Studies Towards Quantum Magnonics

A thesis submitted for the degree of

Doctor of Philosophy

Richard Graham Esdale Morris

University College

Trinity Term 2017

Studies Towards Quantum Magnonics

A thesis submitted for the degree of Doctor of Philosophy

Richard G. E. Morris, University College

Trinity Term 2017

Abstract

This thesis reports on recent results which pave the way for future experiments in the emerging field of quantum magnonics.

Chapter 1 presents a brief outline of the field of magnonics, which provides the context in which quantum magnonics has begun to develop.

Chapter 2 provides an introduction to the theory of spin waves, which is necessary to understand the experiments reported in the thesis.

In Chapter 3, the experimental methods and materials used to carry out the investigations in the thesis are described.

Chapter 4 describes the coupling of resonant magnon modes in a sphere of yttrium-iron garnet to photon modes in a coplanar-waveguide resonator. Strong coupling is achieved to multiple magnon modes, and a theoretical model is used to identify the magnon modes which couple most strongly to the photon mode.

In Chapter 5, the behaviour of propagating magnon modes is investigated in a waveguide formed from a thin film of yttrium-iron garnet. Two different configurations are investigated supporting different types of propagating mode, namely backward-volume and surface spin waves. Simulations are performed which reproduce the main features of the data.

Chapter 6 characterises the effect of the gadolinium-gallium garnet substrate on propagating spin waves. The magnitude of this effect is dependent on both the orientation and temperature of the sample.

Finally, Chapter 7 provides a short summary of the results of the thesis, and speculates on how they may inform future work in the field.

Contents

Chapter 1

Introduction

In the natural world, magnetism is visible to us through only a select few phenomena. Magnetite, an iron-based mineral, can be naturally magnetic, and it is this material which has historically been the source of our experiences of magnetism. Magnets made from magnetite are mentioned in various ancient sources. Some of the earliest writings about magnetism are from the 6^{th} century BC, when Thales of Miletus noted the ability of magnetic stones to attract iron, and suggested this was evidence they possessed some kind of vital force, or 'soul' [1]. In the dialogue Ion, written by Plato around the 4^{th} century BC, Socrates also mentions the ability of magnets to transfer their magnetism to small pieces of iron [2]. These texts demonstrate that some properties of magnetic materials have been known since at least the classical period. However, the mysterious nature of their attractive force has also led to many superstitions surrounding magnetism. One of the earliest uses of magnets was in divination in ancient China, before they were later used for navigation [3]. Others have ascribed magical healing properties to magnets, a belief which continues even today [4].

So whilst some properties of magnets have been known since antiquity, it was not until more recently that systematic investigations of magnetism were performed. Probably the first detailed study of the behaviour of lodestones was by Peter Peregrinus who wrote an extensive letter detailing his studies in 1269 [5], and is notable for the strong emphasis he places on experimentation as the basis for his conclusions. He identified that magnets have north and south poles and that unlike poles attract, and arguably also goes some way to identifying magnetic field lines around a lodestone using small needles of iron. Although much of his writing was concerned with using lodestones to build a perpetual motion machine [6], his is perhaps the first treatise on magnets which may be considered scientific in the modern sense.

Despite this promising start, further progress in the understanding of magnetism was slow. Various experimenters identified the inclination of magnetic compass needles relative to the Earth's surface, which led to the conclusion that the Earth itself was magnetic [6]. However, it was in the 19^{th} century that our understanding of magnetism progressed hugely. Important concepts such as the link between electricity and magnetism, which had been debated since antiquity, were confirmed and mathematically formulated. In 1930, Felix Bloch proposed that the presence of spin waves in magnetically ordered materials could explain the change in their saturation magnetisation close to absolute zero [7]. The first experimental observation of these waves was by Griffiths in 1946 [8], who noticed that ferromagnetic metals displayed peaks in their resistance when subjected to high-frequency radiation and applied magnetic fields. Later, propagating excitations were detected in insulating magnetic materials [9] and directly observed with Brillouin light scattering [10]. Devices based on spin waves have been found

to have many applications in microwave engineering, including filters, oscillators, phase shifters [11], isolators, and circulators [12].

The field of spintronics is concerned with the development of electronic devices which utilise the spin of electrons as well as their charge [13]. Perhaps the best-known product of this work is the effect known as giant magnetoresistance [14, 15], now widely used in hard disc drives. Related to spintronics is magnonics [16], which concentrates on using spin waves (or magnons, the quanta of spin waves) to transfer information in the absence of any corresponding charge current. This field has seen a surge in interest in recent years as a potential candidate for building new information processing devices which do not suffer from problems due to resistive heating [17]. Magnonic crystals [18, 19] have been used to create magnon-based transistors [20], and other devices such as novel logic gates [21, 22] have also been demonstrated.

While magnonic devices present interesting new directions for classical computing architectures, they also have potential applications in quantum information technology. The investigation of these possible applications has created a new field, emerging over the last few years, referred to as quantum magnonics [23]. One popular subject of research is the use of spin ensembles for quantum memory applications [24, 25]. The coupling to such ensembles is enhanced by a factor $\sqrt{N}$ where N is the number of spins [26], and so the magnon modes in ferromagnets, with their high spin density, are a natural candidate. Coupling these ensembles to photonic cavities would provide a way to interface between the magnon modes and quantum bits. After the initial demonstrations of such coupling [27, 28], there has been an explosion of interest in the area, with experiments investigating the details of this coupling in many different geometries and

configurations [29–43]. The coupling of magnons to a superconducting qubit via a cavity has also been investigated [44, 45].

Another potential use of magnonic systems is in the conversion of microwave-frequency excitations to optical excitations. Modern communications networks use both optical signals and microwave-frequency electrical signals. Optical fibres allow signals to be transmitted between nodes faster with lower losses, but electrical processors typically operate at gigahertz frequencies, so conversion between the two regimes is necessary. A similar scheme could be used to create a 'quantum internet' [46], but to be used for quantum information applications, the conversion between microwave and optical frequencies must be coherent [47]. Several experiments have investigated the possibility of using magnon modes in ferromagnetic materials for this purpose [48–51].

This new field of quantum magnonics has grown rapidly in the last few years, and has many potentially interesting applications in quantum information systems. This thesis describes a series of experiments which investigate the behaviour of spin-wave systems at temperatures in the region of 10 mK. This is the regime which is necessary for the operation of superconducting qubits, and so these experiments represent necessary first steps towards building more complex hybrid quantum systems incorporating magnons and superconducting qubits. In Chapter 2, the fundamental theory is laid out which is required to understand the types of spin waves investigated in the later chapters. Chapter 3 describes the experimental setup, techniques and materials which are used to study spin waves at millikelvin temperatures. In Chapter 4 are presented the results of an investigation into the coupling of magnon modes in a ferromagnetic sphere to photons in a planar superconducting resonator, a novel architecture

for magnon-photon coupling. Chapter 5 investigates propagating spin waves in a ferromagnetic waveguide, presenting both continuous-wave transmission measurements and time-resolved analysis of pulsed measurements. In Chapter 6 are identified temperature-dependent effects of the substrate on which the ferromagnetic waveguide is grown. Finally, Chapter 7 provides a short summary of the findings, and an outlook on future directions for this research.

Chapter 2

Theory of spin waves

Spin waves are excitations of a magnetically ordered material. A variety of magnetically ordered structures can support spin waves, including nuclear spins, electron spins or atomic ensembles. However, the studies in this thesis are concerned only with the study of electronic spin waves, and so the term 'spin wave' will be used to refer exclusively to excitations of electronic order. The behaviour of these spin waves depends on the properties, geometry and environment of the materials being considered. To understand these behaviours it is useful to consider under what conditions magnetic order can arise.

This chapter begins with an introduction to magnetism in physics, and moves on to explore the theory of spin waves specifically. The different types of spin wave, and how their characteristics are affected by the geometries in which they are excited, are discussed and an overview is given of how some properties of specific materials can change their behaviour.

2.1 Introduction to magnetism

Magnetism is intriguing as one of the forces that is least obvious in the world around us. Whereas electricity can be experienced through phenomena such as lightning storms and static electricity, magnetism is naturally apparent only in some rare minerals. Despite this, the two are fundamentally linked, and magnetism underlies many important effects, such as the forces experienced by moving charges.

Just as electric fields are created by distributions of charge, magnetic fields can be created by distributions of current (or moving charges). However, unlike electric fields there are no sources or sinks of magnetic flux; all magnetic field lines are closed so that the dipole, rather than the monopole, is the fundamental unit of a magnetic field distribution. The difference between these two cases is illustrated in Fig. 2.1.

Quantum mechanically, magnetism is linked to the concept of angular momen-

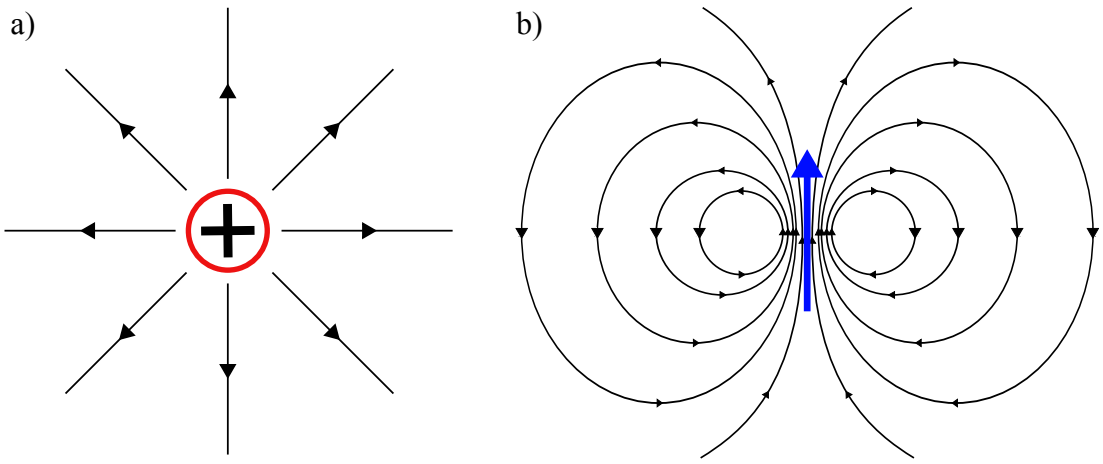

Figure 2.1: a) The electric field around a positively charged particle. The particle is a source of field lines. b) The magnetic field around a magnetic moment. All field lines are closed, so there is no net source of field.

tum, and it is the intrinsic spin angular momentum of electrons which gives them a magnetic dipole moment. This means that atoms and the materials they make up react to externally applied magnetic fields. Their response can be quantified through the relationship between the magnetic flux density $\mathbf{B}$ and magnetic field $\mathbf{H}$. In general

$$\mathbf{B} = \mu_0(\mathbf{H} + \mathbf{M}) \tag{2.1}$$

where μ_0 is the vacuum permeability, $\mathbf{B}$ is the magnetic flux density in Tesla, $\mathbf{H}$ is the magnetic field in amps per metre, and $\mathbf{M}$ is the magnetisation of the material, also in amps per metre.

The response of a material to an externally applied field is expressed through the relationship between $\mathbf{H}$ and $\mathbf{M}$. This is often quantified using a *susceptibility tensor* $\bar{\chi}$, defined such that

$$\mathbf{M} = \bar{\chi}\mathbf{H}. \tag{2.2}$$

The relation between $\mathbf{B}$ and $\mathbf{H}$ then becomes:

$$\mathbf{B} = \mu_0\big(\bar{I} + \bar{\chi}\big)\mathbf{H}, \tag{2.3}$$

where $\bar{I}$ is the three-dimensional identity matrix. This relation is often written in terms of a permeability:

$$\mathbf{B} = \bar{\mu}\mathbf{H}, \tag{2.4}$$

with $\bar{\mu} = \mu_0\bar{\mu}_\mathrm{r} = \mu_0\big(\bar{I} + \bar{\chi}\big)$. The tensor nature of the susceptibility reflects the fact that properties of many materials, such as crystalline anisotropy, can mean that the material responds differently depending on the orientation of the applied magnetic field. Two common simplifications to this relationship are linear materials and isotropic materials. In a linear material, all components of the

susceptibility tensor are independent of the applied field **H**. In an isotropic material the response is independent of direction so the susceptibility tensor $\bar{\chi}$ can be replaced by a scalar susceptibility χ, which may still be a function of field. A material that is both linear and isotropic has a susceptibility which is a scalar and constant as a function of applied field.

The relationship between the macroscopic quantities **M** and **H** allows us to place many materials into one of several categories [52]. Which of these categories a material falls into can give us some information about the structure and interactions of its constituent microscopic particles. The three main categories are diamagnets, paramagnets and ferromagnets.

- **Diamagnetic** materials exhibit the most basic type of magnetic behaviour, and all materials have a diamagnetic component to their magnetic response. The origin of diamagnetism is hard to explain in classical terms, but arises from a change in the shape of filled electron orbitals around the atomic nucleus, resulting in the formation of a net magnetic moment on the atom. Rather like Lenz's law, the moment formed due to this effect will always act to oppose the applied field, so for linear materials this means that $\chi < 0$. Diamagnetism is a very weak effect, however, and is usually dwarfed by any other magnetic response if it exists.

- **Paramagnetic** materials contain unpaired electrons that have a net magnetic moment, but the material overall has no spontaneous order. However, application of an external magnetic field will cause the individual electron moments to align with the field to minimise their interaction energy. This enhances the magnetic field inside the material, reflected in a susceptibility

$\chi > 0$ and giving it an overall net magnetisation. Paramagnets are typically linear at low fields or high temperatures, as will be explored in more detail later.

- **Ferromagnetic** materials are also made up of particles with net magnetic moments, but additional short-range 'exchange' interactions between the particles cause spontaneous ordering. Typically, in their bulk form these materials split into domains in the absence of any external fields, causing the local ordering direction to vary through the material. The short-range order within a domain satisfies the inclination for alignment between neighbours, but the overall arrangement of domains prevents the material having any net magnetisation, which would be energetically unfavourable due to the long-range dipole interaction. Applying an external field rearranges the domains, and the dynamics of this causes ferromagnets to exhibit hysteresis. This means that their magnetisation is a function not only of the field currently applied, but also of their magnetic history. Ferromagnets therefore do not have a simple susceptibility parameter which expresses how their magnetisation changes as a function of applied field.

As well as these major categories, we can identify two more types of magnetic ordering.

- **Antiferromagnetic** materials are similar to ferromagnetic materials, but the exchange interactions between neighbours favour anti-alignment. This means that even in an applied external field, there will be no net magnetisation.

- **Ferrimagnetic** materials are closely related to ferromagnetic and antiferromagnetic materials. The exchange interactions cause neighbouring spins to anti-align, but nearest neighbours have spins of different magnitudes. This means the material can have a net magnetisation despite the antiferromagnetic ordering of its spins.

Other, more exotic, types of order also exist, such as spiral ordering or vortex states such as skyrmions. Some of the most interesting magnetic systems are frustrated systems [53] where two competing interactions, or a combination of magnetic and geometric factors, mean that the system does not have a single obvious ground state of energy. This can lead to a very complex magnetic phase diagram.

2.2 Temperature dependence of magnetisation and saturation

When a magnetic field is applied to a paramagnetic or ferromagnetic material, its magnetisation will increase as the spins of the constituent particles align with the applied field. The stronger the applied field, the greater the energy cost to the spins of being misaligned and the greater the overall degree of alignment. An important factor opposing the alignment of spins is the disordering effect of finite temperature, so the magnetisation of a material will be a function of its temperature, as well as the applied field. Each type of magnetic material described responds differently to these two factors. To simplify the discussion in this section, we will assume materials are isotropic so that a scalar susceptibility

can be used.

2.2.1 Diamagnets

As described above, diamagnetism is caused by the change in shape of electron orbitals due to an applied field. The energy of thermal fluctuations at a given temperature can be approximated as $E_{\mathrm{tf}} \approx k_{\mathrm{B}}T$. Considering that atomic orbitals typically have energies on the order of the Rydberg constant $R \approx 13.6\,\mathrm{eV}$, and that $k_{\mathrm{B}}T \approx 25\,\mathrm{meV}$ at room temperature, thermal fluctuations are several orders of magnitude smaller than the energy scales that dominate diamagnetism. This means that any effect of temperature on diamagnetic behaviour is unlikely to be observed in a laboratory setting.

2.2.2 Paramagnets

In paramagnets, changes in magnetisation are caused by the competition between the two energy scales of the thermal fluctuations and the dipole interaction with the external field. Thermal fluctuations act to promote disorder, whilst the external field makes disorder energetically unfavourable. The stronger the external field, the smaller the degree of disorder and the larger the magnetisation. Lowering the temperature reduces the amount of disorder due to thermal fluctuations, and so at a lower temperatures a greater degree of ordering will be observed for the same applied field. This is reflected in an increased value of susceptibility at lower temperatures. Many materials with weak interactions between their constituent spins display a susceptibility which is inversely proportional to the

temperature, known as the Curie-Weiss law:

$$\chi = \frac{C}{T - T_{\mathrm{C}}}, \tag{2.5}$$

where C is the Curie constant of the material (in T^{-1}) which depends on the density and magnitude of the spins, and T_{C} is the Curie temperature, which can be positive or negative. In the case it is positive, the relationship in Eq. 2.5 only holds above T_{C}. As the strength of the interactions between the constituent spins tends to zero, so does the Curie temperature and we obtain the Curie susceptibility,

$$\chi = \frac{C}{T}. \tag{2.6}$$

However, when the applied field becomes sufficiently strong relative to the disordering effect of temperature, all of the spins are perfectly aligned. At this point the material is said to be saturated and no further increase in its magnetisation is possible. The magnetisation at which this occurs is referred to as the *saturation magnetisation* and is often denoted by M_{s}. This means that the magnetisation has a simultaneous dependence on temperature and field which is determined by the Langevin function

$$M \propto L(x) = \coth(x) - \frac{1}{x}, \tag{2.7}$$

where $x = \frac{\mu_{\mathrm{B}} B}{k_{\mathrm{B}} T}$. This dependence is illustrated as a function of magnetic field in Fig. 2.2 for two different temperatures. At low B or high T where $\frac{\mu_{\mathrm{B}} B}{k_{\mathrm{B}} T}$ is small, this function can be expanded to recover Curie's law. At room temperature the material can be approximated as linear up to very high magnetic fields.

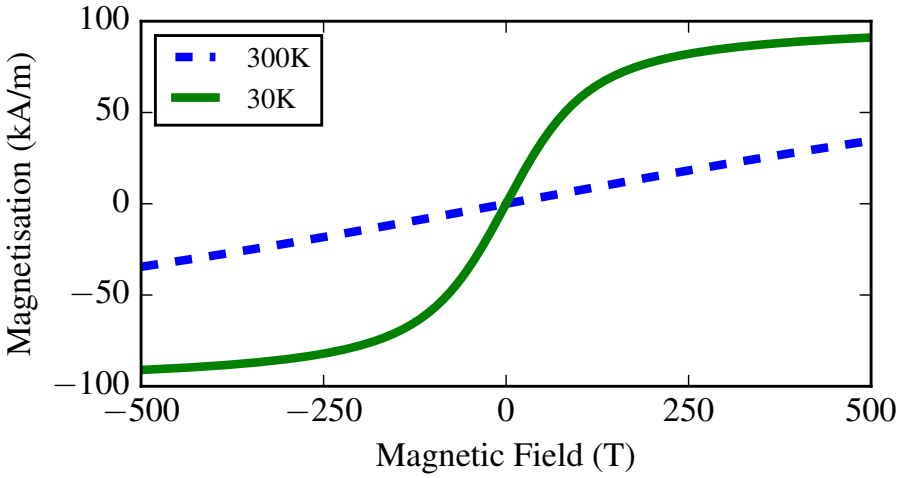

Figure 2.2: Magnetisation according to the Langevin function for a paramagnet with a saturation magnetisation of $100\,\mathrm{kAm^{-1}}$. At 300 K the magnetisation is linear even up to very high field, while at 30 K saturation effects become visible.

2.2.3 Ferromagnets

In ferromagnets the picture is complicated by the existence of a third energy scale. In addition to dipole interactions of spins with each other and the applied field, there is also the short-ranged exchange interaction. This is a quantum-mechanical effect caused by a combination of Pauli exclusion and the Coulomb interaction.

In ferromagnets, the exchange interaction favours the alignment of neighbouring spins, while the dipole interaction favours their anti-alignment. In general, for a single pair of spins the exchange interaction will be many times stronger than the dipole interaction, so alignment wins. However, the exchange interaction has much shorter range than the dipole interaction, so in a bulk material containing many spins the total dipole energy cost incurred by the interaction of the spins with each other will grow significantly. The competition between these

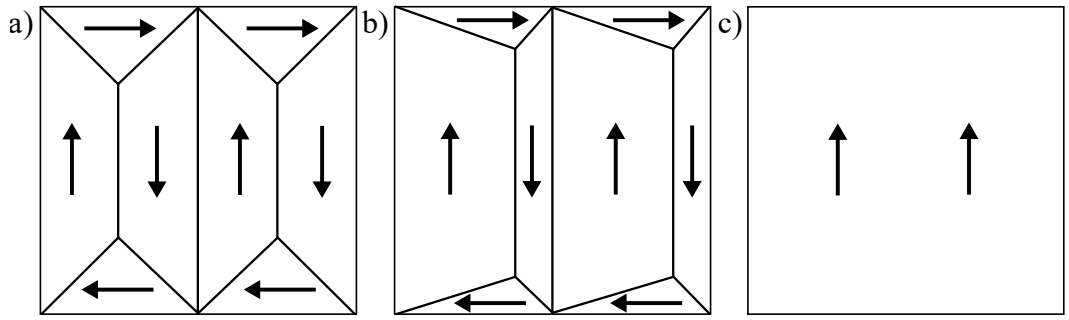

Figure 2.3: a) When there is no external field, the domains in a ferromagnet are aligned so as to produce no net magnetisation. b) As a weak magnetic field is applied in the vertical direction, domains aligned with the field grow at the expense of misaligned domains. c) Eventually when the field is strong enough, the material will be a single domain with a large net magnetisation.

effects causes the formation of domains; regions of material which are locally aligned, but arranged relative to other domains so that the total magnetisation of the material is zero. This is shown schematically in Fig. 2.3a. Spins at the edges of domains, known as 'domain walls', experience an energy penalty from the exchange interaction for being misaligned with their neighbours, but this is outweighed by the favourable dipole energy for minimising the self-interaction of the spins over the whole material.

