

$$M_{\text{ext}} = M_0 + \Delta H_{\text{int}} \frac{dM}{dH_{\text{int}}}, \quad (\text{D.9})$$

with M_0 being one of the measured magnetisation points. This method will lead to a slight underestimation of the magnetisation at high fields, but when compared to a linear extrapolation of the magnetisation, which would overestimate the magnetisation, the difference in the resulting internal field was in the fourth decimal place, which is negligible when compared to the overall magnitude of the field. The results of the extrapolation based on a linear fit to the susceptibility are shown in Fig. D.3.

The sample used in the neutron scattering experiments is a cylinder of length 21.84(1) mm and diameter 7.33(1) mm. There are some small deviations in the sides of the crystal away from a perfect cylinder, but we treat the sample as a perfect cylinder.




Figure D.3: Magnetisation measurements from Fig. D.1, along with data collected using a single crystal at $T = 90$ mK with the field applied along $[001]$ by L’Hotel et al. [37], and the magnetisation for fields above $B = 7$ T, based on a linear fit to the high field susceptibility.

The cylinder was tilted with respect to the applied magnetic field by an angle of 31.6° away from vertical, as measured by a computer controlled measurement device. This was necessary to align the $[001]$ axis of the crystal with the applied magnetic field.

The equation for the volume average demagnetisation of a finite length cylinder with the field applied along the long axis of the cylinder is given by Chen et al. [108]:

$$N_m = 1 - \frac{2lL_s}{\mu_0\pi a^2}, \quad (\text{D.10})$$

where a is the diameter of the cylinder, and l is the length of the cylinder divided by two, and L_s is the self inductance of solenoid representing the surface of the cylinder, which is defined as [108]:

$$L_s = \left(\frac{2\mu_0}{3l^2} \right) \left\{ (a^2 + l^2)^{1/2} \cdot (F(k_s)l^2 + E(k_s)(a^2 - l^2)) - 3a^3 \right\}. \quad (\text{D.11})$$

$F(k_s)$ and $E(k_s)$ are complete elliptic integrals of the first and second kind respectively, and $k_s^2 = a^2/(a^2 + l^2)$.

Using Eqns. (D.10) and (D.11) the longitudinal demagnetisation factor of the

sample is $N_{m,z} = 0.12855$ using the relationship $N_x + N_y + N_z = 1$, which is true when the field is applied along one of the primary axes of the sample, the perpendicular demagnetisation factor is $N_{m,(x,y)} = 0.435725$. This second term is necessary due to the tilt of the cylinder with respect to the applied magnetic field.

In order to account for the cylinder being tilted, it is necessary to take into account both the longitudinal and perpendicular demagnetisation factors, in the following way: the magnetisation along the direction of the applied field is determined from the data in Fig. D.3. This magnetisation is then split into two components associated with the principle axes of the cylinder:

$$M_x = M_s \sin \theta, \quad (\text{D.12})$$

$$M_z = M_s \cos \theta. \quad (\text{D.13})$$

M_x and M_z are then substituted into Eqn. (D.8), using $N_{m,z}$ and $N_{m,(x,y)}$ respectively instead of N_m , to obtain $H_{d,x}$ and $H_{d,z}$ respectively. These two demagnetising field components are then recombined to determine the size of the demagnetising field parallel to the applied field:

$$H_d = H_{d,z} \cos \theta + H_{d,(x,y)} \sin \theta. \quad (\text{D.14})$$

The component of the demagnetising field perpendicular to the applied field, which occurs because in this geometry the demagnetisation tensor is not diagonal, is neglected. Using H_d , the internal field of the sample is determined as a function of applied field, by inverting Eqn. (D.1). The resulting internal fields are given in Table D.2.

B (T)	0.21	0.5	0.75	1.0	1.5	2.0	3.0	4.0	5.0	6.0	7.0	8.0	9.0
$\mu_0 H_{\text{int}}$ (T)	0.166	0.45	0.698	0.946	1.44	1.94	2.94	3.94	4.94	5.93	6.93	7.93	8.93

Table D.2: Values of $\mu_0 H_{\text{int}}$ for corresponding values of the applied magnetic field for neutron scattering experiments in this work.

Appendix E

Background Subtraction

In order to extract the purely magnetic scattering contribution to the inelastic neutron scattering data, especially in the case of continuum scattering, it is essential to subtract an estimate of the non-magnetic background from the data. This was especially important for the 1.34 meV and 2.5 meV data, where the removal of the elastic line allows the use of the inelastic neutron scattering signal down to energies of 0.037 meV and 0.07 meV respectively. The background was constructed by combining energy ranges from various data sets where no magnetic signal is expected to be present.

The area in Fig. E.1(a-c) where magnetic scattering is present has been determined by fitting the maximum and minimum energies of magnetic scattering in the data (see Ch. 5) These fitted boundaries are shown by the blue points, magenta lines show the energy ranges used to “construct backgrounds” for the 1.34 meV, 2.5 meV, and 6.3 meV data.