When a magnetic field is applied, the spins can lower their energy by aligning with the external field. As the field strength increases, the energy cost of being misaligned with the external field approaches the self-interaction energy. Domain walls will move so that domains aligned with the field grow at the expense of misaligned domains, as illustrated in Fig. 2.3b. Eventually, the external interaction energy exceeds the self-interaction energy and the material becomes saturated, as in Fig. 2.3c. Since the domains are already highly ordered, bringing them into alignment by applying a field causes the magnetisation of the material

to increase rapidly. This means that ferromagnets typically saturate at much lower fields than paramagnets.

The effect of temperature is also different in ferromagnets. The exchange interaction is very strong, so local ordering of moments can be energetically favourable even at temperatures significantly above room temperature. However, as temperature increases, so too do the thermal fluctuations which create disorder within domains. Although the energy of the applied field interaction may exceed the self-interaction energy sufficiently to align the domains, it will not compare with the energy scale of the thermal fluctuations unless it is very strong indeed, as seen in the previous section on paramagnets. This means that M_s in the ferromagnet decreases as temperature increases and the domains become less ordered. At sufficiently high temperatures, the energy of thermal fluctuations exceeds even that of the exchange interaction so domains are effectively no longer locally ordered. This critical temperature is the Curie temperature T_C, and above T_C the ferromagnet undergoes a phase change to become paramagnetic, with its susceptibility then obeying the Curie-Weiss law (Eq. 2.5).

In general, the exact dependence of the magnetisation of a ferromagnet on external field and temperature is complicated, depending as it does on a variety of material-specific parameters, as well as the dynamics of domain formation and movement. While an analytic approach to modelling this behaviour is possible in principle, in most cases it is easier to simply measure the relationship between M and H and plot the result.

Though the details of this relationship in a material may be hard to determine analytically, it is still possible to describe some general features of the M-H relationship, one of which is hysteresis. If an M-H curve is hysteretic, the

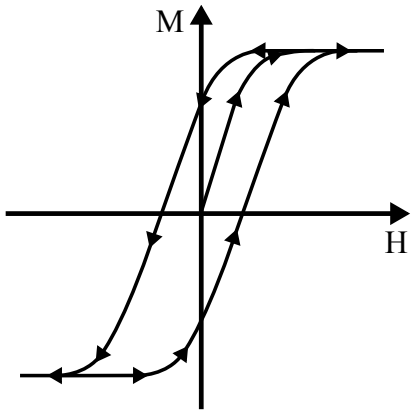

Figure 2.4: Illustration of a hysteresis curve for a ferromagnet. Applying a magnetic field saturates the material at a field which is typically much lower than that required to saturate a paramagnet. Removing the applied field again does not cause the magnetisation to return to zero.

magnetisation depends not only on the current value of the applied field, as for paramagnets and diamagnets, but also on the history of the field. As a ferromagnet is magnetised and its domain walls move, some of the ensuing spin rotations will be due to irreversible processes. These could include local trapping of spins due to defects in the crystal. Overcoming the energy barrier to rotate those spins requires a certain additional cost. Therefore, spins in these sites will not rotate until the energy cost from the applied field overcomes this barrier. Reducing the field again will not cause the spin to rotate back until the energy difference again overcomes the barrier.

This is illustrated in the hysteresis loop in Fig. 2.4. An initially unmagnetised ferromagnet is subjected to a magnetic field strong enough to saturate it. Reducing the field to zero does not return the magnetisation to zero, but leaves a finite magnetisation in the sample. To return the magnetisation to zero requires the application of a magnetic field in the opposite direction to the initial saturating field.

2.3 Magnons

Magnons are excitations of magnetic order, the quanta of spin waves. Since they represent deviations from perfect ordering in a magnetic lattice, a crucial prerequisite for studying magnon propagation is an ordered magnetic material. The most appropriate materials in which to study magnons are therefore ferromagnets, since these exhibit a high degree of magnetic ordering at temperatures and applied fields which are easily accessible experimentally. An important characteristic of ferromagnets is strong exchange coupling between neighbouring spins. We shall use this characteristic to perform a simple derivation of some of the properties of magnons.

As described in the previous section, applying an external magnetic bias field gives a ferromagnet a net magnetisation and causes the individual spins in the material to align with the bias field. This is shown schematically for a one-dimensional chain of spins in Fig. 2.5a. Deviations from this perfect ordering can be described as excitations, and these excitations are spin waves, whose quanta are magnons. A single magnon corresponds to a reduction in the total magnetic moment of the material by a single unit of angular momentum $\hbar$. This is analogous to how the deviations from perfect spatial order in a crystalline material are described using acoustic waves and their quanta, phonons.

To understand the characteristics of magnons, we should examine the behaviour of spins when they experience applied magnetic fields.

2.3.1 Spin precession

Consider a spin which experiences a magnetic field. The Hamiltonian for its interaction with the field is given by [54]:

$$\hat{\mathcal{H}} = -\mu_s \hat{\mathbf{S}} \cdot \mathbf{B} \tag{2.8}$$

where μ_s is the magnetic moment of the spin, $\hat{\mathbf{S}}$ is the spin operator, and $\mathbf{B}$ the applied field. Taking the z-axis to be defined by the direction of the bias field, we have $\mathbf{B} = B_0\hat{\mathbf{z}}$, and

$$\hat{\mathcal{H}} = -\mu_s B_0 \hat{S}^z. \tag{2.9}$$

The eigenstates of this Hamiltonian are those of the $\hat{S}^z$ operator corresponding to a spin aligned or anti-aligned with the z-axis. For a spin-$\frac{1}{2}$ particle, this gives energies:

$$E_\pm = \mp\frac{1}{2}\mu_s B_0. \tag{2.10}$$

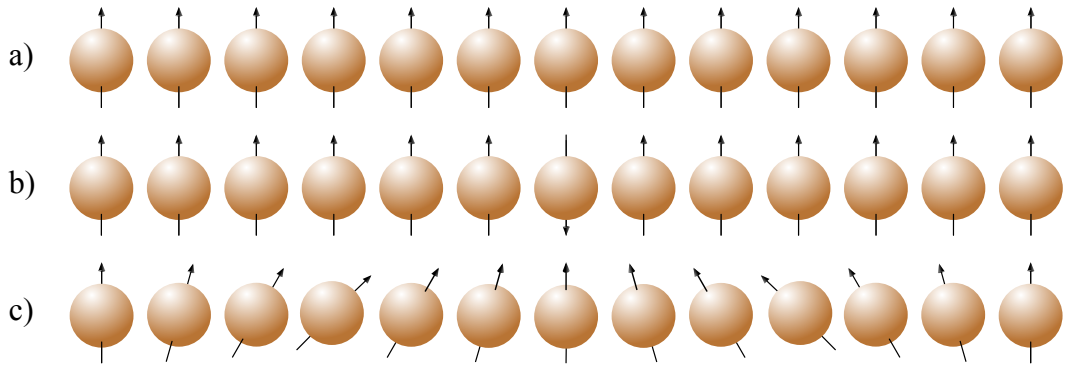

Figure 2.5: a) On application of an external magnetic bias field, the spins in a ferromagnet become aligned with the bias direction. b) The smallest possible excitation can be envisaged as a single spin flip (or a reduction of its total spin by 1 if it has a spin greater than $\frac{1}{2}$). c) The single spin flip is not an eigenstate. To form an eigenstate we 'spread' the excitation over all of the spins, distributing the phase shift across the lattice. Classically this can be visualised as a 'wave' of precessing spins.

A general spin state can be represented as a superposition of the up and down eigenstates. The time evolution of the general state is determined by the energies of the eigenstates in Eq. 2.10 and the time-dependent Schrödinger equation. Defining $\omega = \frac{\mu_s B_0}{\hbar}$, a general spin state is given by

$$|\phi(t)\rangle = a_- e^{\frac{-i\omega t}{2}}|-\rangle + a_+ e^{\frac{i\omega t}{2}}|+\rangle, \tag{2.11}$$

where $|+\rangle$ is the state of a spin aligned with the z-axis and $|-\rangle$ is the state of a spin anti-aligned with the z-axis.

Suppose a spin is initially aligned with the axis $\mathbf{n} = (\sin\theta, 0, \cos\theta)$. Then its wavefunction is given by Eq. 2.11 with $a_- = \sin\frac{\theta}{2}$ and $a_+ = \cos\frac{\theta}{2}$. Now consider the expectation values of the three spin operators $\hat{S}^x$, $\hat{S}^y$, and $\hat{S}^z$:

$$\langle \hat{S}^x \rangle = \frac{1}{2}\sin\theta\cos\omega t, \tag{2.12a}$$

$$\langle \hat{S}^y \rangle = \frac{1}{2}\sin\theta\sin\omega t, \tag{2.12b}$$

$$\langle \hat{S}^z \rangle = \frac{1}{2}\cos\theta. \tag{2.12c}$$

We see that the spin is inclined at an angle θ to the magnetic field axis, and precesses about that axis at a frequency ω. This is known as *Larmor precession*, and ω as the Larmor frequency.

Now consider the relationship between the expectation of the spin and its

time dependence

$$\langle \hat{\mathbf{S}} \rangle = \frac{1}{2} \begin{pmatrix} \sin\theta \cos\omega t \\ \sin\theta \sin\omega t \\ \cos\theta \end{pmatrix}, \tag{2.13a}$$

$$\frac{\mathrm{d}\langle \hat{\mathbf{S}} \rangle}{\mathrm{d}t} = \frac{\omega}{2} \begin{pmatrix} -\sin\theta \sin\omega t \\ \sin\theta \cos\omega t \\ 0 \end{pmatrix}. \tag{2.13b}$$

We notice that

$$\frac{\mathrm{d}\langle \hat{\mathbf{S}} \rangle}{\mathrm{d}t} = \frac{\omega}{2B} \mathbf{B} \times \langle \hat{\mathbf{S}} \rangle. \tag{2.14}$$

Since the magnetisation of a large sample is simply proportional to the average magnetic moment of each of the particles within it, this process gives us a quantum-mechanical justification of the classical magnetic torque equation (see Eq. 2.31), which will be used extensively in later parts of the thesis.

2.3.2 Magnons on an Ising chain

We have seen that isolated spins in a magnetic field can precess at a frequency determined by the strength of the external field. However, this treatment did not take into account any interactions between the spins themselves. We now examine how magnons can arise from spin-spin interactions. To illustrate this we use an Ising model, where spins can only take the values $\pm\frac{1}{2}$. We make the simplifying assumptions of a one-dimensional chain where the only interactions are between nearest neighbours.

Consider the Hamiltonian of this system with exchange interactions between

the nearest neighbours:

$$\hat{\mathcal{H}} = -\sum_{m=0}^{N-1} J\hat{\mathbf{S}}_m \cdot \hat{\mathbf{S}}_{m+1}. \tag{2.15}$$

J is the exchange constant and has units of energy, N is the number of spins in the chain, and the sum is over all pairs of nearest neighbours in the lattice. By expanding the spin operators into their components and using the ladder operators $\hat{S}^{\pm} = \hat{S}^x \pm i\hat{S}^y$ Eq. 2.15 can be rewritten as

$$\hat{\mathcal{H}} = -\sum_{m=0}^{N-1} J\left(\hat{S}_m^z\hat{S}_{m+1}^z + \frac{1}{2}\left(\hat{S}_m^+\hat{S}_{m+1}^- + \hat{S}_m^-\hat{S}_{m+1}^+\right)\right). \tag{2.16}$$

This system has a ground state energy of $E_0 = -\frac{1}{4}JN$. The combination of ladder operators in the second half of Eq. 2.16 will contribute only between pairs of neighbouring spins with opposite alignment. Since the spins have magnitude $\frac{1}{2}$, the smallest possible excitation is when one spin is flipped from an aligned to an anti-aligned state, as visualised in Fig. 2.5b. This can be achieved by applying the appropriate ladder operator to the ground state to flip the spin at a certain position. Each flipped spin will increase the energy of the state by J due to the two unfavourable interactions with its neighbours.

The state with a single spin flipped from up to down at position m in the chain we denote $|m\rangle$. This state is created by applying the spin lowering operator for position m to the ground state, denoted by $|0\rangle$, where all spins point up: $|m\rangle = S_m^-|0\rangle$. However, the combinations of ladder operators in the Hamiltonian also have the effect of 'moving' the flipped spin one space up or down the chain. Therefore, if we apply the Hamiltonian to the state $|m\rangle$, the result is:

$$\hat{\mathcal{H}}|m\rangle = (E_0 + J)|m\rangle - \frac{1}{2}J|m+1\rangle - \frac{1}{2}J|m-1\rangle. \tag{2.17}$$

This is not a multiple of the original state, so the single flip state $|m\rangle$ is not an eigenstate of the Hamiltonian. To create an eigenstate, we form a coherent

superposition of states with the spin flip localised at each position in the lattice with a different phase factor

$$|k\rangle = \frac{1}{\sqrt{N}} \sum_{m=0}^{N-1} e^{ikma} |m\rangle, \tag{2.18}$$

where a is the lattice parameter denoting the spacing of spins in the lattice and k is the wavenumber indicating the length scale over which the phase factor repeats. This wavefunction is an eigenstate, and represents a magnon. Now, moving the flipped spin up or down a place in the chain is just equivalent to phase shifting the wavefunction by a factor e^{ika}, so

$$\hat{\mathcal{H}}|k\rangle = (E_0 + J(1 - \cos ka))|k\rangle, \tag{2.19}$$

and the energy of an excitation with this wavenumber is:

$$\hbar\omega(k) = E(k) - E_0 = J(1 - \cos ka). \tag{2.20}$$

In the regime $ka \ll 1$, the cosine term can be expanded to give rise to the quadratic dispersion of magnons, with $\hbar\omega(k) \approx J\frac{(ka)^2}{2}$. Since a is typically on the scale of nanometres, this approximation is valid over a very wide range of wavenumbers. This model can provide a good description of the behaviour of magnons in ferromagnet, or other sytems which can be approximated as ferromagnetic.

2.4 The magnetostatic approximation

Magnons are excitations describing small perturbations to the magnetisation of a ferromagnetic system. We have examined the microscopic behaviour of spins in a magnetic field, and derived the characteristics of a simple class of magnons in

ferromagnetic materials. However, the studies in this thesis are mostly concerned with magnons whose behaviour is governed by another important concept: the magnetostatic approximation. This approximation defines a class of excitations which do not introduce any curl to the magnetic field of the material in which they exist. It consists of taking $\nabla \times \mathbf{H} = 0$ in Maxwell's equations, which allows the definition of a magnetic potential Ψ through $\mathbf{H} = \nabla\Psi$. This simplifies the analysis of many problems, but we must be aware of the circumstances under which it is valid, and the assumptions which underlie it.

In this section, we shall begin by examining the equations which govern the behaviour of waves in the types of material in which we are interested. Then, we shall derive the magnetic susceptibility when the magnetic dynamics are in the form of small perturbations, to see how an external field affects the magnetisation. Finally, we shall combine these to determine the characteristics of the spin waves which are allowed in this situation. These are magnetostatic spin waves, for which the magnetostatic approximation is valid.

2.4.1 Plane waves in magnetisable materials

Consider Maxwell's equations for polarisable and magnetisable materials [55]:

$$\nabla \cdot \mathbf{D} = \rho_{\mathrm{f}}, \tag{2.21a}$$

$$\nabla \cdot \mathbf{B} = 0, \tag{2.21b}$$

$$\nabla \times \mathbf{E} = -\frac{\partial \mathbf{B}}{\partial t}, \tag{2.21c}$$

$$\nabla \times \mathbf{H} = \mathbf{J}_{\mathrm{f}} + \frac{\partial \mathbf{D}}{\partial t}. \tag{2.21d}$$

Here the $\mathbf{E}$, $\mathbf{D}$, $\mathbf{B}$ and $\mathbf{H}$ fields have their usual meanings, and ρ_f and $\mathbf{J}_\mathrm{f}$ describe, respectively, the free charges and currents in the material. Since we are concerned with the behaviour of wave-like excitations, we consider plane-wave solutions to the Maxwell equations where, for instance, $\mathbf{H} = \mathbf{H_0} e^{i(\mathbf{k}\cdot\mathbf{r}-\omega t)}$. Substituting this dependence into the left-hand sides of the Maxwell equations we find that

$$\nabla \cdot \mathbf{D} = i\mathbf{k} \cdot \mathbf{D}, \tag{2.22a}$$

$$\nabla \cdot \mathbf{B} = i\mathbf{k} \cdot \mathbf{B}, \tag{2.22b}$$

$$\nabla \times \mathbf{E} = i\mathbf{k} \times \mathbf{E}, \tag{2.22c}$$

$$\nabla \times \mathbf{H} = i\mathbf{k} \times \mathbf{H}. \tag{2.22d}$$

Just as the relationship between $\mathbf{B}$ and $\mathbf{H}$ can be described with a relative permeability tensor $\bar{\mu}_\mathrm{r}$ through

$$\mathbf{B} = \mu_0 \bar{\mu}_\mathrm{r} \mathbf{H}, \tag{2.23}$$

so too can $\mathbf{E}$ and $\mathbf{D}$ be related by a relative permittivity tensor $\bar{\epsilon}_\mathrm{r}$ through

$$\mathbf{D} = \epsilon_0 \bar{\epsilon}_\mathrm{r} \mathbf{E}, \tag{2.24}$$

where ϵ_0 is the vacuum permittivity. We assume that the material is electrically linear and isotropic, so that $\bar{\epsilon}_\mathrm{r} \to \epsilon_\mathrm{r}$ and $\mathbf{D} = \epsilon_0 \epsilon_\mathrm{r} \mathbf{E}$ where ϵ_r is a constant. This approximation is valid for most commonly used magnetic materials in spintronics. We also assume that there are no free charges or currents so $\rho_\mathrm{f} = \mathbf{J}_\mathrm{f} = 0$. Under these conditions, and using $\mathbf{B} = \mu_0(\mathbf{H} + \mathbf{M})$ and Eq. 2.22, the Maxwell equations

become

$$\mathbf{k} \cdot \mathbf{D} = 0, \tag{2.25a}$$

$$\mathbf{k} \cdot (\mathbf{H} + \mathbf{M}) = 0, \tag{2.25b}$$

$$\mathbf{k} \times \mathbf{D} = \mu_0 \epsilon_0 \epsilon_\mathrm{r} \omega (\mathbf{H} + \mathbf{M}), \tag{2.25c}$$

$$\mathbf{k} \times \mathbf{H} = -\omega \mathbf{D}. \tag{2.25d}$$

Taking the cross product of $\mathbf{k}$ with Eq. 2.25d and using Eq. 2.25c we arrive at

$$\mathbf{k} \times (\mathbf{k} \times \mathbf{H}) = -\omega^2 \mu_0 \epsilon_0 \epsilon_\mathrm{r} (\mathbf{H} + \mathbf{M}), \tag{2.26}$$

which can be rewritten as

$$(\mathbf{k} \cdot \mathbf{H})\mathbf{k} - k^2 \mathbf{H} = -\omega^2 \mu_0 \epsilon_0 \epsilon_\mathrm{r} (\mathbf{H} + \mathbf{M}). \tag{2.27}$$

Two limits of Eq. 2.27 can be considered [56]. The first is the case of a non-magnetisable material where $\mathbf{M} = 0$, when Eq. 2.25b can be used to reduce it to

$$k^2 \mathbf{H} = \omega^2 \mu_0 \epsilon_0 \epsilon_\mathrm{r} \mathbf{H}, \tag{2.28}$$

yielding, as expected, the normal dispersion relation for plane electromagnetic waves in a linear, isotropic medium,

$$\omega = \frac{k_0}{\sqrt{\mu_0 \epsilon_0 \epsilon_\mathrm{r}}}. \tag{2.29}$$

In the case that $\mathbf{M} \neq 0$ we can rearrange Eq. 2.27 by noticing that $(\mathbf{k} \cdot \mathbf{H})\mathbf{k} = \left(\mathbf{k}\mathbf{k}^\mathbf{T}\right)\mathbf{H}$

$$\left(\mathbf{k}\mathbf{k}^\mathbf{T} - k^2 \bar{I}\right)\mathbf{H} = -\omega^2 \mu_0 \epsilon_0 \epsilon_\mathrm{r} (\mathbf{H} + \mathbf{M}), \tag{2.30}$$

where $\bar{I}$ is the three-dimensional identity matrix. To get further with this equation, we need to find some kind of relationship between $\mathbf{H}$ and $\mathbf{M}$. This we do through the susceptibility, which was introduced in Eq. 2.2.

2.4.2 The Polder susceptibility tensor

We start from the magnetic torque equation [56], whose form we justified quantum mechanically in Section 2.3.1

$$\frac{\partial \mathbf{M}}{\partial t} = \mu_0 \gamma (\mathbf{M} \times \mathbf{H}), \tag{2.31}$$

here γ is the *gyromagnetic ratio*, composed of various fundamental quantities related to the electron and whose value for the materials we consider is $-28.1\,\mathrm{GHz/T}$. The negative sign arises from the negative charge of the electron, and reflects the relative direction of the magnetisation and the torque it experiences.

We now assume that the material being examined is subject to a biasing magnetic field along the z-axis which is large compared to any other magnetic fields, saturating the magnetisation of the material in the direction of the bias field. Therefore the magnetisation and magnetic field consist of a large component aligned with the bias field, and small deviations from it:

$$\mathbf{M} = \begin{pmatrix} m_x \\ m_y \\ M_\mathrm{s} - m_z \end{pmatrix}, \tag{2.32a}$$

$$\mathbf{H} = \begin{pmatrix} h_x \\ h_y \\ H_0 - h_z \end{pmatrix}. \tag{2.32b}$$

Substituting these into Eq. 2.31 and keeping only terms first order in the small quantities, we arrive at

$$\frac{\partial \mathbf{M}}{\partial t} \approx \mu_0 \gamma \begin{pmatrix} m_y H_0 - h_y M_\mathrm{s} \\ h_x M_\mathrm{s} - m_x H_0 \\ 0 \end{pmatrix}. \tag{2.33}$$

As with the plane wave treatment above, we choose an oscillatory solution for $\mathbf{M}$ and $\mathbf{H}$ with a time dependence $e^{-i\omega t}$, so Eq. 2.33 becomes

$$-i\omega \begin{pmatrix} m_x \\ m_y \\ -m_z \end{pmatrix} = \mu_0\gamma \begin{pmatrix} m_y H_0 - h_y M_\mathrm{s} \\ h_x M_\mathrm{s} - m_x H_0 \\ 0 \end{pmatrix}. \tag{2.34}$$

One consequence of this is that, by assuming the magnetisation and field are mostly in the bias direction, we are able to neglect the variation in the z-axis magnetisation and field. This means we can write

$$\mathbf{M} = M_\mathrm{s}\hat{\mathbf{z}} + \mathbf{m}, \tag{2.35a}$$

$$\mathbf{H} = H_0\hat{\mathbf{z}} + \mathbf{h}, \tag{2.35b}$$

where $\mathbf{m}$ and $\mathbf{h}$ only have components in the plane perpendicular to the applied field. Rearranging Eq. 2.34, we can rewrite it in terms of a matrix equation:

$$\mu_0\gamma M_\mathrm{s} \begin{pmatrix} h_x \\ h_y \end{pmatrix} = \begin{pmatrix} \mu_0\gamma H_0 & -i\omega \\ i\omega & \mu_0\gamma H_0 \end{pmatrix} \begin{pmatrix} m_x \\ m_y \end{pmatrix}. \tag{2.36}$$

The problem has reduced to two dimensions since the small quantities $\mathbf{m}$ and $\mathbf{h}$ exist only in the plane perpendicular to the bias field. Inverting this matrix equation gives

$$\mathbf{m} = \bar{\chi}_\mathrm{p}\mathbf{h}, \tag{2.37}$$

where $\bar{\chi}_\mathrm{p}$, the *Polder susceptibility tensor*, is given by:

$$\bar{\chi}_\mathrm{p} = \begin{pmatrix} \chi_1 & i\chi_2 \\ -i\chi_2 & \chi_1 \end{pmatrix}, \tag{2.38}$$

with

$$\chi_1 = \frac{\Omega_\mathrm{H}}{\Omega_\mathrm{H}^2 - \Omega^2}, \tag{2.39a}$$

$$\chi_2 = \frac{\Omega}{\Omega_\mathrm{H}^2 - \Omega^2}, \tag{2.39b}$$

where $\Omega_\mathrm{H} = \frac{H_0}{M_\mathrm{s}}$ and $\Omega = \frac{\omega}{\gamma\mu_0 M_\mathrm{s}}$. This susceptibility describes how the magnetisation of the saturated material responds to small perturbing magnetic fields. Since the plane waves we are considering take the form of small perturbations, we can use this susceptibility and return to the plane wave equations to continue towards the magnetostatic approximation.

2.4.3 Plane waves again

Combining Eq. 2.30 with Eq. 2.3 and Eq. 2.38 [56] we arrive at

$$\left(\mathbf{k}\mathbf{k}^\mathbf{T} - k^2\bar{I} + k_0^2\bar{\mu}\right)\mathbf{h} = 0, \tag{2.40}$$

where we have replaced $\mathbf{H}$ with its small signal equivalent $\mathbf{h}$ and defined $\bar{\mu} = \bar{I} + \bar{\chi}_\mathrm{p}$. We have also used the definition of the free-space wavenumber $k_0 = \omega\sqrt{\mu_0\epsilon_0\epsilon_\mathrm{r}}$ from Eq. 2.29. This equation is of the form $\bar{A}\mathbf{h} = 0$ and only has non-trivial solutions if the determinant of the tensor $\bar{A}$ is equal to zero. Explicitly, this requirement is

$$\det\begin{pmatrix} k_0^2(1+\chi_1) - k_y^2 - k_z^2 & k_x k_y + ik_0^2\chi_2 & k_x k_z \\ k_x k_y - ik_0^2\chi_2 & k_0^2(1+\chi_1) - k_x^2 - k_z^2 & k_y k_z \\ k_x k_z & k_z k_y & k_0^2 - k_x^2 - k_y^2 \end{pmatrix} = 0. \tag{2.41}$$

We shall examine the solutions to this determinant in the special cases where propagation is parallel to the applied field ($\mathbf{k} = k\hat{\mathbf{z}}$), or perpendicular to the applied field ($\mathbf{k} = k\hat{\mathbf{x}}$).

- **Propagation parallel to applied field**

  Solving Eq. 2.41 for $\omega \neq 0$ leads to two possible dispersion relations for waves in the material:

$$k_+^2 = k_0^2(1 + \chi_1 + \chi_2), \tag{2.42a}$$

$$k_-^2 = k_0^2(1 + \chi_1 - \chi_2). \tag{2.42b}$$

- **Propagation perpendicular to applied field**

  Here, solving Eq. 2.41 for $\omega \neq 0$ also yields two solution branches:

$$k_\mathrm{a}^2 = \pm k_0^2, \tag{2.43a}$$

$$k_\mathrm{b}^2 = k_0^2\left(\frac{(1+\chi_1)^2 - \chi_2^2}{1+\chi_1}\right). \tag{2.43b}$$

Both Eq. 2.42 and Eq. 2.43 are plotted in Fig. 2.6. In each case, one of the two solutions is much more strongly affected by the medium than the other. These branches (k_+ and k_b) have dispersions differing significantly from the free-space dispersion $k_0 = \omega\sqrt{\mu_0\epsilon_0\epsilon_\mathrm{r}}$. Solutions for a general propagation angle are more complicated, but deform smoothly between these two limits as the field is rotated relative to the propagation vector. All exhibit the important characteristic of having two solutions: one very different from the free space dispersion, and one very similar to it.

We return to the general case by examining Eq. 2.30 and using Eq. 2.25b to replace $\mathbf{k}\cdot\mathbf{h}$ with $-\mathbf{k}\cdot\mathbf{m}$. Then we can rearrange for $\mathbf{h}$ to obtain:

$$\mathbf{h} = \left(\frac{k_0^2\bar{I} - \mathbf{k}\mathbf{k}^\mathbf{T}}{k^2 - k_0^2}\right)\mathbf{m}. \tag{2.44}$$

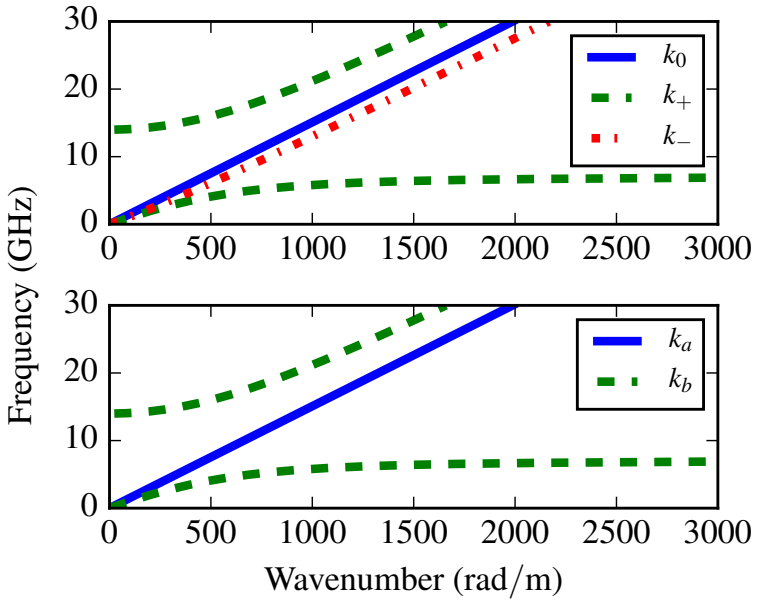

Figure 2.6: Dispersion relations for plane waves in a ferromagnetic material where the bias magnetic field is parallel (upper panel) or perpendicular (lower panel) to the propagation direction.

We can also get an expression for the small-signal electric displacement field $\mathbf{d}$ using Eqs. 2.25c, 2.25a and 2.25d:

$$\mathbf{d} = \left(\frac{\mu_0 \epsilon_0 \epsilon_\mathrm{r} \omega}{k_0^2 - k^2}\right) \mathbf{k} \times \mathbf{m}. \tag{2.45}$$

Since $\nabla \times \mathbf{h} = \frac{\partial \mathbf{d}}{\partial t} = -i\omega \mathbf{d}$:

$$\nabla \times \mathbf{h} = \left(\frac{-ik_0^2}{k_0^2 - k^2}\right) \mathbf{k} \times \mathbf{m}. \tag{2.46}$$

Above, we identified two possible solutions for the dispersion relation of plane waves in a magnetised material; a branch almost identical to normal plane waves in a vacuum, and a branch strongly affected by the magnetic dipoles in the material. It is the second of these two branches which represents the spin-wave dispersion. For the spin-wave branch, we notice that for most of the frequency

range the wavenumber is much larger than the free space wavenumber. This is an important hallmark of the spin wave – the wavelengths of the excitations are much shorter than a free space wave of the same frequency.

Having made this observation about the spin-wave wavelength, we now examine our expressions in the limit of $k \gg k_0$, and see that $\nabla \times \mathbf{h} \to 0$. However, since there are terms proportional to k^2 in the numerator of Eq. 2.44, $\mathbf{h}$ remains finite. Using this approximation we arrive at the characteristic equations of waves in this limit

$$\nabla \times \mathbf{h} = 0, \tag{2.47a}$$

$$\nabla \cdot (\bar{\mu}\mathbf{h}) = 0, \tag{2.47b}$$

$$\nabla \times \mathbf{d} = i\omega\epsilon_0\epsilon_\mathrm{r}\mu_0\bar{\mu}\mathbf{h}. \tag{2.47c}$$

The waves described by Eq. 2.47 are known as *magnetostatic waves*, since the time derivatives normally present in Maxwell's equations no longer appear in their characteristic equations. Only the term in $\mathbf{kk^T}$ in the expression for $\mathbf{h}$ in Eq. 2.44 survives in the magnetostatic limit, this represents the interaction of the spin wave with the dipole moments of the material's magnetisation. For this reason, magnetostatic spin waves are also known as *dipolar spin waves*.

In summary, the expressions in Eqs. 2.47a to 2.47c define the magnetostatic approximation. Combined with the Polder susceptibility tenor, these can be used to solve problems in which a saturated ferromagnetic material experiences small perturbations to its magnetisation from an external field. These spin waves have wavelengths shorter than the free space wavelength of light, but still much larger than the interatomic spacing. As such, their behaviour is dominated by the dipolar interactions between spins.

2.5 The different types of magnon

In the preceding sections, we treated separately the solutions for spin waves arising from the exchange interaction, and spin waves arising from the dipole interaction. In reality, of course, both types of interaction are present at all times, but which of these interactions dominates the behaviour depends on the wavelength of the excitation we are considering. We can identify four main propagation regimes in a material:

- **Electromagnetic waves** have wavelengths comparable to the free-space wavelength of electromagnetic radiation. Such waves are not strongly affected by the magnetisation of the medium and have a dispersion not significantly different from free-space radiation. They were seen as the weakly affected plane wave branch in Fig. 2.6.

- **Dipolar spin waves** are excitations with wavelengths shorter than free space wavelengths, but much larger than the interatomic spacing. For frequencies typically used in experiment, this translates to wavelengths of tens or hundreds of microns. They represent waves governed by the magnetostatic approximation, (Eq. 2.47) and do not include the effects of exchange. This is because the alignments of neighbouring spins in this regime are hardly perturbed by the presence of the excitation, and so the effect of the short-ranged exchange interaction on its total energy is minimal. Their characteristics are dominated by the magnetic dipole interaction.

- **Exchange spin waves** are those with wavelengths short enough that the exchange energy cost due to misalignment of neighbouring spins is much

larger than that of the dipole interaction integrated over the material. In this case, the phase difference between nearby spins can be significant. These are the type of excitation considered in Section 2.3.2, where the dispersion is governed by the strong, short-ranged exchange interaction.

- **Dipole-exchange spin waves** are spin waves with wavelengths intermediate between dipolar and exchange spin waves. In this regime, both dipole and exchange interactions between the spins must be considered. Unsurprisingly, accounting for both interactions simultaneously makes these the most theoretically involved type of spin wave to analyse [57].

In this thesis, experiments are carried out with excitation structures that are on a scale between microns and hundreds of microns. Typically, a magnetic structure will produce fields which change over a length scale comparable to its own size, so it cannot efficiently excite waves with wavelengths much smaller than its own dimensions. This means that dipolar spin waves are the primary excitations measured.

2.6 Spin waves in confined geometries

Having introduced spin waves and their important characteristics, we now examine their behaviour in experimental systems which are not infinite bulk materials As in the analysis of the preceding sections, we consider a saturated ferromagnetic material in a magnetic bias field. We retain the assumptions of electrical linearity and isotropy in the material. The application of a small oscillating magnetic field in the plane perpendicular to the bias field can cause the alignment of

spins to deviate slightly from the bias direction. The bias field exerts a torque on these misaligned spins according to Eq. 2.31, causing the spins to begin precessing around it.

After the treatment of Walker [58], and as in Section 2.4.2, we assume that the oscillating field which disturbs the spins is much smaller than the bias field. This allows us to assume deviations from the bias direction are small and perpendicular to the bias field, so that the total magnetic field inside the sample is given by:

$$\mathbf{H} = H_0\hat{\mathbf{z}} + \mathbf{h}, \tag{2.48a}$$

$$= \begin{pmatrix} h_x \\ h_y \\ H_0 \end{pmatrix}, \tag{2.48b}$$

with $h_x, h_y \ll H_0$. This allows us to employ the magnetostatic approximation derived in Section 2.4.3. Since $\nabla \times \mathbf{h} = 0$ we can define a magnetic potential through $\mathbf{h} = \nabla\Psi$. Substituting this into Eq. 2.47b we arrive at an equation for the potential inside the material which has a similar form to the Laplace equation

$$(1 + \chi_1)\left(\frac{\partial^2\Psi}{\partial x^2} + \frac{\partial^2\Psi}{\partial y^2}\right) + \frac{\partial^2\Psi}{\partial z^2} = 0, \tag{2.49}$$

where χ_1 has the same meaning as in the Polder susceptibility tensor in Eq. 2.38. Eq. 2.49 is known as the Walker equation from its original formulation [59]. Outside the material, the usual Laplace equation, $\nabla^2\Psi = 0$, holds.

In an experiment, samples of ferromagnetic materials have finite size and differing geometries. The boundary conditions imposed on the solutions of Eq. 2.49 by such geometries can significantly alter their behaviour. We shall examine two different geometries which are relevant to the experimental studies in this thesis – spheres and thin films.

2.6.1 Resonant modes of spheres

In the case that the waves are confined to a relatively small region (i.e. small compared to their wavelength or decay length), standing waves can exist. We consider a spherical geometry where the diameter of the sphere is small enough to meet this criterion, and examine the effect that the boundary conditions have on the possible resonances.

In a sphere, we must alter the definition of the field inside the material to account for the demagnetising field – the effect in which, due to boundary conditions or sample geometry, the magnetisation acts to reduce the field inside a material. In general, the demagnetising field can vary through an object, but for a spherical sample it is uniform and will reduce the field inside the sphere from H_0 to $H_0 - \frac{1}{3}M_\mathrm{s}$ [60]. This alters the definition of Ω_H in Eq. 2.49 such that $\Omega_\mathrm{H} \rightarrow \Omega'_\mathrm{H} = \frac{H_0}{M_\mathrm{s}} - \frac{1}{3}$.

We define a magnetic potential using the magnetostatic approximation. Inside the sphere, the potential obeys Walker's equation. Outside, it obeys the Laplace equation, which is identical to Walker's equation with $\Omega_\mathrm{H} \rightarrow 0$.

While solutions to the Laplace equation are well-known (see for example [61]), solutions to the Walker equation in a spherical geometry are complex, especially after the imposition of the relevant boundary conditions. More details of the solutions can be found in [62]. It is found that the allowed frequencies of the resonant modes obey the equation

$$n + 1 + \frac{\frac{d}{d\xi_0}P_n^m(\xi_0)}{P_n^m(\xi_0)} \pm m\chi_2 = 0, \tag{2.50}$$

where $\xi_0 = 1 + \frac{1}{\chi_1}$, $n \geq 1$ is an integer, and m is also an integer obeying $-n \leq m \leq n$. P_n^m are the associated Legendre polynomials. The solutions to this equation

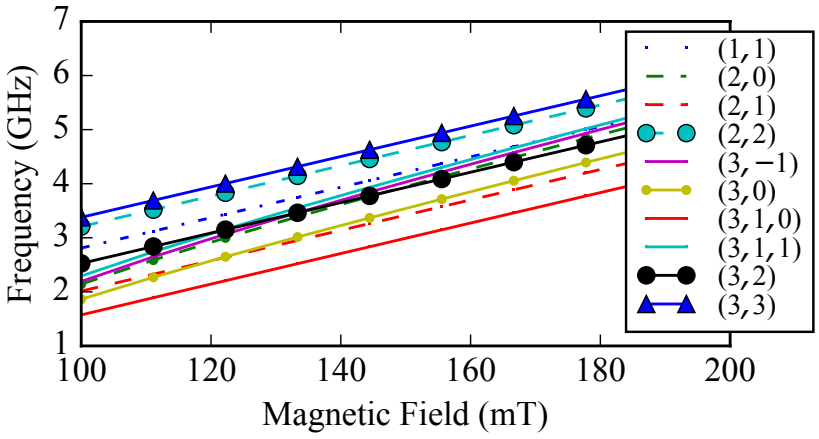

Figure 2.7: Frequencies of modes in the sphere up to $n = 3$. Most modes are linear in applied field with different offsets, but some are quadratic or higher order. In the absence of any magnetocrystalline isotropy, the intercept of the $(1, 1)$ mode is zero. There are two solutions for $n = 3$, $m = 1$.

are labelled with three indices (n, m, r) where n and m have the meaning already described, and r is used to distinguish the cases where multiple solutions exist for the same (n, m) pair. The third index r is often dropped when only a single solution exists. The $(1, 1)$ *ferromagnetic resonance* (FMR) mode, also known as the Kittel mode, is the lowest order standing mode in which all spins in the sphere precess in phase. It has a simple frequency dependence given by [58]

$$\omega_{\mathrm{FMR}} = \mu_0 \gamma H_0. \tag{2.51}$$

A plot of the magnetic field dependence of the frequencies of a variety of possible modes is shown in Fig. 2.7.

As well as the frequencies of allowed modes, we can also extract the relative phases of the spins at different points in the sphere. For a sphere magnetised

along the z-axis, the small-signal approximation to the magnetisation is

$$\mathbf{M}(t) = \begin{pmatrix} m_x e^{i\omega t} \\ m_y e^{i\omega t} \\ M_s \end{pmatrix}, \tag{2.52}$$

where M_s is the saturation magnetisation of the material and m_x and m_y are assumed to be much smaller than M_s. Inserting Eq. 2.52 into Eq. 2.31 and keeping only the terms first-order in the small quantities gives

$$m_x = \frac{h_x \Omega'_\mathrm{H} - i h_y \Omega}{\Omega'^2_\mathrm{H} - \Omega^2}, \tag{2.53a}$$

$$m_y = \frac{h_y \Omega'_\mathrm{H} + i h_x \Omega}{\Omega'^2_\mathrm{H} - \Omega^2}. \tag{2.53b}$$

The magnetic fields inside the sphere h_x and h_y can be obtained from the magnetic potential ($\mathbf{h} = \nabla\Psi$). Examples of the relative phases of spins for two different modes are illustrated in Fig. 2.8.

For modes other than the FMR, spins in different regions of the sphere are orientated in different directions. If the sphere is excited with a uniform field – for instance from an excitation structure much larger than the sphere itself – all the spins will be subject to the same direction of applied field. This makes exciting higher-order modes very difficult. To excite a variety of modes in the sphere therefore requires an excitation structure of a similar or smaller size than the sphere.

Additional effects in real materials

There are a variety of effects in real materials that can alter the behaviour of these modes from the idealised case described here. As the strength of the magnetic

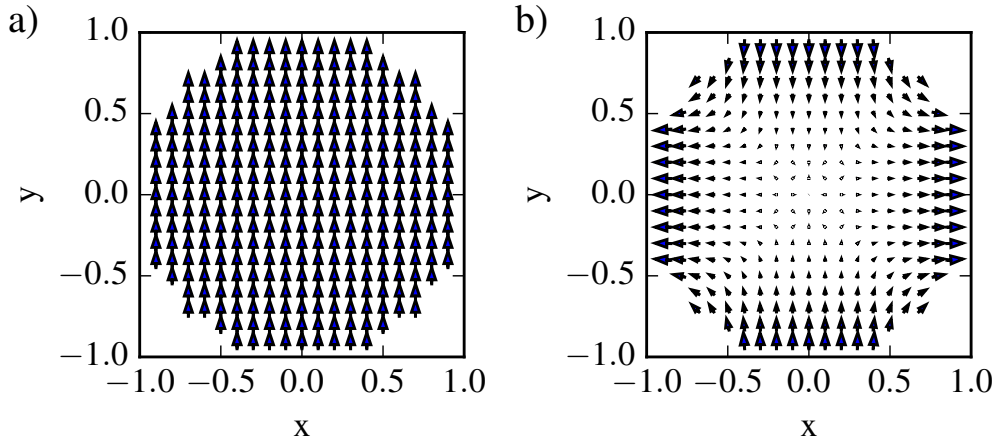

Figure 2.8: The relative orientation of the small-signal magnetisation for two resonant modes in a ferromagnetic sphere of unit radius in a bias field along the z-axis. (x, y) coordinates are relative to the sphere centre in an equatorial plane. a) The $(1, 1)$ mode, also referred to as the ferromagnetic resonance mode, or the Kittel mode. b) The $(4, 2)$ mode (right panel).

bias field increases, so too do the frequencies of the resonant modes. Eventually, the wavelengths of the spin waves reduce to similar scales as that of the sphere itself. In this case, adjustments must be made for their finite velocity [63, 64], which shifts the frequencies of the resonances. This effect tends to reduce the frequency of a resonant mode at high fields, and the shift due is stronger for lower-n modes [65].