Some of the energy ranges used to construct the “background” can be seen to overlap (Fig. E.1(b,c)) with the magnetic scattering. This overlap is within the uncertainty of the magnetic boundary for these data sets. Also, the magnetic scattering in these energy ranges is continuum scattering that is very weak, and tapers to zero gradually, so the effect of any overlap will be minimal. This has been confirmed by comparing these energy ranges to other data sets where no magnetic scattering is present and confirming that the neutron scattering intensities in these ranges are




Figure E.1: (a), (b), and (c) show the energy ranges of data sets collected at different fields used to construct a representation of the non-magnetic backgrounds for the 1.34 meV, 2.5 meV, and 6.3 meV data sets respectively. The magenta lines indicate the energy ranges taken from single angle data sets for each field. The red lines in panels (b) and (c) indicate the energy ranges extracted from a high temperature $T > 900$ mK data set. The dashed line in panel (c) indicates an energy range where interpolated data was used to form the background. All of these data sets were combined then smoothed to form a representation of the non-magnetic background. The shaded blue area in each panel shows the energy ranges containing magnetic scattering, the blue points are the data points used to determine the boundaries of this area.

within each others uncertainty. The data files have also been checked for any evidence of over-subtraction after the background has been removed.

For energy ranges where no data was available that was free of magnetic scattering, two different solutions were used. In some cases, high temperature data ($T > 900$ mK) was used, where magnetic scattering will be diffused throughout the Brillouin zone, and diffused in energy by thermal excitations. The energy ranges where high temperature data was incorporated into the 2.5 meV and 6.3 meV backgrounds are indicated by red lines in Fig. E.1(b,c). For energy ranges where even the high temperature data contained magnetic scattering it was necessary to interpolate the background so that it smoothly connected the background above and below the energy window. The energy range where this was necessary is shown by a dashed line near zero field and 2 meV in Fig. E.1(c).

Using the data ranges shown in Fig. E.1, background data sets were constructed for each of the three incident neutron energies used in the experiments (1.34, 2.5, and 6.3 meV) by combining the data from runs at each field in the energy ranges shown, weighted according to the counting time. Afterwards, the resulting background was smoothed between neighbouring energy bins and detectors to ensure that the statistics of the data sets were not made worse when the background was subtracted. These resulting backgrounds are shown in Fig. E.2. The obvious features in the data are the intense band of elastic scattering around $E = 0$ meV, the elastic line, which results from elastic scattering from the sample and parts of the instrument, and the sharp quadratically dispersing feature in all of the data sets, which comes from elastic scattering from the tails of the superconducting magnet used during the experiment. Since this is a relatively sharp feature compared to the underlying smooth background, the pixels near this feature were not averaged, but were kept as originally measured in the contributing data sets. The quality of the background subtraction can be seen in Fig. 5.2(m,n), note the very clean region with no signal up to the Zeeman energy gap above 1.5 meV. In particular note that the quadratic dispersive feature seen in Fig. E.2(b) is completely absent in the background subtracted data. Also note in Fig.

5.4(a), all of the high field data points below 1.7 meV are indistinguishable from zero energy with their uncertainty, giving confidence to the idea that all of the low field intensity is genuine magnetic scattering. This “background” was subtracted from all data sets to obtain the purely magnetic scattering.




Figure E.2: Slices showing the (a) 1.34 meV, (b) 2.5 meV, and (c) 6.3 meV non-magnetic background.

Appendix F

Neutron Absorption Correction

Absorption correction of the inelastic neutron scattering data was necessary as several of the data sets showed strong absorption effects. This occurred in two situations, the 1.34 meV data showed significant effects for all energies, as the energy of the incident neutrons is quite low, and at higher energy transfers in all of the inelastic data sets, due to the scattered neutrons having very low energies.

The attenuation of a beam of neutrons through a material is given by:

$$I = I_0 \exp(-\mu l), \quad (\text{F.1})$$

where I_0 is the incident neutron intensity, I is the intensity of the neutrons after passing through the material, μ is the linear attenuation coefficient for the material, and l is the distance through the material travelled by the neutron. For simplicity, the absorption factor is defined as $\alpha \equiv \exp(-\mu l)$. The linear attenuation coefficient has two contributions, absorption by nuclei, and incoherent scattering, which are characterised by σ_{abs} , the absorption cross section, and σ_{inc} , the incoherent scattering cross section. These two cross sections are given for all of the elements in $\text{Yb}_2\text{Ti}_2\text{O}_7$ in Table F.1. The incoherent scattering cross section is independent of the energy of

the neutron, while the absorption cross section has energy dependence:

$$\sigma_{\text{abs}}(E) = \sigma_{0,\text{abs}} \sqrt{\frac{E_0}{E}}, \quad (\text{F.2})$$

where $E_0 = 25.299$ meV.

Table F.1: The absorption and incoherent scattering cross sections of the elements in $\text{Yb}_2\text{Ti}_2\text{O}_7$. $\sigma_{0,\text{abs}}$ is given for thermal neutrons with energy $E_0 = 25.299$ meV.