Real materials also exhibit magnetic damping, which causes the precession of spins to decay. This is usually accounted for by introducing an additional term $\frac{\alpha}{M_s}\mathbf{M} \times \frac{\partial \mathbf{M}}{\partial t}$ to the magnetic torque equation [56]. The parameter α is a phenomenological constant describing the strength of the damping. The resulting equation is known as the Landau-Lifshitz-Gilbert (LLG) equation. This effect gives resonances a non-zero linewidth but, at least to first order, does not shift their frequencies.

Some materials also exhibit magnetocrystalline anisotropy. This property creates preferred directions in the material, breaking the isotropy of the plane perpendicular to the applied field. This will affect the frequencies and magnetisation profiles of the resonance modes. For instance, in the case of cubic anisotropy, aligning the bias field with the hard or easy axes of the crystal simply causes a constant shift to all of the resonance frequencies [66]. However, if the field is not aligned along one of these special directions, the general effects become too complicated to calculate analytically beyond the first few lowest-order resonances [67].

2.6.2 Propagating spin waves in thin films

In thin films, we use the same approach as for spheres; constructing a magnetic potential which obeys the Walker equation (Eq. 2.49) inside the film and the Laplace equation outside. The difference is that the planar geometry imposes different boundary conditions on the problem. In films, we often study propagating waves rather than resonant modes, since the structures are typically large in the dimensions other than the thickness. The characteristics of these propagating magnetostatic spin waves depend strongly on the relative orientation of the propagation direction, the magnetic bias field, and the film.

We consider an isolated slab of ferromagnetic material with the geometry shown in Fig. 2.9 and examine three special cases of the relative alignment of bias field and film. In all cases we assume that the thickness of the film is much smaller than the other two dimensions, so propagation of spin waves occurs only in the plane of the film. The dispersions for all three cases are plotted in Fig. 2.11. The

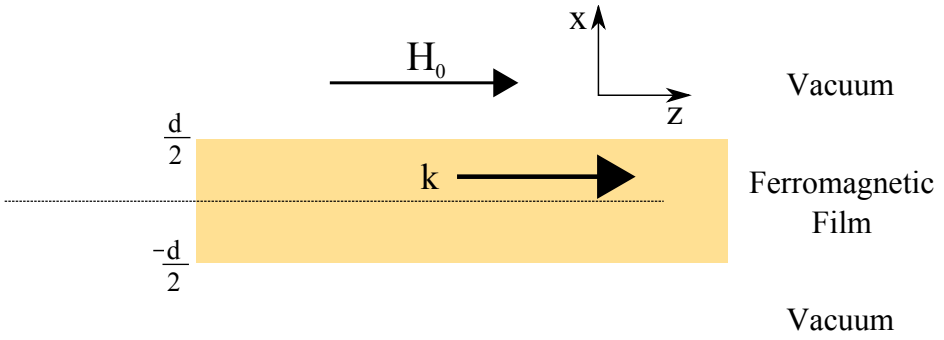

Figure 2.9: Schematic of a ferromagnetic film in the configuration supporting backward-volume waves. The film is assumed to be infinite in the z and y dimensions.

form of the magnetic potentials for the three cases is illustrated in Fig. 2.10.

Forward-volume magnetostatic spin waves

If the bias field is applied normal to the plane of an infinite film, then the propagation characteristics will be the same in all in-plane directions. The modes observed in this configuration are known as *forward-volume* waves because their group and phase velocities have the same direction, and their magnetic potential is distributed sinusoidally through the thickness of the film (see Fig. 2.10). Forward-volume waves propagate as plane wavefronts in the film, but the sinusoidal distribution of the magnetic potential in the thickness direction makes multiple harmonics possible, creating a potentially infinite spectrum of modes.

The full dispersion relation for all the forward-volume thickness modes is a transcendental equation whose solutions must be obtained numerically [68]. However, a useful approximation for the frequency of the lowest order mode is given by [69]

$$\omega = \mu_0\gamma\sqrt{H_\mathrm{i}\left(H_\mathrm{i} + M_\mathrm{s}\left(1 - \frac{1 - e^{-kd}}{kd}\right)\right)}. \tag{2.54}$$

In an experimental context, it is very often the lowest order mode which dominates the observed behaviour, making Eq. 2.54 valuable in the analysis of many systems. As in the treatment of spheres in Section 2.6.1, the internal magnetic field must be adjusted to account for the demagnetising field of the film. In this case, the appropriate substitution is $H_0 \rightarrow H_\mathrm{i} = H_0 - M_\mathrm{s}$.

The range of allowed frequencies for forward-volume waves is given by

$$\mu_0\gamma H_\mathrm{i} \leq \omega \leq \mu_0\gamma\sqrt{H_\mathrm{i}(H_\mathrm{i} + M_\mathrm{s})} \tag{2.55}$$

where the $k = 0$ (ferromagnetic resonance) frequency corresponds to the lower frequency limit.

Backward-volume magnetostatic spin waves

If the biasing magnetic field is applied in the plane of the film, then the spin-wave dispersion depends strongly on the propagation direction. *Backward-volume* waves are observed when the propagation direction is parallel to the applied field. They are so named because the group and phase velocities of the waves are in opposite directions.

As for forward waves, the full dispersion for a general thickness mode is obtained from a transcendental equation, but an approximation to the lowest order mode is given by [69]

$$\omega = \mu_0\gamma\sqrt{H_0\left(H_0 + M_\mathrm{s}\left(\frac{1 - e^{-kd}}{kd}\right)\right)}. \tag{2.56}$$

The associated magnetic potential is distributed sinusoidally through the film thickness, leading to a spectrum of thickness modes. Their 'backward' nature can lead to surprising results, such as negative Doppler shift [70] where the frequency

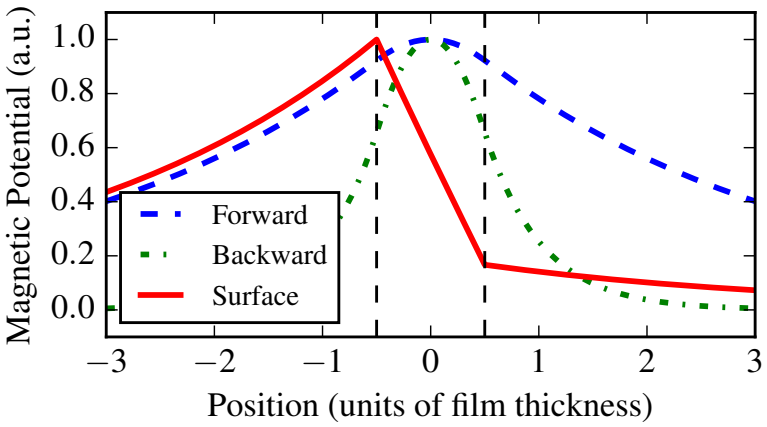

Figure 2.10: The potentials for the three special cases of spin waves which exist in ferromagnetic films. The film extends from $x = -0.5$ to $x = 0.5$. All three functions are plotted for an applied field of 150 mT in a film of YIG with $M_s = 196\,\mathrm{kAm^{-1}}$. The potential functions are for even modes in the case of backward and forward waves. They are plotted for $k = 6.9 \times 10^4$ in the case of forward and surface waves, and $k = 3.1 \times 10^6$ in the case of backward waves.

of radiation emitted by a source appears to decrease, rather than increase, as it moves towards an observer.

When the bias field is in the plane of the film, there is no demagnetising field, so the allowed range of backward-volume frequencies is

$$\mu_0\gamma H_0 \leq \omega \leq \mu_0\gamma\sqrt{H_0(H_0 + M_s)}. \tag{2.57}$$

However, in the case of backward-volume waves, the ferromagnetic resonance frequency is the upper limit in frequency due to the negative gradient of the dispersion relation.

Magnetostatic surface spin waves

Also called Damon-Eshbach modes after the authors of the paper which initially described their theory [68], *magnetostatic surface waves* are of different character

to the volume waves already described. They exist when the applied field is in the plane of the film, but the direction of propagation of spin waves is perpendicular to the applied field.

While the range of allowed frequencies for volume waves coincides with the allowed range for the bulk plane waves that were analysed in Section 2.4.3, surface waves exist entirely outside of this band. Their name arises from the fact that they are localised to the film surfaces, with a magnetic potential that decays exponentially from the surface rather than being sinusoidally distributed through the entire thickness, as seen in Fig. 2.10. Consequently, there is only a single surface-wave band rather than a spectrum of thickness modes, and an analytic expression exists for the spin-wave dispersion:

$$\omega = \mu_0\gamma\sqrt{H_0(H_0 + M_\mathrm{s}) + \frac{1}{4}M_\mathrm{s}^2(1 - e^{-2kd})}. \tag{2.58}$$

The allowed range of frequencies for surface waves is given by

$$\mu_0\gamma\sqrt{H_0(H_0 + M_\mathrm{s})} \leq \omega \leq \mu_0\gamma\left(H_0 + \frac{1}{2}M_\mathrm{s}\right). \tag{2.59}$$

Another interesting property of surface waves is their non-reciprocity: surface waves propagating in opposite directions are localised on opposite surfaces of the film. Depending on the experimental configuration, this can lead to different propagation characteristics in different directions.

Surface waves have a positive dispersion, like forward waves, and so the low frequency limit is also the ferromagnetic resonance frequency. Note also that there is no overlap between the backward and surface-wave bands. This means that as the bias field is rotated in plane, the frequency width of one band will grow while the other shrinks. This is different to the behaviour observed when rotating

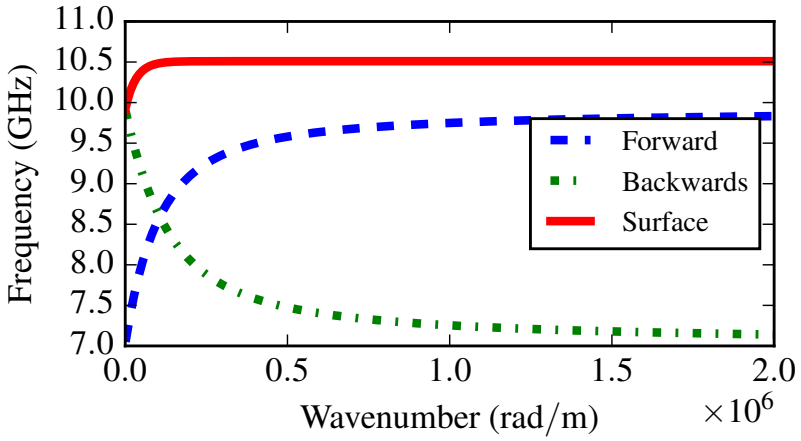

Figure 2.11: Dispersion relations for the three varieties of spin wave in ferromagnetic films. Dispersions are shown here for a film of 5 µm thickness and a saturation magnetisation of 196 $\mathrm{kAm^{-1}}$, the literature value for the material yttrium-iron garnet at 4 K [56]. An applied field of 250 mT is used for backward and surface waves, and a field of $(250 + \mu_0 M_\mathrm{s})$ mT for forward waves to compensate for the demagnetising field.

between forward and backward waves where one dispersion deforms smoothly into the other.

Additional effects in real materials

All of the dispersions described for the three configurations of spin waves above assume that the waves exist in an isolated slab of ferromagnetic material which has infinite extent in two dimensions. However, typical experimental geometries do not use films large enough to be considered infinite.

Spin waves are often excited in waveguides composed of thin films of a magnetic material which have one axis much longer than the other. The finite width along the shorter axis can distort the dispersion relation, or allow standing waves to form in the waveguide.

As in the spherical geometry, magnetocrystalline anisotropy will also play a

role and alter the propagation characteristics relative to a preferred axis.

Another possible complication is the effect of other dielectrics or ground planes in proximity to the film [71, 72] which alter the electromagnetic boundary conditions. This is relevant to many experimental situations, where sample holders might create a ground plane, or metallic layers may be deposited onto a ferromagnetic material. In Fig. 2.12, the dispersion relation for surface waves is plotted for the case where ground planes are present at unequal distance from each surface of a ferromagnetic film. Due to the non-reciprocity of surface waves and their localisation at the film surfaces, the effect of the ground planes is different depending on the direction of propagation. We note that as well as the distortion of the shape of the dispersion curve, the FMR frequency is shifted by the presence of the ground planes.

One ferromagnetic material particularly important to this thesis is yttrium-iron garnet ($Y_3Fe_5O_{12}$), films of which are usually grown epitaxially on substrates of gadolinium-gallium garnet ($Gd_3Ga_5O_{12}$). The presence of this substrate can introduce additional losses and different dielectric properties that affect spin-wave behaviour. Some of these effects will be explored in more detail in later chapters.

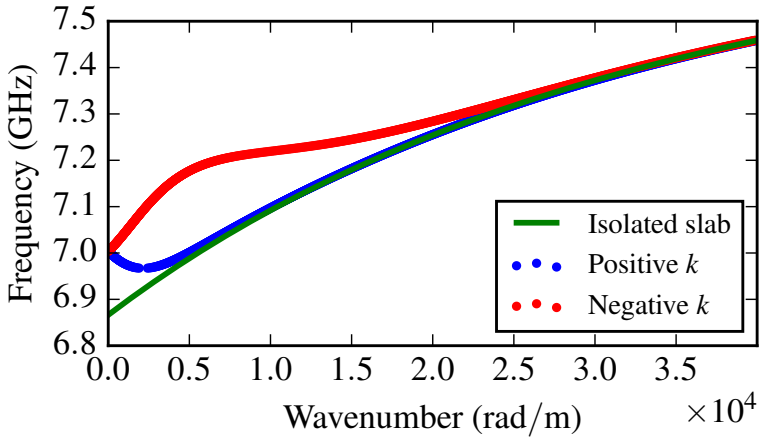

Figure 2.12: The surface-wave dispersion of a 5 µm film of yttrium-iron garnet in a bias field of 150 mT with ground planes 0.1 mm above the film and 0.5 mm below it. These spacings are smaller than those used in the experiments described in this thesis, and have been chosen to exaggerate the effect on the dispersion. Due to the asymmetry of the geometry, surface waves propagating in opposite directions are affected differently by the proximity of the ground planes. The dispersion for the isolated slab is also plotted for comparison.

Chapter 3

Experimental setup

The eventual aim of our research in quantum magnonics is to interface magnonic components with superconducting quantum bits, or qubits. Superconducting qubits have a number of features which lend them to this application. Their main advantage is that their design parameters are much more easily altered than other types of qubit. Since they are produced using lithography, their characteristics and their coupling to excitation and detection structures can be easily controlled during design and fabrication. However, the fact that magnons and superconducting qubits normally have frequencies in the microwave regime also makes it necessary to work at very low temperatures. At room temperature, typical thermal fluctuations have energies orders of magnitude larger than that of magnons or qubits, making microwave-frequency quantum experiments impossible.

Using the Boltzmann distribution, it is possible to calculate the expected occupancy of an energy level with a given equivalent frequency ($E = hf$) at a certain temperature. In Fig. 3.1 the average occupation is plotted as a function of frequency for three temperatures. For frequencies in the GHz regime it is evi-

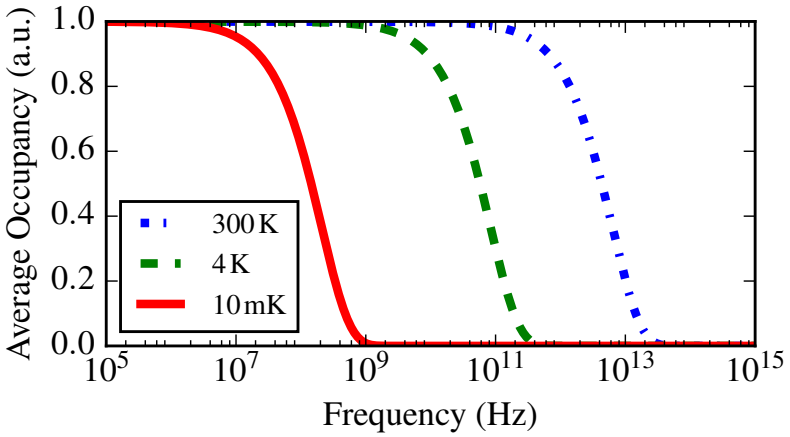

Figure 3.1: At room temperature, modes with equivalent frequencies up to 10^{13} Hz are still highly occupied, equivalent to radiation in the mid infra-red. The equivalent temperature for 4 GHz, roughly in the range of our experiments, is around 200 mK. Accordingly, only by reducing the temperature by a further order of magnitude, to approximately 10 mK, can the average occupancy of gigahertz frequency modes be brought close to zero.

dent that even liquid helium temperature is insufficient to bring the background thermal population close to zero. Temperatures on the order of 10 mK are necessary to ensure that thermal excitations are reduced to a level where experiments with single microwave excitations are practical. This makes the use of dilution refrigerators necessary for experiments with superconducting qubits.

3.1 Dilution refrigerator

A dilution refrigerator uses a mixture of the rare isotope helium-3 (He-3) and the more common helium-4 (He-4) to achieve cooling down to millikelvin temperatures [73]. Such refrigerators are typically organised in two stages; the first providing cooling down to approximately 4 K , and the second cooling from 4 K

to its base temperature. The first stage is achieved either by a bath of liquid He-4 or through the use of a closed cycle of compression and expansion of helium gas in a 'pulse tube'. The compression and expansion of gas in different regions of the tube transports heat from the cold end of the tube to the hot end [74]. The former are often colloquially known as 'wet' fridges and the latter as 'dry' fridges. Dry fridges have the advantage that they do not require a steady supply of liquid helium and liquid nitrogen for pre-cooling and maintaining the 4 K stage. However, careful design is necessary in dry fridges to avoid transmitting vibrations from the pulse tube to the sample.

The second stage, known as the dilution unit, is similar in both types of system and relies on two key properties of the mixture of He-3 and He-4: a unique form of phase separation at temperatures below the triple point of about 0.7 K, and the large difference in the vapour pressure of the two isotopes at the same temperature.

3.1.1 The dilution unit

A schematic of a dilution unit is shown in Fig. 3.2. Before entering the dilution unit the He-3/He-4 mixture is first liquefied. In a wet fridge this is achieved by heat exchange with a small reservoir of liquid He-4, above which the pressure is reduced to suppress its boiling point to around 1.5 K. In a dry fridge it can be provided by the second cold head of a two-stage pulse tube, typically at around 2.9 K. The mixture is forced into the mixing chamber through a flow impedance designed to keep the gas pressure in the condensing line relatively high. It is then cooled to around 0.7 K by a heat exchanger in contact with the still, and

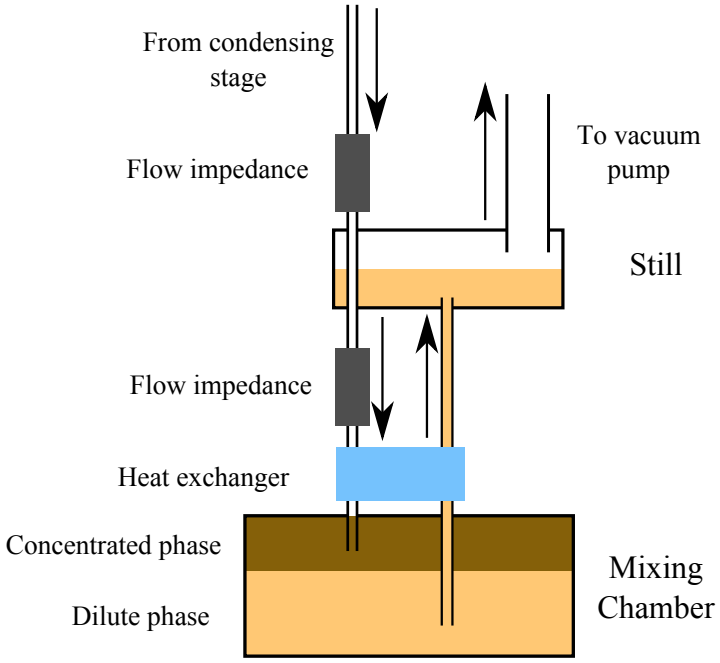

Figure 3.2: A schematic of the different parts of the lowest temperature stage of a dilution refrigerator. The still is typically held at around 0.7 K where the difference in vapour pressure of helium-3 and helium-4 is large. Depending on the exact design parameters of the system, temperatures in the mixing chamber can be continuously maintained at or below 10 mK.

passes through a second flow impedance to the mixing chamber. Before it enters, another heat exchanger cools the incoming liquid to near base temperature using the cold liquid leaving the mixing chamber.

The mixing chamber is maintained at the base temperature of the system, which can reach ten millikelvin or below depending on the design of the dilution unit. At these temperatures, the helium mixture undergoes a phase separation into a less dense, 'concentrated' phase composed almost entirely of He-3, and a more dense 'dilute' phase containing only about 6.6% He-3. A line placed at the bottom of the mixing chamber draws some of the more dense phase up into the still, pre-cooling the incoming mixture as it leaves.

The still temperature is around 0.7 K and at this temperature the vapour pressure of He-3 is much higher than that of He-4. This means that when gas is pumped out of the still by an external pump, mostly He-3 boils out of the mixture in the still. Since He-3 is being drawn out of the dilute phase, more must be drawn in to maintain the equilibrium between the two phases in the mixing chamber. As a result He-3 moves from the concentrated phase across the phase boundary into the dilute phase in the mixing chamber. Due to the entropy difference between the phases, this process is endothermic and causes cooling of the mixing chamber as energy is drawn from the surroundings to allow the He-3 to cross the boundary. Cooling power can be increased by applying a small amount of power to heat the still and more rapidly evaporate the He-3 out of the He-4.

Continuous cooling is therefore achieved as He-3 is drawn across the phase boundary in the mixing chamber due to the action of the vacuum pump on the mixture in the still. The He-3 pumped out of the still is continuously cycled back into the concentrated phase through the condensing line to ensure the correct ratio of phases in the mixing chamber is maintained. A useful conceptual analogy is to evaporative cooling – the evaporation of the He-3 from the still causes cooling of the remaining liquid. However, due to the quantum mechanical nature of the phase separation, the cooling occurs at the phase boundary in the mixing chamber, rather than at the surface of the liquid as would be the case for normal evaporation.

For the experiments in this thesis, an Oxford Instruments Triton 200 dilution refrigerator was used with an unloaded base temperature of approximately 7 mK. This is a 'dry' system with a pulse tube cooler providing cooling to 4 K. A picture

of the system is shown in Fig. 3.4a.

3.2 Sample space and microwave setup

All of the experiments reported in this thesis are conducted using microwave transmission measurements of various structures containing the ferrimagnetic insulator yttrium-iron garnet. More information about yttrium-iron garnet (commonly known as YIG) is given in Section 3.3.1. The apparatus used to conduct our measurements can be broadly divided into two parts: the sample to be tested (which changes for each experiment) and the microwave setup, which remains broadly the same for all experiments.

3.2.1 Sample setup

Samples to be tested in the fridge are mounted inside a copper sample box, such as that seen in Fig. 3.4b, which is thermally anchored to the mixing chamber stage of the fridge.

Microwave inputs are provided to the sample via MMPX connectors soldered to coplanar waveguide structures on a printed circuit board (PCB) inside the sample box. The centre lines of the PCB coplanar waveguides are wire-bonded to the excitation and detection structures on a chip in the centre of the box. The YIG is placed on top of the chip in appropriate positions relative to the excitation and detection structures.

As discussed in Chapter 2, an external magnetic field must be applied to the magnetic material in order to allow the spin waves of interest to propagate. This field is supplied by a home-built superconducting magnet shown in Fig. 3.4c. The

magnet is fixed to the 4 K stage of the dilution refrigerator to keep it below the transition temperature of the niobium-titanium wire from which it is wound. A Helmholtz arrangement of two complementary sets of windings on the magnet provides a highly uniform magnetic field. A calculation of the field around the magnet is plotted in Fig. 3.3. The field inside a cylindrical region of diameter and height 20 mm at the centre of the magnet varies by less than 0.2%. The magnet provides a magnetic field in its central region of 79 mT/A. This was determined using a measurement of the ferromagnetic resonance (FMR) frequency of a sphere of yttrium-iron garnet as a function of the current in the magnet. Since the FMR frequency of a YIG sphere increases at 28 GHz/T (see Eq. 2.51), measuring the frequency change as a function of current gives an equivalence between magnet current and magnet field. Using a current-limited power supply to pass a current through the superconducting coil, we can apply fields up to

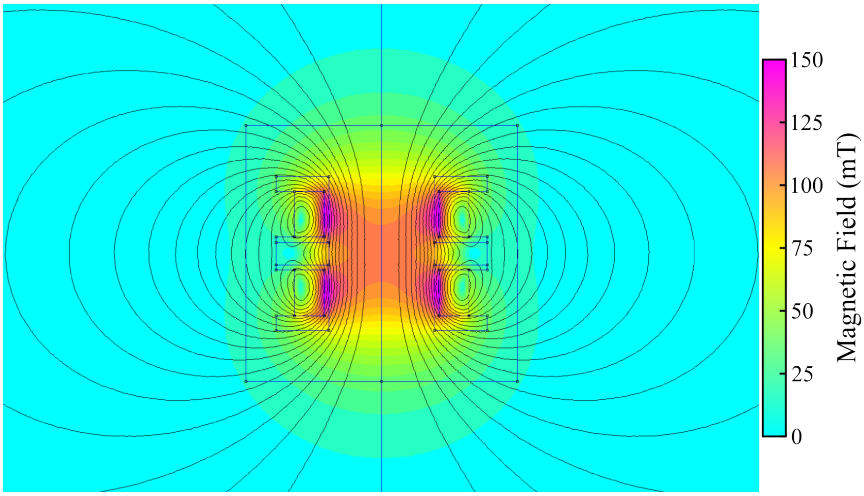

Figure 3.3: A simulation of the shape of the magnetic field produced by the superconducting magnet, when the field in the sample space is approximately 100 mT. The calculation was performed in FEMM [75], a finite-element magnetostatics solver. The Helmholtz arrangement of coils provides a uniform field in the bore of the magnet.

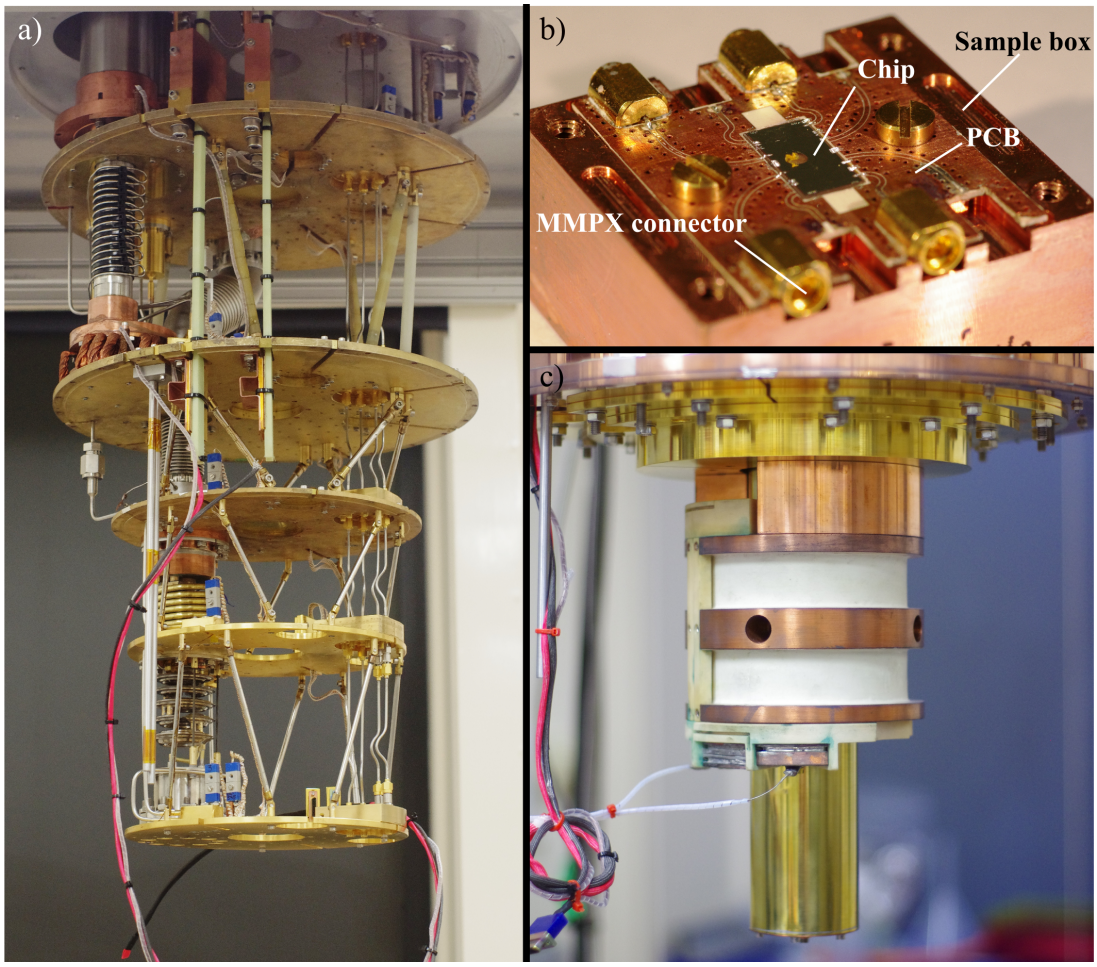

Figure 3.4: a) The Triton 200 dilution refrigerator with all of its radiation shields removed. No sample is mounted in this picture. b) An open sample box being prepared to go into the fridge. MMPX connectors, soldered to the PCB at the edges of the sample box, are used to provide microwave inputs to the system under test. A chip containing excitation and detection structures is glued in the centre of the PCB and connected to it using wire bonds. The magnetic samples are then glued on top of the chip. c) The superconducting magnet mounted in the fridge. It hangs below the lowest plate shown in a.

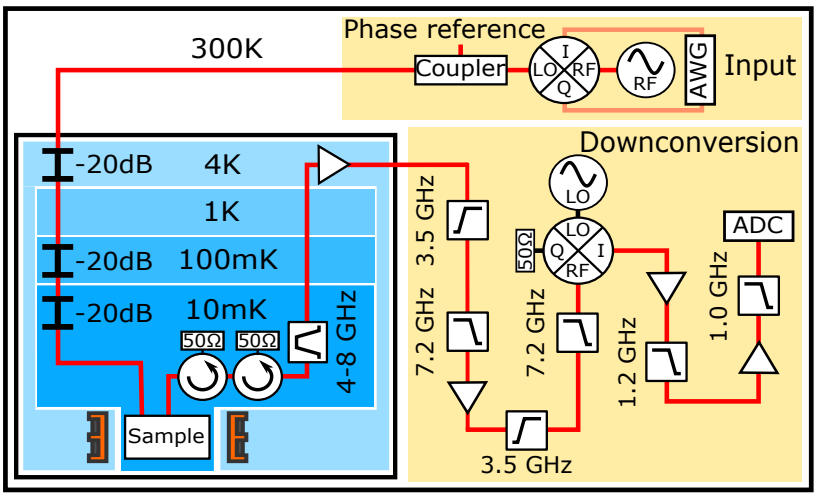

Figure 3.5: A schematic of the microwave setup used for experiments in the fridge. The input can be in the form of continuous microwaves or shaped pulses. To measure the phase of signals in continuous wave measurements, a -10 dB directional coupler splits off a fraction of the input signal before it enters the fridge. The carrier frequency of the output signals is reduced through heterodyne down-conversion and the signals are then digitised by a 2.5 GHz analogue-to-digital converter.

400 mT with a resolution of approximately 10 μT.

3.2.2 Microwave setup

A schematic of the microwave setup is shown in Fig. 3.5. It can be roughly divided into three parts: generation of the input signals, cabling to and from the sample in the fridge, and capture and processing of the output signals. Both the input generation and output processing make use of microwave mixers. A microwave mixer combines two signals by multiplying them together, an example of which is given in Fig. 3.6a. In the frequency domain, this corresponds to a convolution of the two signals. A mixer has three ports; the local oscillator (LO) and radio-frequency (RF) ports operate at high frequencies, typically in the GHz range, and the intermediate frequency (IF) port operates at lower frequencies,

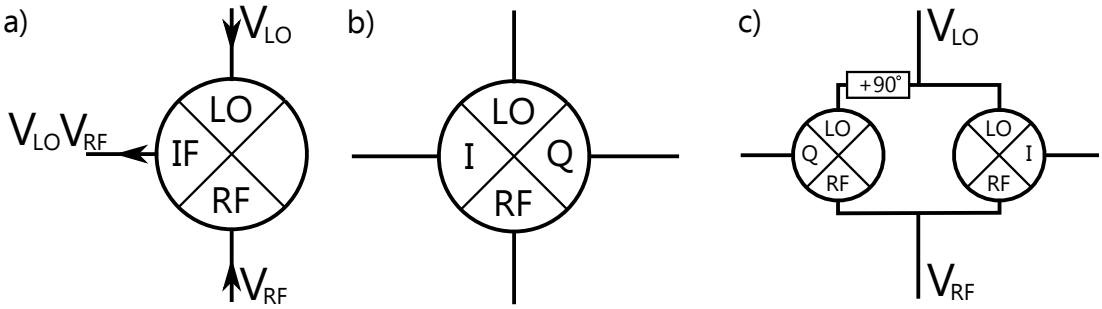

Figure 3.6: a) A microwave mixer multiplies an RF input with a reference LO signal to produce its output. In this diagram the mixer is being used for downconversion. b) An IQ mixer produces two outputs which are, ideally, 90° out of phase with one another. c) A schematic of an IQ mixer, which is composed of two single-output mixers with the LO of one phase shifted by 90°.

typically DC to hundreds of MHz [76]. Two ports are used as inputs, one of which is always the LO port, and the third provides the output. A high-power, continuous-wave signal is put into the LO port. If input is provided to the IF port, the RF port is the output and $V_{\mathrm{RF}} = V_{\mathrm{IF}}V_{\mathrm{LO}}$. This mode of operation is known as 'upconversion', since it is usually used to add a carrier frequency from the LO port to a low-frequency signal from the IF port. If input is provided to the RF port, then the IF is the output and $V_{\mathrm{IF}} = V_{\mathrm{RF}}V_{\mathrm{LO}}$. This mode is known as 'downconversion', since it can be used to reduce the carrier frequency of a signal. Both of these modes of operation are examined in more detail in the following sections. An IQ mixer (as shown in Fig. 3.6b and c has two IF ports, labelled I and Q, which are 90° out of phase with one another. The use of this type of mixer will also be examined in a later section.

Input generation

A continuous-wave microwave source with variable output frequency (1 to 40 GHz) and power is used to generate input signals. To measure the complex transmission, where the phase as well as the amplitude of the transmission is to be determined, a -10 dB directional coupler is used to divert 10% of the input signal to the data capture instrumentation as a phase reference. The details of this are explained further below.

For some experiments, the continuous-wave signal is shaped into pulses using an arbitrary waveform generator (Textronix AWG520) and IQ microwave mixer. The AWG has a clock frequency of 1 GHz and can generate waveforms with rise times as short as 2.5 ns. Its output is used as the input to the IF ports of an IQ microwave mixer, so that the continuous output of the microwave source can be shaped into an arbitrary envelope.

Given a time-dependent pulse shape produced by the AWG, V_{AWG}, we have

$$V_{\mathrm{RF}} = V_{\mathrm{AWG}} V_{\mathrm{LO}} \tag{3.1a}$$

$$= V_{\mathrm{AWG}}(t)\cos(\omega_{\mathrm{LO}} t). \tag{3.1b}$$

In reality, however, the isolation between mixer ports is never perfect; some signal will always leak from the LO and IF ports to the RF port, resulting in a small amount of uncombined signal on the RF port. A more realistic output is

$$V_{\mathrm{RF}} = V_{\mathrm{LO}} V_{\mathrm{AWG}} + \epsilon_1 V_{\mathrm{LO}} + \epsilon_2 V_{\mathrm{AWG}}, \tag{3.2a}$$

$$= V_{\mathrm{AWG}}(t)\cos(\omega_{\mathrm{LO}} t) + \epsilon_1 \cos(\omega_{\mathrm{LO}} t) + \epsilon_2 V_{\mathrm{AWG}}(t), \tag{3.2b}$$

where the parameters ϵ_1 and ϵ_2 represent the leakage from the LO and IF ports respectively. The leakage of the AWG signal from the IF port can be ignored

because it is at a relatively low frequency, outside of the band of possible excitations, and will be filtered out by the sample and other components such as DC blocks in the fridge. However, the leakage of the LO signal can be a problem because it means that there is a constant input to the experiment at frequencies which could cause unwanted excitation of the sample. To counteract this, a small DC offset is added to the AWG output, so $V_{\mathrm{AWG}}(t) \to V_{\mathrm{AWG}}(t) - V_{\mathrm{DC}}$. Then

$$V_{\mathrm{RF}} = V_{\mathrm{LO}}(V_{\mathrm{AWG}} - V_{\mathrm{DC}}) + \epsilon_1 V_{\mathrm{LO}}, \tag{3.3a}$$

$$= V_{\mathrm{AWG}}(t)\cos(\omega_{\mathrm{LO}}t) + (\epsilon_1 - V_{\mathrm{DC}})\cos(\omega_{\mathrm{LO}}t). \tag{3.3b}$$

If $V_{\mathrm{DC}} = \epsilon_1$ it will cancel with the leakage signal. The correct value of V_{DC} is determined by measuring the output of the mixer with $V_{\mathrm{AWG}}(t) = 0$ and sweeping V_{DC} to minimise V_{RF}.

As shown in Fig. 3.6c, an IQ mixer is composed of two mixers in parallel, with I and Q ports being the IF ports of the two component mixers. The DC offset to both I and Q ports must be determined separately, since even for high-quality mixers the characteristics of the two will differ slightly. The optimum offset voltage will also vary as a function of the LO frequency, so in experiments where several significantly different carrier frequencies are used the offsets applied to the mixer must be recalibrated for each frequency.

Cabling in the fridge

Any signal will carry an electrical noise spectrum reflecting the temperature at which it was generated. Thermal fluctuations cause random motion of the charge carriers in a resistive material which results in small voltage fluctuations. This type of noise is known as Johnson-Nyquist noise. As would be expected, the

average value of the voltage produced is zero, but the rms value is non-zero and can be calculated from the black-body radiation spectrum [76]

$$V_{\mathrm{RMS}}^2 = \frac{4hfBR}{e^{\frac{hf}{k_{\mathrm{B}}T}} - 1}, \tag{3.4}$$

where f is the centre frequency of the measurement band, B its bandwidth, and R the resistance of the component. For microwave frequencies in a sample at any temperature above about 1 K, $hf \ll k_{\mathrm{B}}T$ so Eq. 3.4 can be simplified to:

$$V_{\mathrm{RMS}}^2 \approx 4k_{\mathrm{B}}TBR. \tag{3.5}$$

From Eq. 3.5 we see that the square of the rms noise voltage (which is proportional to the noise power) is directly proportional to the temperature of the resistor in the microwave regime. We can therefore easily characterise the noise power produced by a component with an effective temperature. If $hf > k_{\mathrm{B}}T$, then Eq. 3.5 will give a higher noise power than the correct expression in Eq. 3.4, so the linear approximation will always overestimate the noise. In the following, we assume that the bandwidth and resistance in Eq. 3.5 are constant, while the temperature varies.

The temperature of the noise added to a signal by passive components such as resistors will be determined by the thermodynamic temperature of the component. However, for active components such as sources and amplifiers the temperature of the noise can easily be greater than the thermodynamic temperature.

The input signals from the room-temperature microwave setup will carry a noise power which at least equal to the temperature at which it was generated. If high-quality sources are used, the noise temperature should be only slightly above room temperature. Were this signal to be transmitted unaltered to the cold

stages of the fridge, its electrical noise temperature would be much greater than the thermodynamic temperature of the sample, and would lead to large output voltage fluctuations. To prevent this, the input signal is heavily attenuated as it is transmitted down to the sample.

Looking at Fig. 3.5, we see that the input signal will contain noise from the room temperature setup and the attenuators at the 4 K and 100 mK stages. To this will be added the noise due to components on the mixing chamber itself. If the signal into an attenuator has an amplitude V_s, then the output signal will be $\sqrt{A}V_\mathrm{s}+\sqrt{1-A}V_\mathrm{n}$, where A is the attenuation factor and V_n is a noise amplitude. This means that the effective temperature of the noise added by the attenuator will be $T_\mathrm{eff} = (1-A)T_\mathrm{A}$ where T_A is the temperature of the noise signal V_n, which is simply given by the thermodynamic temperature of the attenuator. In our setup we use $-20\,\mathrm{dB}$ attenuators, for which $A = \frac{1}{100}$, so $T_\mathrm{eff} \approx T_\mathrm{A}$.

Returning to the setup in Fig. 3.5, the room temperature noise passes through three $-20\,\mathrm{dB}$ attenuators, so is reduced in power by a factor 10^6. Similarly the 4 K noise is attenuated by 10^4 and the 100 mK noise by 10^2. We can then approximate the total noise temperature as:

$$T_\mathrm{tot} = (300\,\mathrm{K}\times 10^{-6}) + (4\,\mathrm{K}\times 10^{-4}) + (100\,\mathrm{mK}\times 10^{-2}) + 10\,\mathrm{mK}, \tag{3.6a}$$

$$= 0.3\,\mathrm{mK} + 0.4\,\mathrm{mK} + 1\,\mathrm{mK} + 10\,\mathrm{mK}, \tag{3.6b}$$

$$\approx 12\,\mathrm{mK}. \tag{3.6c}$$

The dominant source of noise experienced by the sample will therefore be the noise generated by the components on the mixing chamber. The room temperature signals are generated at high enough power that the signal to noise ratio at the sample is still high even after attenuation.

Another factor to consider is the thermal isolation between different parts of the dilution refrigerator. In Fig. 3.2 we saw that helium gas must be exchanged between stages at different temperatures, and in Fig. 3.4a it is obvious that different parts of the fridge at different temperatures must be physically connected. If the heat conducted from a higher temperature stage to a lower one were to exceed the cooling power at the lower temperature, the fridge would not be able to reach its base temperature. Care must be taken, therefore, that components mounted in the fridge do not create good thermal contact between stages at different temperatures, for instance by making contact with the radiation shields.

This consideration is particularly important in the case of the coaxial cables which carry signals from the room temperature microwave setup to the sample. It would be desirable to use coaxial cables made from materials which are good electrical conductors, so that any losses in signal, especially on the output lines, are minimised. However, for most materials, good electrical conductivity correlates with good thermal conductivity. This means a compromise must be made between between thermal resistance and electrical conductivity. The coaxial cables connecting different stages of the fridge are therefore made with an outer conductor of stainless steel, which has low thermal conductivity, and an inner conductor made from copper, which has high electrical conductivity. This reduces the heat conducted by the lines at the cost of increased attenuation.

On the input line, increased loss is not too problematic, since the signal is already heavily attenuated to reduce the noise temperature. However, on the output line, cables with high attenuation can mean the loss of valuable signal strength. For the initial section of the output line between the sample and the first amplifier, this problem is avoided in our setup by using cables made from a

superconducting niobium-titanium alloy. Superconductors are unique in combining low thermal conductivity with high electrical conductivity, but the cables can only be used where they are at temperatures below the transition temperature of the superconductor. After this section of superconducting cable, the signal is amplified at 4 K by a low-noise HEMT amplifier with a characteristic noise temperature of less than 5 K. A pair of circulators between the sample and the amplifier act as isolators to prevent any thermal noise from the amplifier being transmitted back into the sample box. After amplification, the signal is carried to the data capture apparatus at room temperature.

Heterodyne downconversion and data capture

Magnetostatic magnons typically have frequencies in the microwave regime, and the experiments in this thesis are conducted on magnon systems operating at frequencies between 4 and 8 GHz. According to the Nyquist-Shannon sampling theorem, a signal with its highest frequency components at f must be sampled at a rate of at least $2f$ in order to completely resolve the signal [77]. Sampling at any lower rate can lead to an inaccurate representation of the original signal due to aliasing. For our experiments, this would mean that directly capturing the signal on the output line would require an analogue-to-digital converter (ADC) with a sampling rate of around 15 GHz, a very demanding requirement.

Instead, we use heterodyne downconversion to reduce the frequency of the carrier signal to a range which is accessible to an ADC. This is achieved through the use of a microwave mixer such as the one depicted in Fig. 3.6.