Element	$\sigma_{0,\text{abs}}$ (10^{-24} cm ²)	σ_{inc} (10^{-24} cm ²)
Yb	34.803	4.00
Ti	6.091	2.87
O	0.0	0.001

The linear attenuation coefficient is defined in terms of σ_{inc} and σ_{abs} as:

$$\mu(E) = \frac{N}{v_0} \sum_{a=\text{Yb,Ti,O}} n_a (\sigma_{\text{inc},a} + \sigma_{\text{abs},a}(E)) \quad (\text{F.3})$$

which has units of cm^{-1} . v_0 is the volume of the cubic unit cell of $\text{Yb}_2\text{Ti}_2\text{O}_7$, $N = 8$ is the number of formula units in the unit cell, a indexes all of the elements in the material, and n_a is the number of times an element appears in the chemical formula.

For inelastic neutron scattering the energy of the neutron passing through the material will be different before and after the scattering event. This is accounted for by setting:

$$\alpha = \exp(-\mu_1 l_1 - \mu_2 l_2), \quad (\text{F.4})$$

where μ_1 and μ_2 are the attenuation coefficients, and l_1 and l_2 are the distances travelled through the sample, of the neutron before and after scattering respectively. μ_1 and μ_2 are computed using the energy of the neutron before and after scattering respectively. The absorption coefficient depends on the the energy transfer, the detector position, and the sample orientation, therefor the absorption coefficient was calculated individually for each sample orientation, detector, and energy bin in the data. The absorption coefficient was computed by Monte Carlo sampling l_1 and l_2

over the volume of the sample to calculate the average $\langle\alpha\rangle$. All intensities, after background subtraction, were then divided by the resulting absorption coefficient.

References

- [1] S. J. Blundell. *Magnetism in Condensed Matter*. Oxford University Press, Oxford, 2001.
- [2] C Castelnovo, R Moessner, and S L Sondhi. Magnetic monopoles in spin ice. *Nature*, 451(7174):42–5, 2008.
- [3] M. J. P. Gingras and P. A. McClarty. Quantum spin ice: a search for gapless quantum spin liquids in pyrochlore magnets. *Reports on Progress in Physics*, 77(5):056501, 2014.
- [4] S. T. Bramwell and M. J. P. Gingras. Spin ice state in frustrated magnetic pyrochlore materials. *Science*, 294(5546):1495–501, 2001.
- [5] J. A. Hodges, P. Bonville, A. Forget, M. Rams, K. Krolas, and G. Dhalenne. The crystal field and exchange interactions in $\text{Yb}_2\text{Ti}_2\text{O}_7$. *Journal of Physics: Condensed Matter*, 13:9301–9310, 2001.
- [6] B. Malkin, A. Zakirov, M. Popova, S. Klimin, E. Chukalina, E. Antic-Fidancev, Ph. Goldner, P. Aschehoug, and G. Dhalenne. Optical spectroscopy of $\text{Yb}_2\text{Ti}_2\text{O}_7$ and $\text{Y}_2\text{Ti}_2\text{O}_7\text{:Yb}^{3+}$ and crystal-field parameters in rare-earth titanate pyrochlores. *Physical Review B*, 70(7):075112, 2004.
- [7] H. B. Cao, A. Gukasov, I. Mirebeau, and P. Bonville. Anisotropic exchange in frustrated pyrochlore $\text{Yb}_2\text{Ti}_2\text{O}_7$. *Journal of Physics: Condensed Matter*, 21(49):492202, 2009.

- [8] J. Gaudet, D. D. Maharaj, G. Sala, E. Kermarrec, K. A. Ross, H. A. Dabkowska, A. I. Kolesnikov, G. E. Granroth, and B. D. Gaulin. Neutron spectroscopic study of crystalline electric field excitations in stoichiometric and lightly stuffed $\text{Yb}_2\text{Ti}_2\text{O}_7$. *Physical Review B*, 92(13):134420, 2015.
- [9] H. W. J. Blöte, R. F. Wielinga, and W. J. Huiskamp. Heat-Capacity Measurements on Rare-Earth Double Oxides $\text{R}_2\text{M}_2\text{O}_7$. *Physica*, 43:549–568, 1969.
- [10] M. A. Subramanian, G. Aravamudan, and G.V. Subba Rao. Oxide pyrochlores A review. *Progress in Solid State Chemistry*, 15(2):55–143, 1983.
- [11] B. D. Dunlap, G. K. Shenoy, J. M. Friedt, M. Meyer, and G. J. McCarthy. Crystal Field Properties of $\text{Yb}_2\text{Ti}_2\text{O}_7$ using ^{170}Yb Mossbauer Spectroscopy. *Physical Review B*, 18(4):1936–1942, 1978.
- [12] A. Sengupta. Crystal field studies on magnetic and nuclear properties of $\text{Yb}_2\text{Ti}_2\text{O}_7$. *Journal of Physics and Chemistry of Solids*, 60(3):331–336, 1999.
- [13] R. Siddharthan, B. Shastri, A. Ramirez, A. Hayashi, R. Cava, and S. Rosenkranz. Ising Pyrochlore Magnets: Low-Temperature Properties, Ice Rules, and Beyond. *Physical Review Letters*, 83(9):1854–1857, 1999.
- [14] S. T. Bramwell, M. N. Field, M. J. Harris, and I. P. Parkin. Bulk magnetization of the heavy rare earth titanate pyrochlores-a series of model frustrated magnets. *Journal of Physics*, 12:483–495, 2000.
- [15] J. A. Hodges, P. Bonville, A. Forget, and G. André. First-order transition in frustrated $\text{Yb}_2\text{Ti}_2\text{O}_7$ without long-range order. *Canadian Journal of Physics*, 79:1373–1380, 2001.
- [16] J. A. Hodges, P. Bonville, A. Forget, A. Yaouanc, P. Dalmás De Réotier, G. André, M. Rams, K. Królas, C. Ritter, P. C. M. Gubbens, C. T. Kaiser, P. J. C. King, and C. Baines. First-Order Transition in the Spin Dynamics of Geometrically Frustrated $\text{Yb}_2\text{Ti}_2\text{O}_7$. *Physical Review Letters*, 88(7):077204, 2002.