Analogue downconversion Consider a signal consisting of an envelope function with a carrier frequency of ω_{RF}. We can write this as a sinusoidal voltage with a time-dependent amplitude:

$$V_{\mathrm{i}} = A(t)\sin(\omega_{\mathrm{RF}}t), \tag{3.7}$$

where the initial phase of the signal is defined as zero for simplicity. This signal passes into the experimental apparatus where it is modified by interaction with the sample. It is then bandpass filtered to remove noise outside the measurement band, and arrives at the data capture instruments with the form

$$V_{\mathrm{o}} = A'(t)\sin(\omega_{\mathrm{RF}}t + \phi_{\mathrm{e}}). \tag{3.8}$$

The shape of the envelope has been modified, $A(t) \rightarrow A'(t)$, and a phase shift ϕ_{e} has been introduced. It is this modification of the envelope and phase shift which we wish to use to study the behaviour of the sample, so we would like to extract them from the captured data. Inside the mixer, the output signal will be multiplied with another signal whose frequency ω_{LO} is close to the carrier frequency. The resulting output from the mixer is given by

$$V_{\mathrm{m}} = A'(t)\sin(\omega_{\mathrm{RF}}t + \phi_{\mathrm{e}})\sin(\omega_{\mathrm{LO}}t + \phi_{\mathrm{m}}), \tag{3.9}$$

where we have allowed for a potentially non-zero phase of the mixing signal, ϕ_{m}. Eq. 3.9 can be rewritten as

$$V_{\mathrm{m}} = \frac{1}{2}A'(t)(\cos[(\omega_{\mathrm{RF}} - \omega_{\mathrm{LO}})t + \phi_{\mathrm{e}} - \phi_{\mathrm{m}}] - \cos[(\omega_{\mathrm{RF}} + \omega_{\mathrm{LO}})t + \phi_{\mathrm{e}} + \phi_{\mathrm{m}}]). \tag{3.10}$$

We see that the output of the mixer provides two versions of the input signal, one with a carrier frequency of $\omega_{\mathrm{IF}} = \omega_{\mathrm{RF}} - \omega_{\mathrm{LO}}$ and one at the sum frequency of

$\omega_{\mathrm{RF}} + \omega_{\mathrm{LO}}$. Passing the signal through a low-pass filter with a cut-off frequency between ω_{IF} and ω_{LO} eliminates the high-frequency component of the signal as well as any potential leakage through the mixer. This means that by choosing the mixing frequency ω_{LO} relative to the input signal frequency ω_{RF} we can reduce the signal to a lower carrier frequency ω_{IF} which is accessible to the ADC, whilst still maintaining access to the envelope in which we are interested. These steps are illustrated in panels a and b of Fig. 3.7. Ideally, we would choose this lower carrier frequency such that $\omega_{\mathrm{IF}} = \frac{1}{4}\omega_{\mathrm{SF}}$, where ω_{SF} is the sampling frequency of the ADC. This is because the maximum frequency resolvable by the ADC is, according to the Nyquist-Shannon theorem, $\frac{1}{2}\omega_{\mathrm{SF}}$. Centring the signal in the middle of the frequency range available to the ADC will place it as far away as possible in frequency from any low-frequency noise, or aliased higher-frequency signals. This, in turn, will make it easier to separate out these sources of noise from the signal through filtering and digital down-conversion. In our setup we instead use $\omega_{\mathrm{IF}} = \frac{1}{5}\omega_{\mathrm{SF}}$. This is because the ADC sampling frequency can only be set to values $f_{\mathrm{s}} = \frac{5}{2^n}\,\mathrm{GHz}$. The IF frequency must be an integer multiple of the repetition rate of the experiment, such that subsequent measurements do not interfere destructively when averaged. For the sake of simplicity, we would also often like to use repetition rates of $2^m\,\mathrm{MHz}$. This means that, for instance, if we have $\omega_{\mathrm{SF}} = 2.5\,\mathrm{GHz}$ and $\omega_{\mathrm{IF}} = \frac{1}{4}\omega_{\mathrm{SF}} = 625\,\mathrm{MHz}$, then any repetition rate above 1 MHz must be carefully chosen to avoid destructive interference between measurements. Using $\omega_{\mathrm{IF}} = \frac{1}{5}\omega_{\mathrm{SF}} = 500\,\mathrm{MHz}$ gives us the freedom to also choose a repetition rate of 2 MHz or 4 MHz if desired.

The microwave source providing the signal at ω_{LO} is frequency locked to the input source and all other instruments by a rubidium clock and distribution

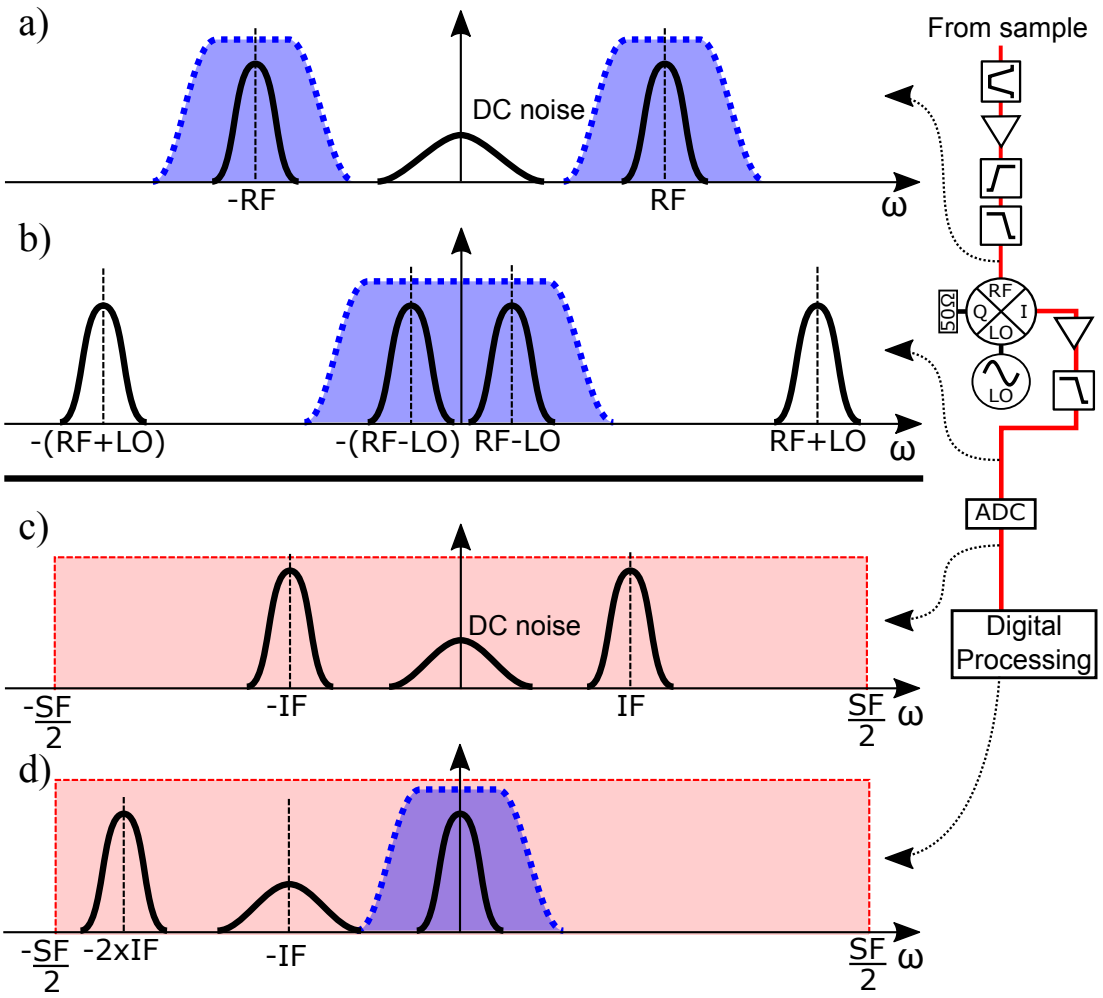

Figure 3.7: The steps in the downconversion process depicted in frequency space. On the right-hand side, a simplified schematic of the downconversion setup from Fig. 3.5 shows which part of the chain each step refers to. a) The signal from the sample is bandpass filtered around ω_{RF} to remove noise outside the measurement band. b) Analogue mixing of the input signal with the LO tone produces two images of the signal at the sum and difference frequencies. The sum frequencies are rejected by a lowpass filter. c) The signal is digitised in preparation for further digital processing. Only frequencies up to $\frac{1}{2}\omega_{\mathrm{SF}}$ can be captured. Higher frequency signals will be aliased back into the window. Additional noise can enter the signal at this stage, especially at low frequencies, added, for example, by the ADC. d) Digital mixing is done with both an I and Q channel, which are combined to produce the output. Producing the final signal in this way means that the signal is shifted in frequency, rather than producing two images at both positive and negative frequencies. The final output is digitally lowpass filtered to remove the high-frequency image and upconverted DC noise.

amplifier. This ensures that the signals produced by all parts of the apparatus are exactly synchronised. Without this frequency locking, it would be impossible to perform measurements with high numbers of averages. Any drift in phase or frequency of one source relative to the other would cause destructive interference of the signals over many repetitions.

After these analogue downconversion and filtering steps, the signal entering the ADC is

$$V_{\mathrm{ADC}} = \frac{1}{2} A'(t) \cos(\omega_{\mathrm{IF}} t + \phi_{\mathrm{e}} - \phi_{\mathrm{m}}). \tag{3.11}$$

Since the carrier frequency ω_{IF} is chosen to be less than half the sampling frequency, it can be accurately digitised by the ADC.

Averaging After the downconversion stage, the ADC is typically used to perform a high number of averages of the signal. The measurement is repeated many times and the results of all the measurements added together in the ADC memory. This reduces the amplitude of the noise on the signal by a factor $\sqrt{N}$ (with N the number of averages), since random fluctuations have a zero expectation value and will average to zero over many repetitions. However, it also has two further, less obvious effects which limit the types of measurement possible with the setup. Firstly, it removes contributions from any incoherent processes in the sample. Some processes, such as inelastic scattering, have no fixed phase relationship between incoming and outgoing excitations, so the phase becomes randomised. Any such signal will not contribute to the final averaged signal as the random initial phase means that signals from multiple measurements will cancel out. Secondly, it removes non-linear contributions to the signal, where an input at one frequency produces an output at another frequency. This is, again,

due to the phase relationship between subsequent measurements. The repetition rate of measurements is chosen to be equal to an integer number of periods of the IF frequency, so that measurements add constructively when summed in the ADC memory. This will in general not be the case for signals of other frequencies, so unless the output frequency is, after downconversion, also correctly related to the measurement repetition rate, each new measurement will have a different phase, and will cancel with previous measurements.

Digital downconversion After averaging, the final signal acquisition step is digital downconversion to remove the carrier frequency entirely, leaving only the envelope function $A'(t)$. The process is illustrated in panels c and d of Fig. 3.7. To recover the envelope, we must produce two versions of the signal, which are analogous to the I and Q channels on an analogue IQ mixer, as shown in Fig. 3.6b. The first version of the signal is

$$V_{\mathrm{Q}} = \frac{1}{4} A'(t) \cos(\phi_{\mathrm{e}} - \phi_{\mathrm{m}} - \phi_{\mathrm{d}}), \tag{3.12}$$

where ϕ_{d} allows for the phase of the digital mixing signal. This voltage has no time dependence other than that given by the envelope function, but is still multiplied by a cosine of the phase factors accumulated from the sample and the signal processing. Some of these phase factors are essentially random, depending as they do on the relative timings of the original input and the moment the first measurement is made. To eliminate them, we also multiply a copy of the signal by another sinusoidal wave $\sin \omega_{\mathrm{IF}} t$ which is 90 degrees out of phase with the first. This gives a second DC signal

$$V_{\mathrm{I}} = \frac{1}{4} A'(t) \sin(\phi_{\mathrm{e}} - \phi_{\mathrm{m}} - \phi_{\mathrm{d}}). \tag{3.13}$$

These two channels can be combined to give the envelope function independent of the phase factors

$$A' = 4\sqrt{V_{\mathrm{I}}^2 + V_{\mathrm{Q}}^2}. \tag{3.14}$$

In some situations, it is desirable to know the phase introduced to the signal by the sample as well. This can also be extracted from the I and Q voltages

$$\phi_{\mathrm{s}} = \phi_{\mathrm{e}} - \phi_{\mathrm{m}} - \phi_{\mathrm{d}} = \arctan\left(\frac{V_{\mathrm{I}}}{V_{\mathrm{Q}}}\right). \tag{3.15}$$

However, it is difficult to disentangle which part of the phase is due to the sample, and which has been introduced by the measurement apparatus. To do so requires comparison to a reference signal, which is split off from the original input before it enters the fridge, as seen in Fig. 3.5. Passing this signal through the same process of analogue and digital downconversion produces reference I and Q signals

$$V_{\mathrm{Qref}} = \frac{1}{4}A(t)\cos(\phi_{\mathrm{m}} + \phi_{\mathrm{d}}), \tag{3.16a}$$

$$V_{\mathrm{Iref}} = -\frac{1}{4}A(t)\sin(\phi_{\mathrm{m}} + \phi_{\mathrm{d}}), \tag{3.16b}$$

then

$$\phi_{\mathrm{ref}} = -(\phi_{\mathrm{m}} + \phi_{\mathrm{d}}) = \arctan\left(\frac{V_{\mathrm{Iref}}}{V_{\mathrm{Qref}}}\right), \tag{3.17}$$

and the experimental phase shift can be calculated from $\phi_{\mathrm{s}} - \phi_{\mathrm{ref}}$.

Another factor to consider is the fact that the digital filtering frequency must be chosen to remove the high-frequency image of the signal without affecting the envelope function $A'(t)$ which is itself a function of time. In our experiments, the bandwidth of the AWG which shapes the input pulses is approximately 300 MHz. Assuming the effect of the sample is linear, this means we do not expect the envelope function to contain any frequency components higher than this. The

intermediate frequency ω_{IF} is chosen to be 500 MHz and the cutoff frequency of the digital filtering is $\approx$ 300 MHz. This is sufficiently above the characteristic frequency of the input pulses to be able to resolve the features of the envelope. It also removes any low-frequency noise from the input of the ADC, which has been upconverted to around 500 MHz by the digital mixing.

3.3 Experimental materials

It was established in Chapter 2 that spin waves exist in ferromagnetic materials, and a wide variety of such materials exist in nature. Which is used for experimental purposes depends on the requirements of the particular investigation to be carried out. The two most commonly used materials are permalloy, a metallic alloy commonly composed of 80% iron and 20% nickel, and yttrium-iron garnet (or YIG), an insulating crystalline material. Both are notable for their low magnetic damping at room temperature [16, 17, 78].

Permalloy has the advantage of being easily patterned using techniques such as optical lithography. It is well understood, having been in use for both experimental and commercial purposes for over one hundred years [79] and still has one of the lowest spin-wave damping parameters of metallic ferromagnets [80]. Since it is easily manipulated, it is often used for studies of small [81] or layered [82] magnetic structures. It is also a conductor, making it useful for studies of the interconversion between spin currents, where spin is transported by a flow of spin-polarised electrons, and spin waves, where the transfer of angular momentum is not accompanied by any corresponding flow of charge [83].

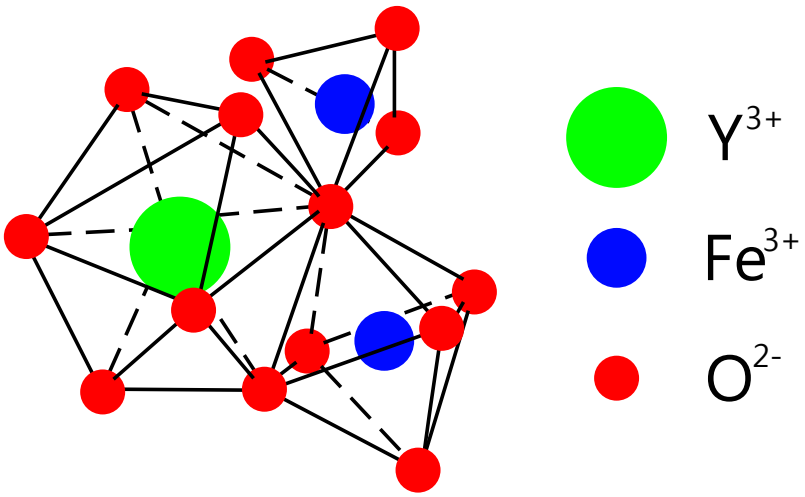

Figure 3.8: Yttrium-iron garnet has a cubic crystal structure with eighty atoms in the unit cell. The magnetic behaviour comes from the Fe^{3+} ions which are distributed between two inequivalent sites with differing oxygen coordination. Although these two sublattices are antiferromagnetically coupled, the number of iron ions on each is not equal, leading to a net magnetisation.

3.3.1 Yttrium-iron garnet

Whilst low for a metallic material, the damping of spin waves in permalloy is still high in comparison to YIG, in which spin waves can propagate over distances of millimetres or more. YIG is an artificial, electrically insulating, ferrimagnetic garnet with the formula unit $Y_3Fe_5O_{12}$. Its structure is complex, with four formula units per unit cell, making a total of eighty atoms. The magnetic behaviour comes from the Fe^{3+} ions, which are distributed on two inequivalent sites in the unit cell; eight are on octahedrally coordinated sites, and twelve on tetrahedrally coordinated sites. A schematic of the arrangement of atoms around each yttrium site is shown in Fig. 3.8. The tetrahedral and octahedral sublattices are antiferromagnetically coupled to each other, but since there are more iron ions on one sublattice than the other, the crystal still develops a net magnetic moment. The

spin of each iron ion is $\frac{5}{2}\,\mu_B$, giving a total moment per unit cell of $10\,\mu_B$. Despite this complexity on an atomic scale, YIG can be adequately approximated as a cubic ferromagnet for the purposes of studying magnetostatic spin waves, since their wavelengths are much larger than the scale of a unit cell.

High-quality YIG crystals have the narrowest ferromagnetic resonance linewidth of any known magnetic material [84]. This makes them ideal for investigating a wide range of magnetic behaviours, and makes them very attractive for use in microwave engineering. Commercial filters and oscillators based on YIG have existed for many years [85], and thin films of YIG are widely used as spin-wave waveguides both for commercial and research applications [11].

YIG films are most commonly grown by liquid-phase epitaxy onto a substrate of gadolinium-gallium garnet (GGG) with the formula unit $Gd_3Ga_5O_{12}$. GGG is chosen as the substrate because its lattice parameter (1.2383 nm) is mismatched from that of YIG (1.2376 nm) by less than 0.06%. This means that low-strain, single-crystal films can be grown with very high quality and low loss.

The saturation magnetisation of YIG has a maximum literature value at 0 K of $196\,\mathrm{kAm^{-1}}$ [86], which decreases to zero at its Curie temperature of 559 K [87]. This high Curie temperature is another attractive feature of YIG, enabling experiments to be conducted at room temperature, where its saturation magnetisation is $140\,\mathrm{kAm^{-1}}$. The crystal also exhibits cubic anisotropy which is weak at room temperature, but increases threefold at low temperatures. These properties are summarised in Table 3.1.

Table 3.1: Some properties of yttrium-iron garnet $Y_3Fe_5O_{12}$ [56]

Physical properties		
Lattice constant (room temperature)		1.2376 nm
Mass density (room remperature)		5172 kg/m^3
Magnetic properties		
Saturation magnetisation at 4 K	$M_s(4\,\mathrm{K})$	196 kAm^{-1}
Saturation magnetisation at 298 K	$M_s(298\,\mathrm{K})$	140 kAm^{-1}
Curie temperature	T_C	559 K
Anisotropy constants		
First-order cubic anisotropy at 4 K	$H_{c1}(4\,\mathrm{K})$	-20.2 kAm^{-1}
First-order cubic anisotropy at 298 K	$H_{c1}(298\,\mathrm{K})$	-6.95 kAm^{-1}

3.3.2 Superconducting thin films

Our research aims to build systems hybridising magnonic elements with devices derived from circuit quantum electrodynamics (or cQED). Structures made from superconducting thin films are very often used in cQED experiments to couple to qubits and other circuit elements [88, 89]. Using superconducting materials makes it possible to achieve very high quality factors in resonant circuit elements and makes integrating the fabrication of superconducting qubits easier. The most commonly used material for these structures in cQED is aluminium, which has a relatively low bulk critical field of 10.5 mT [90]. This would make it unsuitable for use in experiments involving ferromagnetic materials, where applied bias fields commonly exceed 100 mT.

However, it is also known that, when deposited in thin film form, all superconducting materials behave like type-2 superconductors [91]. This permits flux penetration into the material, which lowers the diamagnetic energy associated with flux expulsion. If the magnetic field is aligned parallel to the superconducting film this can enhance the critical field significantly above its bulk value by a

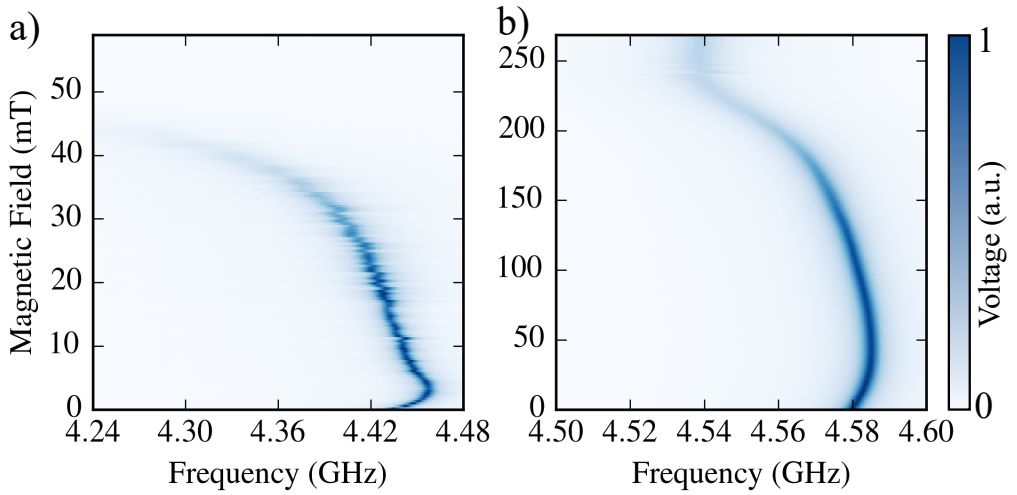

Figure 3.9: Spectroscopy measurements as a function of in-plane magnetic field of an aluminium co-planar waveguide resonator (CPWR) (panel a), and a niobium CPWR manufactured to the same design (panel b). The intensity of the plot represents the voltage measured at the resonator output for a constant power at the input. The absolute measured voltage varied between the resonators, so both plots are normalised to the their peak values to make comparison easier. The quality factor of the aluminium resonator decays rapidly as the field is increased, but the niobium resonator is robust up to much higher magnetic fields.

factor

$$\frac{B_{c||}}{B_{cb}} = \frac{2\sqrt{6}\lambda}{d} \tag{3.18}$$

where λ is the penetration depth of the material and d the thickness of the thin film. In aluminium, the penetration depth is approximately 49 nm [92], and so for thin films it would be possible in principle to realise a ten-fold enhancement of the critical field.

To test this effect, an aluminium resonator formed from a section of coplanar waveguide was fabricated. A sapphire wafer was coated with resist and exposed using optical lithography and a titanium mask. A 50 nm film of aluminium was then deposited by an electron-beam evaporator and the unwanted regions of metal

removed by lift-off of the resist.

The resonator was capacitively coupled at its two ends to an input and output line, and mounted in the sample box shown in Fig. 3.4b. Once cooled to the base temperature of the dilution refrigerator, the transmission of the resonator was measured as a function of frequency and applied in-plane magnetic field. The results of this measurement are shown in Fig. 3.9a.

The linewidth of the resonator remains small well above the bulk critical field of aluminium, indicating that the film is still superconducting. The film has a thickness of 50 nm and the critical field is enhanced to approximately 50 mT, which is consistent with Eq. 3.18. However, even the enhanced critical field was lower than required for magnon experiments, and so an alternative material had to be selected.

An resonator identical to the aluminium prototype was fabricated from a 100 nm film of niobium. As a type-II superconductor, niobium has a much higher critical field than aluminium – around 200 mT [93]. Spectroscopy as a function of in-plane magnetic field for the niobium resonator is shown in Fig. 3.9b. The niobium resonator remains superconducting up to around 200 mT – several times higher than the aluminium resonator – making it much more suitable for interfacing with magnonic structures.

Chapter 4

Strong coupling of resonant magnons and photons

Having described the theory underlying spin waves, and the experimental techniques we use to investigate them, we now present measurements of a resonant magnonic system. Dilute ensembles of electron spins, such as nitrogen vacancy centres, have been proposed as a form of memory for a quantum computer [25, 26]. The weak interactions between the spins improve their coherence times, potentially allowing storage of quantum information over useful time scales. However, this comes at the cost of weak coupling to external fields due to the low spin density of such systems. Ferromagnetic spheres represent the opposite limit; they have very high spin densities and strong spin-spin interactions. This potentially allows information to be encoded in collective excitations of the spins, whilst providing high coupling strengths due to the high spin density. Ferromagnetic spin ensembles also have potential in microwave to optical conversion [51], which could form the basis of various kinds of quantum networks [46]. These reasons illustrate

the potential utility of magnonic systems in a quantum information context.

Resonant structures are widely used in superconducting quantum circuits to read out qubit states and to improve coherence times by reducing the available final states for decay [94]. They represent the natural choice to couple microwave-frequency excitations between electron spin ensembles and quantum information devices. If we want to build quantum magnonic devices it is therefore important to verify that these structures are compatible with magnonic components at the low temperatures and input powers that are necessary for the operation of superconducting qubits. Previous investigations have established the feasibility of coupling magnonic components with three-dimensional cavities of various types [28, 29, 31–33, 35, 36, 38]. However, there has been only one study investigating the coupling to planar structures such as a coplanar waveguide resonator (CPWR) [27], and this was performed at a relatively high electrical temperature with high input powers incompatible with the operation of a superconducting qubit. These planar structures have many desirable qualities, such as the ability to include qubit-specific control lines [95] and fabrication techniques which simplify integration with other components.

In this chapter we demonstrate coupling of the fundamental photon mode of a superconducting CPWR to resonant magnon modes in a sphere of YIG. We also discuss the theoretical approach used to identify which magnon modes couple to the CPWR, and what behaviour we expect to see. The findings in this chapter were published in [96].

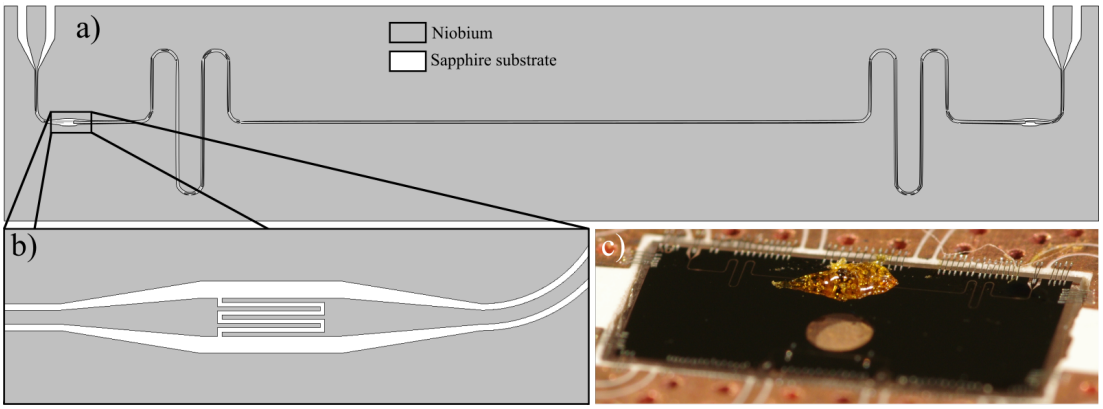

Figure 4.1: a) Schematic of the CPWR. The meander structures at either end increase the length of the CPWR within the allowable space, such that it has the desired resonance frequency. The centre section is straight to provide a region of uniform electromagnetic field. b) Close up of the coupling structures at either end of the resonator. The width of the centre conductor is 10 μm and the gap between centre conductor and ground plane is 6 μm. These structures connect the input and output lines to the resonator capacitively. The fingers increase the effective area of the capacitors, thus increasing their capacitance and reducing their impedance. This provides stronger coupling to the input and output lines. c) A picture of the resonator with a YIG sphere placed on the centre in preparation for the coupling experiment. The YIG sphere is inside the GE varnish used to fix it in place.

4.1 Sample design

Our experiment consists of two components: a microwave-frequency photonic resonator and a magnonic resonator. These are coupled together through their magnetic fields.

The magnonic resonator is formed by a 250 μm-diameter YIG sphere, obtained from Ferrisphere Inc. [97]. As described in Section 2.6.1, these spheres can support a wide variety of resonant modes corresponding to different distributions of the magnetisation through the volume of the sphere. As seen in Fig. 2.7, the frequencies of these modes increase as the strength of the applied magnetic field is increased. The photonic resonator is a coplanar waveguide resonator (CPWR),

which was fabricated from a 100 nm film of niobium. Such resonators are very simple to analyse, since their resonance frequency depends only on their length and an effective permittivity parameter ϵ_{eff}.

The design of the CPWR is shown in Fig. 4.1a. Since the structure is symmetric, it forms a $\frac{1}{2}\lambda$ resonator. In order to perform spectroscopy on the CPWR, it is capacitively coupled at each end to coplanar waveguide transmission lines, as shown in Fig. 4.1b. Assuming that the input and output coupling is weak enough not to shift the resonance, its frequency in zero magnetic field is given by:

$$f_0 = \frac{c}{2l\sqrt{\epsilon_{\text{eff}}}} \tag{4.1}$$

where l is the length of the CPWR, c the speed of light in vacuum, and ϵ_{eff} an effective permittivity which accounts for both the permittivity of the substrate and the geometry of the coplanar waveguide [89].

Initially, the CPWR was measured in the absence of any magnonic component to determine how it behaved in the presence of an applied magnetic field. The transmission of the CPWR was measured by applying a signal to one end and measuring the voltage produced at the other. The CPWR acts as a bandpass filter and will only transmit signals at frequencies close to its resonance frequency. The results of this measurement are shown in Fig. 3.9. We can investigate how the properties of the CPWR change as a function of field by fitting a Lorentzian lineshape to its transmission at each value of applied magnetic field. From this, the centre frequency and linewidth can be extracted. In Fig. 4.2 we plot the change in the centre frequency and quality factor of the CPWR as a function of field. The quality factor, Q, is a measure of the lossiness of the structure, with a higher Q reflecting lower loss. It is calculated from $Q = \frac{f_0}{\Delta f}$, where f_0 is the centre

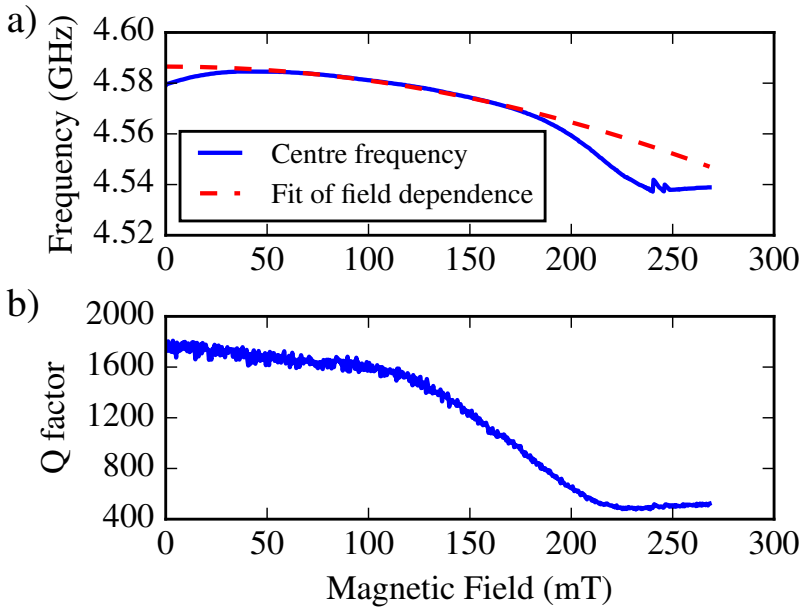

Figure 4.2: a) The centre frequency of the niobium CPWR as a function of the strength of a magnetic field applied in the plane of the resonator. At zero field it has a centre frequency of 4.579 GHz. A fit of the expected behaviour of the centre frequency is also plotted. There is a non-zero shift of the centre frequency at zero field due to trapped flux in the film and magnet. The fit breaks down at fields approaching the critical field. b) The quality factor of the resonator as a function of magnetic field. The quality factor decreases as field increases, but only begins decreasing sharply above about 120 mT.

frequency of the resonator and Δf its full-width at half-maximum. For $B \ll B_c$ (where B_c is the critical field of the superconductor) the centre frequency of the CPWR is expected to change quadratically with magnetic field [98] according to

$$f(B) = f_0\left(1 - \frac{B^2}{B_s}\right), \tag{4.2}$$

where B_s is a constant determined from a combination of the critical field of the superconductor and its inductance. In Fig. 4.2a the centre frequency is fitted using Eq. 4.2. The fit is good over a large part of the measured magnetic field range. At low fields, flux can become trapped in both the superconducting film of the CPWR ground plane and the coils of the superconducting magnet providing

the applied field. This means that even when there is no current in the magnet, the film still experiences a non-zero effective field [98]. Consequently, there is a non-zero deviation of the centre frequency at low fields. At high fields we approach the critical field of niobium of 200 mT [93] and the fit becomes invalid.

The quality factor of the CPWR reduces as applied field increases, due to the increased level of flux penetration into the film. The motion of flux vortices causes dissipation in the film and as the number of vortices increases, so too does the loss [99]. However, the CPWR is relatively unaffected by fields up to approximately 120 mT, a range large enough to allow it to interact with magnons in the GHz frequency regime. We also note that the magnetic field dependence of the centre frequency is much weaker than that of the magnonic modes in the YIG sphere, as can be seen from the comparison with Fig. 2.7.

Another important characteristic of the CPWR for understanding how it will interact with magnons in a ferromagnetic material is the shape of its associated magnetic field. To determine this, we use the high-frequency simulation software HFSS [100], a finite-element solver from ANSYS Inc. [101]. An example of a calculation performed in HFSS is shown in Fig. 4.3a. As seen in Fig. 4.3b, the field shape near the centre of the CPWR is very non-uniform over a scale of tens of microns.

The YIG sphere is glued using GE varnish, a thermally conductive material with good low-temperature microwave properties, onto the centre of the CPWR where its magnetic field is maximal, as shown in Fig. 4.1c. The chip carrying the CPWR is placed into an oxygen-free copper sample box, and connected to the room temperature electronics described in Chapter 3.

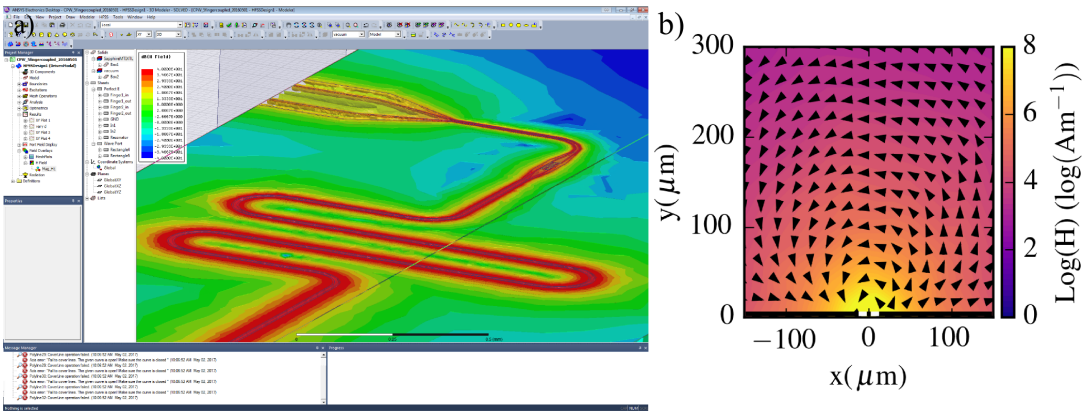

Figure 4.3: a) A screenshot of the HFSS software, displaying the magnitude of the magnetic field around the meander structure at one end of the resonator. b) A plot of the magnetic field around the centre of the resonator, calculated using data from the HFSS simulation shown in a.

4.2 Coupling of magnonic and photonic modes

4.2.1 Coupled simple harmonic oscillators

Both the magnon modes in the YIG sphere and the CPWR are resonators that can be modelled as simple harmonic oscillators. Their individual Hamiltonians are given by

$$\hat{\mathcal{H}}_{\mathrm{p}} = \hbar\omega_{\mathrm{p}}\left(\hat{a}_{\mathrm{p}}^{\dagger}\hat{a}_{\mathrm{p}} + \frac{1}{2}\right), \tag{4.3a}$$

$$\hat{\mathcal{H}}_{\mathrm{m}} = \hbar\omega_{\mathrm{m}}\left(\hat{a}_{\mathrm{m}}^{\dagger}\hat{a}_{\mathrm{m}} + \frac{1}{2}\right). \tag{4.3b}$$

In Eq. 4.3a, the creation and annihilation operators $\hat{a}_{\mathrm{p}}^{\dagger}$ and $\hat{a}_{\mathrm{p}}$ are formed analogously to the normal quantum harmonic oscillator, with the operators for the electric and magnetic field for each resonant mode taking the place of the position and momentum of the oscillator [102]. In the case of the magnon mode, the creation and annihilation operators in Eq. 4.3b are formed by a Holstein-Primakoff transformation of the spin raising and lowering operators [103]. Our experimental

system consists of these two oscillators coupled together via their magnetic fields. To explore its expected behaviour, we consider transitions between the ground state, where neither oscillator is excited, and the one-excitation manifold, where one excitation is shared between the two oscillators.

When a single excitation is present in the system composed of the two independent oscillators, two states are possible: Ψ_{10} where the photon mode is excited, and Ψ_{01} where the magnon mode is excited. Using these two states as the basis of the Hamiltonian, and recognising that $\hat{a}^\dagger\hat{a}$ is equal to the number operator $\hat{n}$, the total Hamiltonian for the system of independent oscillators is given by

$$\hat{\mathcal{H}}_\mathrm{T} = \hat{\mathcal{H}}_\mathrm{p} + \hat{\mathcal{H}}_\mathrm{m} = \begin{pmatrix} \hbar\omega_\mathrm{p} & 0 \\ 0 & \hbar\omega_\mathrm{m} \end{pmatrix}, \tag{4.4}$$

where the zero-point energy has been neglected for simplicity. This means that the excited-state eigenvalues of the system are just the frequencies $\hbar\omega_\mathrm{p}$ and $\hbar\omega_\mathrm{m}$ of the two individual oscillators. If we introduce a single photon to the system initially in the ground state Ψ_{00}, two transitions are possible corresponding to $\Psi_{00} \rightarrow \Psi_{10}$ or $\Psi_{00} \rightarrow \Psi_{01}$. Conducting a spectroscopy measurement will reveal two peaks: one at ω_m and another at ω_p.

We now introduce an interaction term to the Hamiltonian of the one-excitation manifold. The interaction is mediated through the magnetic fields of the two oscillators and its strength is characterised by the parameter g. Under the rotating wave approximation, the interaction term is given by

$$\hat{\mathcal{H}}_\mathrm{int} = \hbar g\big(\hat{a}_\mathrm{p}^\dagger\hat{a}_\mathrm{m} + \hat{a}_\mathrm{p}\hat{a}_\mathrm{m}^\dagger\big). \tag{4.5}$$

The interaction either transfers an excitation from the photon mode to the

magnon mode or vice versa. Its action is described by

$$\hat{\mathcal{H}}_{\text{int}}\Psi_{10} = \hbar g\Psi_{01}, \tag{4.6a}$$

$$\hat{\mathcal{H}}_{\text{int}}\Psi_{01} = \hbar g\Psi_{10}. \tag{4.6b}$$

The total Hamiltonian now becomes

$$\hat{\mathcal{H}}_{\text{T}} = \hat{\mathcal{H}}_{\text{p}} + \hat{\mathcal{H}}_{\text{m}} + \hat{\mathcal{H}}_{\text{int}} = \begin{pmatrix} \hbar\omega_{\text{p}} & \hbar g \\ \hbar g & \hbar\omega_{\text{m}} \end{pmatrix}. \tag{4.7}$$

The states of the system will be modified by the interaction, and to find the energies of the excited levels of the coupled system, we diagonalise the Hamiltonian, yielding new eigenvalues

$$\omega_{\pm} = \frac{\omega_{\text{p}} + \omega_{\text{m}}}{2} \pm \frac{1}{2}\sqrt{4g^2 + (\omega_{\text{p}} - \omega_{\text{m}})^2}. \tag{4.8}$$

When the frequencies of the two oscillators are very different ($\omega_{\text{p}} - \omega_{\text{m}} \gg g$), this reduces to the original frequencies of the two separate oscillators. However, when the frequencies of the separate oscillators are close together ($\omega_{\text{p}} - \omega_{\text{m}} \ll g$), the levels of the coupled system display an avoided crossing. The minimum gap associated with the crossing is $2g$. In Fig. 4.4 this is demonstrated theoretically. The energy levels are plotted in the case that the frequency of one oscillator is constant, while the frequency of the other is a function of applied magnetic field.

This result can be generalised to the case where an arbitrary number of excitations are present in the system. For n excitations, we have $n + 1$ possible states of the combined two-oscillator system: Ψ_{n0}, $\Psi_{(n-1)1}$, ..., $\Psi_{1(n-1)}$, Ψ_{0n}. When the interaction term is introduced to the Hamiltonian, the new eigenstates are a ladder of states, each separated from its neighbours by an energy gap

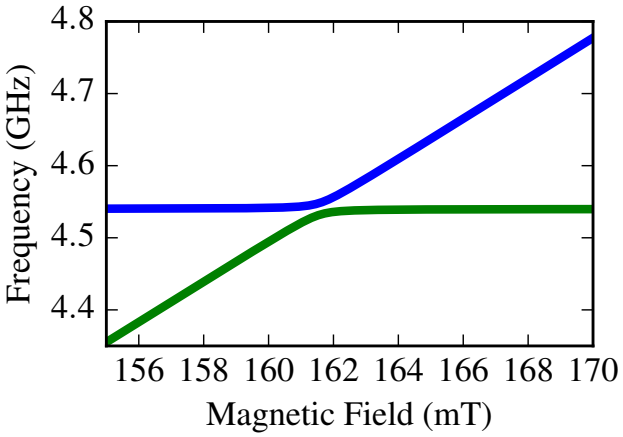

Figure 4.4: The energy levels of a system of two coupled simple harmonic oscillators, where the frequency of one of the oscillators changes as a function of magnetic field. The avoided crossing is visible in the centre where the two uncoupled levels would cross. The gap between the levels at this point has a minimum value of $2g$ determined by the coupling strength.

$\Delta = \sqrt{4g^2 + (\omega_\mathrm{p} - \omega_\mathrm{m})^2}$. This is illustrated for the possible levels up to $n = 3$ in Fig. 4.5.