- [17] Y. Yasui, M. Soda, S. Iikubo, M. Sato, N. Hamaguchi, T. Matsushita, N. Wada, T. Takeuchi, N. Aso, and K. Kakurai. Ferromagnetic Transition of Pyrochlore Compound $\text{Yb}_2\text{Ti}_2\text{O}_7$. *Journal of the Physics Society Japan*, 72(11):3014–3015, 2003.
- [18] P. Bonville, J. A. Hodges, J.-Ph. Bouchard, E. Bertin, M. Ocio, P. Dalmas de Réotier, P. Regnault, J. P. Sanchez, S. Sosin, A. Yaouanc, M. Rams, K. Krolas, and H. M. Rønnow. Transition and Spin Dynamics at very low temperature in the pyrochlore $\text{Yb}_2\text{Ti}_2\text{O}_7$ and $\text{Gd}_2\text{Sn}_2\text{O}_7$. *Zakopane School of Physics*, (1):1–20, 2003.
- [19] J. S. Gardner, G. Ehlers, N. Rosov, R. W. Erwin, and C. Petrovic. Spin-spin correlations in $\text{Yb}_2\text{Ti}_2\text{O}_7$: A polarized neutron scattering study. *Physical Review B*, 70(18):180404, 2004.
- [20] P. Bonville, J. A. Hodges, E. Bertin, J.-Ph. Bouchaud, P. Dalmas De Réotier, L.-P. Regnault, H. M. Rønnow, J.-P. Sanchez, S. Sosin, and A. Yaouanc. Transitions and Spin Dynamics at Very Low Temperature in the Pyrochlores $\text{Yb}_2\text{Ti}_2\text{O}_7$ and $\text{Gd}_2\text{Sn}_2\text{O}_7$. *Hyperfine Interactions*, 156/157(1-4):103–111, 2004.
- [21] P. Dalmas de Réotier, V. Glazkov, C. Marin, A. Yaouanc, P. C. M. Gubbens, P. Sakarya, P. Bonville, A. Amato, C. Baines, and P. J. C. King. Studies of $\text{R}_2\text{Ti}_2\text{O}_7$ (R=Gd and Yb); new results. *Physica B: Condensed Matter*, 374-375:145, 2006.
- [22] K. A. Ross, J. P. C. Ruff, C. P. Adams, J. S. Gardner, H. A. Dabkowska, Y. Qiu, J. R. D. Copley, and B. D. Gaulin. Two Dimensional Kagome Correlations and Field Induced Order in the Ferromagnetic XY Pyrochlore $\text{Yb}_2\text{Ti}_2\text{O}_7$. *Physical Review Letters*, 103:227202, 2009.
- [23] H. B. Cao, A. Gukasov, I. Mirebeau, P. Bonville, C. Decorse, and G. Dhalenne. Ising versus XY Anisotropy in Frustrated $\text{R}_2\text{Ti}_2\text{O}_7$ Compounds as Seen by Polarized Neutrons. *Physical Review Letters*, 103:056402, 2009.

- [24] B. Z. Malkin, T. T. A. Lummen, P. H. M. van Loosdrecht, G. Dhalenne, and A. R. Zakirov. Static magnetic susceptibility, crystal field and exchange interactions in rare earth titanate pyrochlores. *Journal of Physics: Condensed Matter*, 22(27):276003, 2010.
- [25] K. A. Ross, L. R. Yaraskavitch, M. Laver, J. S. Gardner, J. a. Quilliam, S. Meng, J. B. Kycia, D. K. Singh, Th. Proffen, H. a. Dabkowska, and B. D. Gaulin. Dimensional evolution of spin correlations in the magnetic pyrochlore $\text{Yb}_2\text{Ti}_2\text{O}_7$. *Physical Review B*, 84(17):174442, 2011.
- [26] A. Yaouanc, P. Dalmas de Réotier, C. Marin, and V. Glazkov. Single-crystal versus polycrystalline samples of magnetically frustrated $\text{Yb}_2\text{Ti}_2\text{O}_7$: Specific heat results. *Physical Review B*, 84(17):1–3, 2011.
- [27] K. Ross, L. Savary, B. Gaulin, and L. Balents. Quantum Excitations in Quantum Spin Ice. *Physical Review X*, 1(2):021002, 2011.
- [28] S. Onoda. Effective quantum pseudospin-1/2 model for Yb pyrochlore oxides. *Journal of Physics: Conference Series*, 320:012065, 2011.
- [29] J. Thompson, P. A. McClarty, H. Rønnow, L.-P. Regnault, A. Sørge, and M. J. P. Gingras. Rods of Neutron Scattering Intensity in $\text{Yb}_2\text{Ti}_2\text{O}_7$: Compelling Evidence for Significant Anisotropic Exchange in a Magnetic Pyrochlore Oxide. *Physical Review Letters*, 106(18):187202, 2011.
- [30] J. D. Thompson, P. A. McClarty, and M. J. P. Gingras. Local susceptibility of the $\text{Yb}_2\text{Ti}_2\text{O}_7$ rare earth pyrochlore computed from a Hamiltonian with anisotropic exchange. *Journal of physics. Condensed matter*, 23(16):164219, 2011.
- [31] N. R. Hayre, K. A. Ross, R. Applegate, T. Lin, R. R. P. Singh B. D. Gaulin, and M. J. P. Gingras. Thermodynamic properties of the $\text{Yb}_2\text{Ti}_2\text{O}_7$ pyrochlore as a function of temperature and magnetic field: validation of a quantum spin ice exchange Hamiltonian. *Physical Review B*, 87:184423, 2013.

- [32] R. Applegate, N. Hayre, R. Singh, T. Lin, A. Day, and M. Gingras. Vindication of $\text{Yb}_2\text{Ti}_2\text{O}_7$ as a Model Exchange Quantum Spin Ice. *Physical Review Letters*, 109(9):097205, 2012.
- [33] L.-J. Chang, S. Onoda, Y. Su, Y.-J. Kao, K.-D. Tsuei, Y. Yasui, K. Kakurai, and M. R. Lees. Higgs transition from a magnetic Coulomb liquid to a ferromagnet in $\text{Yb}_2\text{Ti}_2\text{O}_7$. *Nature Communications*, 3:992, 2012.
- [34] K. A. Ross, Th. Proffen, H. A. Dabkowska, J. A. Quilliam, L. R. Yaraskavitch, J. B. Kycia, and B. D. Gaulin. Lightly stuffed pyrochlore structure of single-crystalline $\text{Yb}_2\text{Ti}_2\text{O}_7$ grown by the optical floating zone technique. *Physical Review B*, 86(17):174424, 2012.
- [35] A. Bertin, Y. Chapuis, P. Dalmas de Réotier, and A. Yaouanc. Crystal electric field in the the $\text{R}_2\text{Ti}_2\text{O}_7$ pyrochlore compounds. *Journal of Physics: Condensed Matter*, 24(25):256003, 2012.
- [36] R. M. D’Ortenzio, H. A. Dabkowska, S. R. Dunsiger, B. D. Gaulin, M. J. P. Gingras, T. Goko, J. B. Kycia, L. Liu, T. Medina, T. J. Munsie, D. Pomaranski, K. A. Ross, Y. J. Uemura, T. J. Williams, and G. M. Luke. Unconventional magnetic ground state in $\text{Yb}_2\text{Ti}_2\text{O}_7$. *Physical Review B*, 88(13):134428, 2013.
- [37] E. L’hotel, S. R. Giblin, M. R. Lees, G. Balakrishnan, L. J. Chang, and Y. Yasui. First-order magnetic transition in $\text{Yb}_2\text{Ti}_2\text{O}_7$. *Physical Review B*, 89(22):224419, 2014.
- [38] L. Pan, S. K. Kim, A. Ghosh, C. M. Morris, K. A. Ross, E. Kermarrec, B. D. Gaulin, S. M. Koohpayeh, O. Tchernyshyov, and N. P. Armitage. Low-energy electrodynamics of novel spin excitations in the quantum spin ice $\text{Yb}_2\text{Ti}_2\text{O}_7$. *Nature Communications*, 5:4970, 2014.
- [39] L.-J. Chang, M. R. Lees, I. Watanabe, A. D. Hillier, Y. Yasui, and S. Onoda. Static magnetic moments revealed by muon spin relaxation and thermo-

- dynamic measurements in the quantum spin ice $\text{Yb}_2\text{Ti}_2\text{O}_7$. *Physical Review B*, 89(18):184416, 2014.
- [40] J. Robert, E. L’Hotel, G. Remenyi, S. Sahling, I. Mirebeau, C. Decorse, B. Canals, and S. Petit. Spin dynamics in the presence of competing ferromagnetic and antiferromagnetic correlations in $\text{Yb}_2\text{Ti}_2\text{O}_7$. *Physical Review B*, 92(6):064425, 2015.
- [41] S. Bhattacharjee, S. Erfanifam, E. L. Green, M. Naumann, Z. Wang, S. Granovski, M. Doerr, J. Wosnitza, A. A. Zvyagin, R. Moessner, A. Maljuk, S. Wurmehl, B. Büchner, and S. Zherlitsyn. Acoustic signatures of the phases and phase transitions in $\text{Yb}_2\text{Ti}_2\text{O}_7$. *Physical Review B*, 93(14):144412, 2016.
- [42] L. D. C. Jaubert, O. Benton, J. G. Rau, J. Oitmaa, R. R. P. Singh, N. Shannon, and M. J. P. Gingras. Are Multiphase Competition and Order by Disorder the Keys to Understanding $\text{Yb}_2\text{Ti}_2\text{O}_7$? *Physical Review Letters*, 115(26):267208, 2015.
- [43] S. J. Blundell. Personal communication on muon experiments, 2015.
- [44] S. Curnoe. Quantum spin configurations in $\text{Tb}_2\text{Ti}_2\text{O}_7$. *Physical Review B*, 75(21):212404, 2007.
- [45] S. Onoda and Y. Tanaka. Quantum fluctuations in the effective pseudospin-1/2 model for magnetic pyrochlore oxides. *Physical Review B*, 83(9):094411, 2011.
- [46] P. A. McClarty, S. H. Curnoe, and M. J. P. Gingras. Energetic selection of ordered states in a model of the $\text{Er}_2\text{Ti}_2\text{O}_7$ frustrated pyrochlore XY antiferromagnet. *Journal of Physics: Conference Series*, 145:012032, 2009.
- [47] L. Onsager. Electric moments of molecules in liquids. *Journal of the American Chemical Society*, 58:1486–1493, 1936.

- [48] D. Hohlwein. Magnetic interactions and Onsager reaction field from paramagnetic diffuse neutron scattering. *Applied Physics A: Materials Science & Processing*, 742:740–742, 2002.
- [49] M. Hermele, Matthew P. A. Fisher, and L. Balents. Pyrochlore photons: The U(1) spin liquid in a S=1/2 three-dimensional frustrated magnet. *Physical Review B*, 69(6):064404, 2004.
- [50] S. J. Li, Z. Y. Zhao, C. Fan, B. Tong, F. B. Zhang, J. Shi, J. C. Wu, X. G. Liu, H. D. Zhou, X. Zhao, and X. F. Sun. Low-temperature thermal conductivity of $\text{Dy}_2\text{Ti}_2\text{O}_7$ and $\text{Yb}_2\text{Ti}_2\text{O}_7$. *Physical Review B*, 92(9):094408, 2015.
- [51] L. Pan, N. J. Laurita, K. A. Ross, B. D. Gaulin, and N. P. Armitage. A measure of monopole inertia in the quantum spin ice $\text{Yb}_2\text{Ti}_2\text{O}_7$. *Nature Physics*, 12(4):361–366, 2015.
- [52] Y. Tokiwa, T. Yamashita, M. Udagawa, S. Kittaka, T. Sakakibara, D. Terazawa, Y. Shimoyama, T. Terashima, Y. Yasui, T. Shibauchi, and Y. Matsuda. Thermal conductivity of quantum magnetic monopoles in the frustrated pyrochlore $\text{Yb}_2\text{Ti}_2\text{O}_7$. *Nature Communications*, 7:10807, 2015.
- [53] D. J. P. Morris, D. A. Tennant, S. A. Grigera, B. Klemke, C. Castelnovo, R. Moessner, C. Czternasty, M. Meissner, K. C. Rule, J.-U. Hoffmann, K. Kiefer, S. Gerischer, D. Slobinsky, and R. S. Perry. Dirac strings and magnetic monopoles in the spin ice $\text{Dy}_2\text{Ti}_2\text{O}_7$. *Science*, 326(5951):411–4, 2009.
- [54] L. Savary and L. Balents. Coulombic quantum liquids in Spin-1/2 pyrochlores. *Physical Review Letters*, 108(3):037202, 2012.
- [55] S. Lee, S. Onoda, and L. Balents. Generic quantum spin ice. *Physical Review B*, 104412:13, 2012.
- [56] Y. Wan and O. Tchernyshyov. Quantum Strings in Quantum Spin Ice. *Physical Review Letters*, 108(24):247210, 2012.

- [57] O. Benton, O. Sikora, and N. Shannon. Seeing the light: Experimental signatures of emergent electromagnetism in a quantum spin ice. *Physical Review B*, 86(7), 2012.
- [58] H. Yan, O. Benton, L. Jaubert, and N. Shannon. Living on the edge: ground-state selection in quantum spin-ice pyrochlores. *arXiv:1311.3501v1*, 2013.
- [59] M. Elhajal, B. Canals, R. Sunyer, and C. Lacroix. Ordering in the pyrochlore antiferromagnet due to Dzyaloshinsky-Moriya interactions. *Physical Review B*, 71(9):094420, 2005.
- [60] S. Palmer and J. Chalker. Order induced by dipolar interactions in a geometrically frustrated antiferromagnet. *Physical Review B*, 62(1):488–492, 2000.
- [61] S. T. Bramwell, M. J. P. Gingras, and J. N. Reimers. Order by disorder in an anisotropic pyrochlore lattice antiferromagnet. *Journal of Applied Physics*, 75(10):5523, 1994.
- [62] J. Champion, M. Harris, P. Holdsworth, A. Wills, G. Balakrishnan, S. Bramwell, E. Čížmár, T. Fennell, J. Gardner, J. Lago, D. McMorro, M. Orendáč, A. Orendáčová, D. Paul, R. Smith, M. Telling, and A. Wildes. $\text{Er}_2\text{Ti}_2\text{O}_7$: Evidence of quantum order by disorder in a frustrated antiferromagnet. *Physical Review B*, 68(2):1–4, 2003.
- [63] M. G. Townsend and W. A. Crossley. Magnetic susceptibility of rare-earth compounds with the pyrochlore structure. *Journal of Physics and Chemistry of Solids*, 29(4):593–598, 1968.
- [64] A. Abragam and B. Bleaney. *Electron Paramagnetic Resonance of Transition Ions*. Clarendon Press, Oxford, 1970.
- [65] S. J. Blundell and K. M. Blundell. *Concepts in Thermal Physics*. Oxford University Press, Oxford, second edition, 2010.

- [66] G. Ventura and M. Perfetti. *Thermal Properties of Solids at Room and Cryogenic Temperatures*. London, 2014.
- [67] E. S. R. Gopal. *Specific Heats at Low Temperatures*. Plenum Press, New York, 1966.
- [68] F. J. Morin and J. P. Maita. Specific Heats of Transition Metal Superconductors. *Physical Review*, 129(3):1115–1120, 1963.
- [69] R. Bachmann, F. J. DiSalvo, T. H. Geballe, R. L. Greene, R. E. Howard, C. N. King, H. C. Kirsch, K. N. Lee, R. E. Schwall, H. U. Thomas, and R. B. Zubeck. Heat capacity measurements on small samples at low temperatures. *Review of Scientific Instruments*, 43(2):205–214, 1972.
- [70] P. Sullivan and G. Seidel. Steady-State, ac-Temperature Calorimetry. *Physical Review*, 173(3):679–685, 1968.
- [71] A. W. Rost. *Magnetothermal Properties near Quantum Criticality in the Itinerant Metamagnet $Sr_3Ru_2O_7$* . PhD thesis, University of St. Andrews, 2009.
- [72] F. Pobell. *Matter and Methods at Low Temperatures*. Springer-Verlag, London, third edition, 2007.
- [73] S. Rojstaczer, D. Cohn, and G. Marom. Thermal expansion of Kevlar fibres and composites. *Journal of Materials Science Letters*, 4:1233–1236, 1985.
- [74] H. J. Schink and H. v. Lohneysen. Specific heat of Apiezon N grease at very low temperatures. *Cryogenics*, 21(10):591–592, 1981.
- [75] Lake Shore Cryotronics Inc. Lakeshore Cryotronics Ruthenium Oxide Product Overview, 2016.
- [76] R. G. Goodrich, D. Hall, E. Palm, and T. Murphy. Magnetoresistance below 1K and temperature cycling of ruthenium oxide-bismuth ruthenate cryogenic thermometers. *Cryogenics*, 38(2):221–225, 1998.

- [77] D. L. Martin. Specific Heats below 3 K of Pure Copper, Silver, and Gold, and of Extremely Dilute Gold-Transition Metal Alloys. *Physical Review*, 170:651–655, 1968.
- [78] H. A. Algra, L. J. de Jongh, H. W. J. Blöte, W. J. Huiskamp, and R. L. Carlin. Heat Capacity of Cs_2CoCl_4 Below 1 K, Compared with the $S = 1/2$ Linear Chain XY Model. *Physics B+C*, 82B:239–246, 1976.
- [79] O. Breunig, M. Garst, A. Rosch, E. Sela, B. Buldmann, P. Becker, L. Bohatý, R. Müller, and T. Lorenz. Low-temperature ordered phases of the spin-1/2 XXZ chain system Cs_2CoCl_4 . *Physical Review B*, 91(2):024423, 2015.
- [80] C. J. Mukherjee, R. Coldea, D. A. Tennant, M. Koza, M. Enderle, K. Habicht, P. Smeibidl, and Z. Tylczynski. Field-induced quantum phase transition in the quasi 1D XY-like antiferromagnet Cs_2CoCl_4 . *Journal of Magnetism and Magnetic Materials*, 272-276:920–921, 2004.
- [81] J. de Nobel and F.J. du Chatenier. Specific heats of dilute alloys of manganese in silver and copper at low temperatures and in magnetic fields. *Physica*, 25(7-12):969–979, 1959.
- [82] R. B. Griffiths. Ferromagnetic Heat Capacity in an External Magnetic Field near the Critical Point. *Physical Review*, 188(2):942, 1969.
- [83] G. L. Squires. *Introduction to the Theory of Thermal Neutron Scattering*. Cambridge University Press, Cambridge, 1978.
- [84] S. W. Lovesey. *Theory of Neutron Scattering from Condensed Matter*. Clarendon Press, Oxford, 1984.
- [85] R. I. Bewley, J. W. Taylor, and S. M. Bennington. LET, a cold neutron multi-disk chopper spectrometer at ISIS. *Nuclear Instruments and Methods in Physics Research A*, 637(1):128–134, 2011.

- [86] A. Furrer, J. Mesot, and T. Strassle. *Neutron Scattering in Condensed Matter Physics*. World Scientific, London, 2009.
- [87] R. A. Ewings, A. Buts, M. D. Le, J. van Duijn, I. Bustinduy, and T. G. Perring. HORACE: software for the analysis of data from single crystal spectroscopy experiments at time-of-flight neutron instruments. *arXiv*, 834:1–14, 2016.
- [88] T.-H. Han, J. S. Helton, S. Chu, D. G. Nocera, J. S. Rodriguez-Rivera, C. Broholm, and Y. S. Lee. Fractionalized excitations in the spin-liquid state of a kagome-lattice antiferromagnet. *Nature*, 492(7429):406–10, 2012.
- [89] M. Punk, D. Chowdhury, and S. Sachdev. Topological excitations and the dynamic structure factor of spin liquids on the kagome lattice. *Nature Physics*, 10:289–293, 2014.
- [90] A. G. Del Maestro. *Quantum Spin Fluctuations in the Heisenberg-Like Pyrochlore Antiferromagnet Gadolinium Titanate*. M.Sc., University of Waterloo, 2003.
- [91] Jordan D. Thompson. *Searching for the Magnetic Interactions in the Rare Earth Pyrochlore Oxide $\text{Yb}_2\text{Ti}_2\text{O}_7$* . M.sc., University of Waterloo, 2011.
- [92] R. Coldea, D. Tennant, K. Habicht, P. Smeibidl, C. Wolters, and Z. Tylczynski. Direct Measurement of the Spin Hamiltonian and Observation of Condensation of Magnons in the 2D Frustrated Quantum Magnet Cs_2CuCl_4 . *Physical Review Letters*, 88(13):137203, 2002.
- [93] L. Savary, K. A. Ross, B. D. Gaulin, J. P. C. Ruff, and L. Balents. Order by quantum disorder in $\text{Er}_2\text{Ti}_2\text{O}_7$. *Physical Review Letters*, 109(16):167201, 2012.
- [94] A. L. Chernyshev and M. E. Zhitomirsky. Spin waves in a triangular lattice antiferromagnet: Decays, spectrum renormalization, and singularities. *Physical Review B*, 79(14):1–28, 2009.

- [95] O. Benton, L. D. C. Jaubert, H. Yan, and N. Shannon. A spin-liquid with pinch-line singularities on the pyrochlore lattice. *Nature Communications*, 7:11572, 2016.
- [96] S. Isakov, K. Gregor, R. Moessner, and S. Sondhi. Dipolar Spin Correlations in Classical Pyrochlore Magnets. *Physical Review Letters*, 93(16):1–4, 2004.
- [97] C. L. Henley. Power-law spin correlations in pyrochlore antiferromagnets. *Physical Review B*, 71(1):014424, 2005.
- [98] L. Messio, B. Bernu, and C. Lhuillier. Kagome Antiferromagnet: A Chiral Topological Spin Liquid? *Physical Review Letters*, 108(20):207204, 2012.
- [99] A. Wietek, A. Sterdyniak, and A. M. Läuchli. Nature of chiral spin liquids on the kagome lattice. *Physical Review B*, 92(12):125122, 2015.
- [100] Appendix D: Sensor Calibration Accuracies. Technical report, 2000.
- [101] M. Watanabe, M. Morishita, and Y. Ootuka. Magnetoresistance of RuO_2 -based resistance thermometers below 0.3 K. *Cryogenics*, 41:143–148, 2001.
- [102] D. Prabhakaran and A.T. Boothroyd. Floating-zone growth of $\text{Na}_{0.8y}\text{A}_y\text{CoO}_2$ (A=Ca, Sr) single crystals. *Journal of Crystal Growth*, 318(1):924–926, 2011.
- [103] N. W. Ashcroft and N. D. Mermin. *Solid State Physics*. Saunder College Publishing, Orlando, 1976.
- [104] S. Wilkins. QLaue <https://github.com/stuwilkins/QLaue>.
- [105] W. G. Jenks, I. M. Thomas, and J. P. Wikswo Jr. SQUIDS, 1997.
- [106] M. McElfresh. *Fundamentals of Magnetism and Magnetic Measurements Featuring Quantum Design’s Magnetic Property Measurement System*. Quantum Design, Inc., San Diego, 1994.

- [107] D. X. Chen, E. Pardo, and A. Sanchez. Demagnetizing factors of rectangular prisms and ellipsoids. *IEEE Transactions on Magnetism*, 38(4 II):1742–1752, 2002.
- [108] D. X. Chen, J. A. Brug, and R. B. Goldfarb. Demagnetizing factors for cylinders. *IEEE Transactions on Magnetism*, 27(4):3601–3619, 1991.