Introducing an additional excitation induces transitions between the set of states for $n = m$ and the set of states for $n = m + 1$. Going from the ground state $n = 0$ to one of the first excited states for $n = 1$ there are two possible transitions with the frequencies given by the eigenvalues from Eq. 4.8. However, for higher transitions there are potentially many more possible transition frequencies between the sets of levels. For instance, in Fig. 4.5 we can see that going from $n = 1$ to $n = 2$ excitations in the coupled system there are six separate transitions possible. Some of these will have the same frequencies as the transitions from ground to first excited states, but others will not. However, in experiment, we see only the two frequencies corresponding to the two eigenvalues given above. This is due to the form of the eigenfunctions for the hybridised states; these ensure

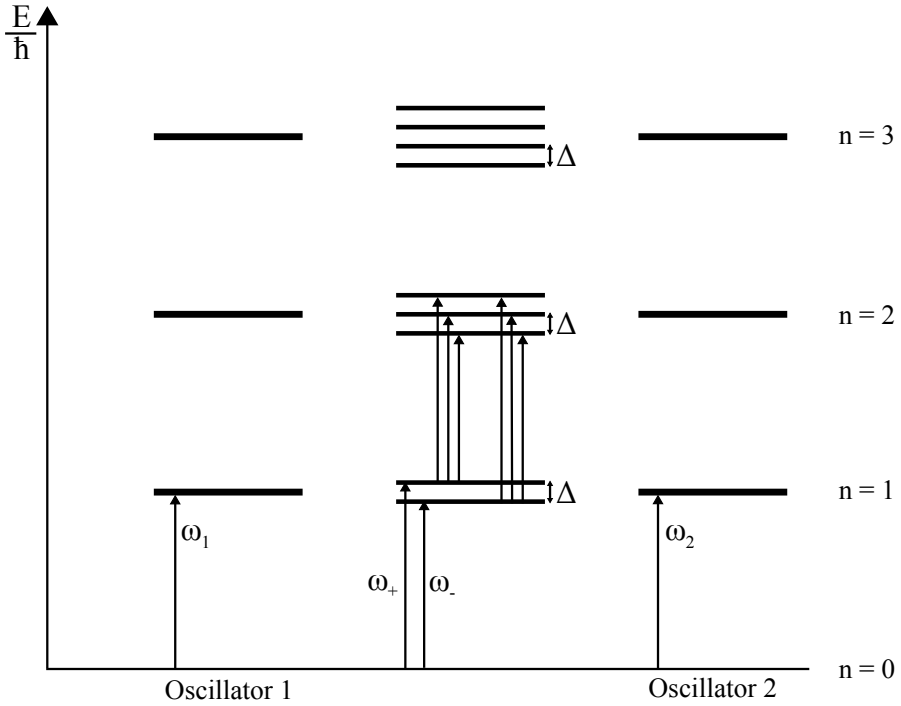

Figure 4.5: The ladder of states for a hybridised two-oscillator system, where both oscillators have the same resonance frequency, so $\omega_1 = \omega_2$ and $\Delta = 2g$. The levels of the two individual oscillators are shown on the left and right, and the levels of the hybridised system are shown in the centre. The two transitions from $n = 0$ to $n = 1$ correspond to the observed frequencies plotted in Fig. 4.4. The six possible transitions from $n = 1$ to $n = 2$ are also marked.

that the matrix element for transitions is only non-zero when the energy gap is equal to $\omega_\pm$. This means that we still only see two peaks in the spectrum even when an arbitrary number of excitations is present.

In our experimental system, the magnetic field dependence of the magnon mode frequencies is much stronger than that of the CPWR centre frequency (compare Fig. 2.7 to Fig. 4.2). This means that by varying the magnetic field applied to the sample, we can control the difference in frequency between the CPWR resonance and the magnon resonances. From this analysis we know that, if the magnon and photon modes couple sufficiently strongly, an avoided crossing

will be visible as this difference is reduced. If the coupling is strong enough, then around the region of the crossing, we should be able to measure the two modes and their splitting. Even though only the CPWR is experimentally accessible, both resonant lines will be visible around the crossing point, since the hybridisation of the modes means both are partially magnon-like and partially photon-like.

On account of their non-zero losses, both the magnon mode of the YIG sphere and the photon mode of the CPWR have a non-zero linewidth. In order for the states of the two resonators to be properly hybridised they must be able to exchange excitations. However, if the rate of energy loss from the system is greater than the rate of energy exchange between the resonators, then a single excitation will, on average, be lost before it can be passed back and forth. This means that instead of an avoided crossing, we see only enhanced damping, as the coupling provides an additional means for the oscillator to lose energy. This enhanced damping is known as the Purcell effect [104]. In practice, the regime where the avoided crossing will be visible (the so-called 'strong coupling' regime) is achieved if the linewidths of each subsystem are smaller than the coupling strength between them.

4.2.2 Mode matching

As was demonstrated in Section 2.6.1, there are a wide variety of possible resonant modes in a ferromagnetic sphere, and it would be useful to be able to predict which will couple most strongly to the CPWR. We can estimate the relative strength of the coupling to different modes through the interaction energy of the magnetic field of the CPWR with the magnetisation of the sphere.

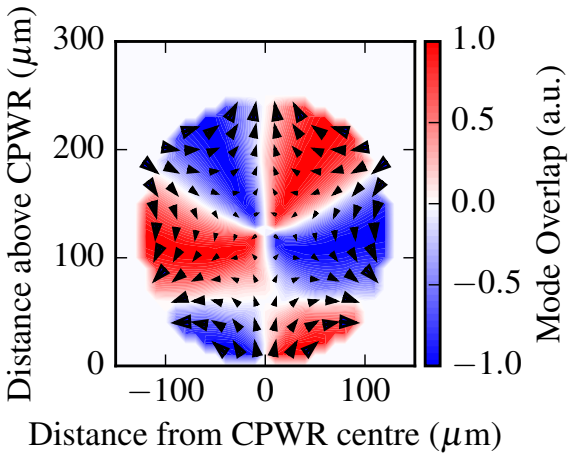

Figure 4.6: Overlap of the $(4, 3)$ resonant mode in the sphere with the magnetic field of the CPWR. The colour map is calculated from the dot product of the CPWR magnetic field direction and sphere magnetisation at each point. Note that, because only the CPWR field direction is used, this plot does not account for the variation in the CPWR field intensity, as plotted in Fig. 4.3b. This is done to improve visibility of the overlap values. In the experiment, the region of the sphere closest to the CPWR (i.e. at the bottom of the plot) will contribute most of the coupling between the sphere and the CPWR.

The energy associated with a magnetic moment $\boldsymbol{\mu}$ placed in a magnetic field with flux density $\mathbf{B}$ is given by

$$E = -\boldsymbol{\mu} \cdot \mathbf{B}. \tag{4.9}$$

As described in Section 2.4.2 we assume that the magnetisation in the sphere is given by a large constant component aligned with the magnetic bias field and a small time-dependent component in the plane perpendicular to the bias field. The CPWR is aligned such that the centre conductor is parallel to the bias field. This means that the magnetic field of the CPWR shown in Fig. 4.3b is in the plane perpendicular to the bias field and therefore in the same plane as the small-signal magnetisation associated with the resonant sphere modes.

By comparing the magnetisation profiles in the sphere (calculated in the man-

ner described in Section 2.6.1) with the CPWR magnetic field we can estimate the coupling strength. An example of such a calculation is plotted in Fig. 4.6 for the $(4, 3)$ mode.

The total energy of interaction of a particular resonant mode of the sphere with the magnetic field it experiences can be calculated by integrating all of these individual points over the volume of the sphere:

$$E = \int_V \mathbf{M} \cdot \mathbf{B} \, \mathrm{d}V = \int_V (M_0 \hat{\mathbf{z}} + \mathbf{m}_{nm}) \cdot (B_0 \hat{\mathbf{z}} + \mathbf{b}_\mathrm{r}) \, \mathrm{d}V, \tag{4.10}$$

where $\mathbf{m}_{nm}$ and $\mathbf{b}_\mathrm{r}$ are the small-signal components from the sphere and CPWR respectively, which are in the xy-plane perpendicular to the bias field. Since the large dc components will be approximately the same regardless of which sphere mode is being considered, we can neglect their contribution to the energy as a constant offset, leaving

$$E = \int_V \mathbf{m}_{nm} \cdot \mathbf{b}_\mathrm{r} \, \mathrm{d}V. \tag{4.11}$$

Calculating this integral for a variety of different modes in the sphere will give an indication of how closely the magnon mode shape matches the field distribution of the CPWR, and therefore how strongly it couples to the resonator. It is for this reason that the inhomogeneity of the CPWR field on the scale of tens of microns is significant: if the field were highly uniform it would couple efficiently only to the ferromagnetic resonance mode where the magnetisation of the sphere is uniform (c.f. Fig. 2.8a).

The magnetisation distributions for a wide range of modes have been tabulated by Fletcher and Bell [62]. Using these, we can calculate the overlap for different mode families as a function of n. The results of these calculations for

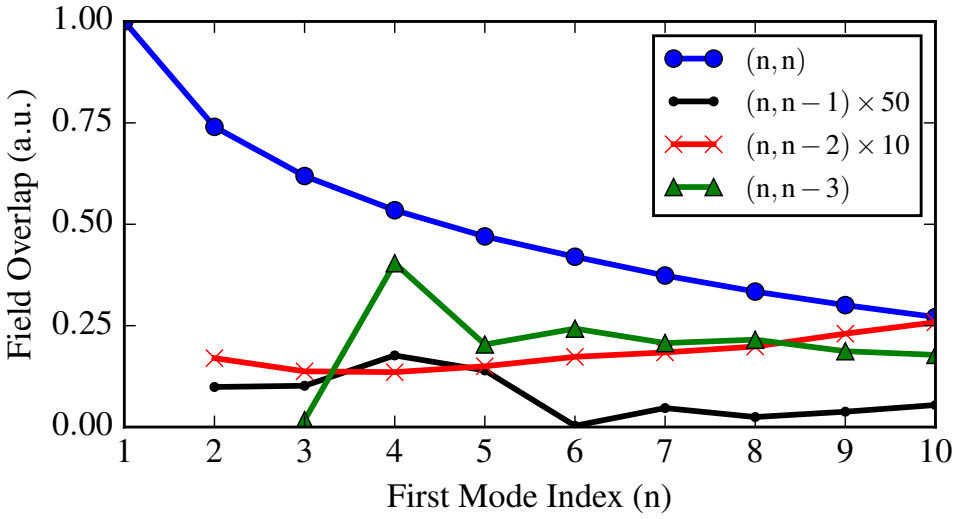

Figure 4.7: Overlap of different resonant modes in the YIG sphere with the resonator field for a range of values of (n, m). The overlap is normalised such that the $(1, 1)$ (ferromagnetic resonance) mode has a value of 1, since it has the largest coupling strength. The values for the $(n, n-1)$ and $(n, n-2)$ families are multiplied by the numbers shown in the legend to increase their visibility in the plot.

four different sets of modes are plotted in Fig. 4.7. The (n, n) modes consistently couple most strongly up to a large value of n.

4.3 Experimental results

The results of transmission spectroscopy of the CPWR as a function of magnetic field are shown in Fig. 4.8. Above about 120 mT a series of avoided crossings are visible as modes in the sphere are brought into resonance with the CPWR. A large number of weaker lines are also visible.

On top of the spectroscopy measurement are plotted the frequencies of the first six (n, n) modes of the YIG sphere. They are calculated for a saturation magnetisation of 169 $\mathrm{kAm^{-1}}$ and a constant offset has been added to the frequen-

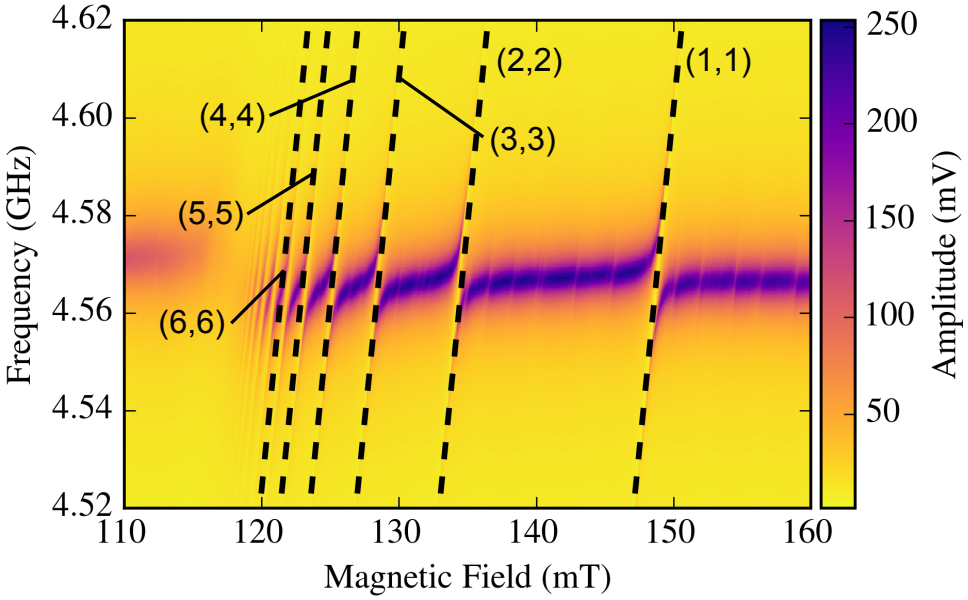

Figure 4.8: Voltage measured at the output of the CPWR as a function of frequency and applied magnetic field for a constant input power to the resonator of -80 dBm. Each point in the measurement is averaged 2000 times. Approximately ten avoided crossings are visible between the CPWR and modes in the YIG sphere, along with weaker features due to less strongly coupled sphere modes. Six avoided crossings are fitted with the frequencies of the first six (n, n) modes for a saturation magnetisation of $169\,\mathrm{kAm^{-1}}$.

cies of all six modes. The lines calculated in this way correspond well to the largest avoided crossings seen in the measurement. The frequencies of the (n, n) modes converge as the magnetic field is reduced, giving rise to the broad dip in CPWR transmission around 120 mT as the modes become almost continuous. The weaker lines are likely due to other families of modes in the sphere with $n \neq m$, such as those plotted in Fig. 4.7.

The constant offset was added to the mode frequencies in Fig. 4.8 to account for the effect of magnetocrystalline anisotropy. Yttrium-iron garnet exhibits cubic magnetocrystalline anisotropy [105], and the strength of this anisotropy increases

three-fold between room temperature and 4 K (see Table 3.1). As mentioned in Section 2.6.1, if the biasing magnetic field is along one of the hard or easy axes of the crystal, the effect of the anisotropy is to shift the entire spectrum of modes by a constant factor [106]. However, in the case of an arbitrarily-directed bias field, the problem immediately becomes highly complex, and exact solutions for the mode frequencies exist for only the first few lowest-order magnetostatic modes [66].

In our experiment it was not possible to align the crystal axes of the sphere with the bias field. However, adding a constant offset to all the modes did allow us to get good agreement between the observed and predicted modes. The offset was calculated from the difference between the theoretical and observed FMR frequency. For modes above $n = 6$, the simple offset is no longer sufficient and the observed frequencies begin to deviate from the predicted frequencies. This is likely due to the more complex behaviour of the mode frequencies for a general alignment of the magnetic field relative to the crystal axes.

The value of $169\,\mathrm{kAm^{-1}}$ for the saturation magnetisation which was used to fit the mode spacings is significantly different from the literature value of $196\,\mathrm{kAm^{-1}}$ given in Chapter 3 which would normally be expected for yttrium-iron garnet at temperatures below 4 K. This value was obtained from the observed spacing of the $(1, 1)$ and $(2, 2)$ modes; a measure which has been used previously to determine the saturation magnetisation of YIG [107]. Magnetocrystalline anisotropy can affect the spacing of sphere modes as well as their absolute frequencies, and for these two low-order modes analytic expressions exist for their frequencies in the

presence of anisotropy [66]

$$\omega_{11} = \gamma\sqrt{(H_0 + T_x)(H_0 - T_y)}, \tag{4.12a}$$

$$\omega_{22} = \gamma\sqrt{\left(H_0 + \frac{M_\mathrm{s}}{15} + T_x\right)\left(H_0 + \frac{M_\mathrm{s}}{15} - T_y\right) - \frac{2M_\mathrm{s}^2}{25}}, \tag{4.12b}$$

where T_x and T_y are the torque constants for a cubic crystal with H_0 in the [110] plane of the crystal. If we neglect the second-order cubic anisotropy, which is approximately 20 times smaller than the first-order anisotropy in YIG, the equations become

$$T_x = \frac{K_1}{M_\mathrm{s}}\big(2 - \sin^2\theta - 3\sin^2 2\theta\big), \tag{4.13a}$$

$$T_y = -\frac{2K_1}{M_\mathrm{s}}\left(1 - 2\sin^2\theta - \frac{3}{8}\sin^2 2\theta\right). \tag{4.13b}$$

Here $\frac{K_1}{M_\mathrm{s}}$ is the first-order anisotropy field for YIG at temperatures below 4 K as given in Table 3.1. The angle θ is that between the bias field H_0 and the [100] axis in the [110] plane. Since the [110] plane contains both easy and hard axes of the crystal, we should see the full range of possible values by rotating the crystal in this plane.

Using these expressions, we can examine to what extent the magnetocrystalline anisotropy is able to explain the unexpected spacing of the first two (n, n) modes. The theoretical frequency of the $(1, 1)$ mode, given by Eq. 4.12a was calculated as a function of the angle θ between the hard crystal axis and the applied field for a field strength of 148.8 mT. Figure 4.9a is a plot of the difference between this frequency and the experimentally observed $(1, 1)$ frequency at the same applied field (taken from Fig. 4.8). The angles where the line crosses zero indicate alignments where the theoretical frequency is consistent with the observed frequency.

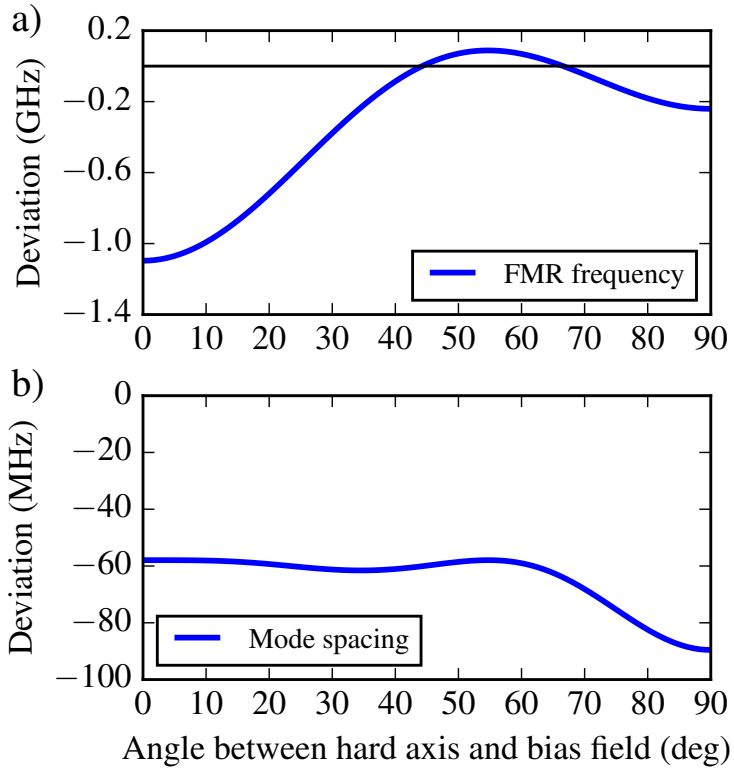

Figure 4.9: a) A plot of the difference between the experimentally observed and theoretically calculated FMR frequencies for a YIG sphere with cubic anisotropy using the literature value of saturation magnetisation of $196\,\mathrm{kAm^{-1}}$ in a magnetic field of 148.8 mT. The difference is plotted as a function of the angle between the applied field and the hard axis of the crystal. When the line crosses zero, the theoretical FMR frequency matches the observed frequency. This occurs at two angles around 44° and 66°. b) A plot of the difference between the observed spacing of the $(1, 1)$ and $(2, 2)$ modes, and the theoretical spacing for the same sphere as in a. The difference is again plotted as a function of angle between the applied field and the crystalline hard axis. There are no angles where the line crosses zero, indicating that for this value of saturation magnetisation there are no alignments of the sphere where anisotropy can account for the observed mode spacing.

The theoretical spacing between the $(1,1)$ and $(2,2)$ modes was also calculated as a function of θ by taking the difference between Eq. 4.12a and Eq. 4.12b. In Fig. 4.9b, the difference between the calculated and observed spacing is plotted as a function of angle. As in Fig. 4.9a, any point where the line crosses zero indicates an angle where the theoretically calculated spacing is consistent with the experimentally observed spacing. Both of these panels are calculated using the literature value of saturation magnetisation of $196\,\mathrm{kAm^{-1}}$. There are two possible angles where the observed FMR frequency matches the theoretical frequency. However, there are no angles where the spacing of the first two modes matches the theoretical spacing.

It is for this reason that a value of $169\,\mathrm{kAm^{-1}}$ was used for fitting the lines in Fig. 4.8; this is the largest value of saturation magnetisation for which it is possible to find an angle where the theoretical spacing of the $(1,1)$ and $(2,2)$ modes matches the observed spacing.

Non-sphericity of the YIG sample may also contribute to the smaller spacing of the modes, since it has been shown that the spacing of modes in an ellipsoidal sample is smaller than that in a perfectly spherical sample [59]. However, the calculations in the literature are performed for an ellipsoid of revolution with the bias field aligned along its semi-major axis. Since it was neither possible to align the sphere, nor guarantee that any deviations from perfect sphericity are of this kind, we cannot make any quantitative statements about the effect of this imperfection, other than to say that it may contribute to the smaller than expected mode spacings observed in the experiment.

Extracting the coupling parameters

In Fig. 4.10 the amplitude and phase data for a high-resolution measurement of the $(1,1)$ crossing are fitted using an expression for the expected transmission derived from input-output theory [28]

$$S_{21}(\omega) = \frac{\kappa_{\text{ext}}}{i(\omega - \omega_{\text{r}}) - \kappa_{\text{ext}} - \frac{\kappa_{\text{int}}}{2} + \frac{g_{\text{m}}^2}{i(\omega - \omega_{\text{m}}) - \frac{\gamma_{\text{m}}}{2}}}, \tag{4.14}$$

where ω_{r} and ω_{m} are the frequencies of the resonator and magnon mode respectively; κ_{int} and κ_{ext} are the internal and external linewidths of the resonator, corresponding to its intrinsic losses and the additional losses introduced by coupling to the input and output lines; γ_{m} is the linewidth of the FMR mode; and g_{m} is the coupling strength between the sphere and resonator. The frequency of the resonator is $\omega_{\text{r}} = 4.5679\,\text{GHz}$, as determined from its value away from the crossing, and the FMR frequency was calculated using Eq. 4.12a. From the fit we can extract the coupling parameters, which are given in Table 4.1.

The fit in Fig. 4.10 shows that the coupling strength between the two resonators is larger than the linewidths of either of the individual resonators. This is important because it means that energy can be exchanged between the systems faster than it decays away, demonstrating that coherent exchange of excitations is possible. For potential quantum applications, this establishes that information

Table 4.1: The extracted properties of the coupled sphere-CPWR system.

Parameter	**Value** (MHz)
κ_{ext}	1.34
κ_{int}	1.39
γ_{m}	2.97
g_{m}	8.17

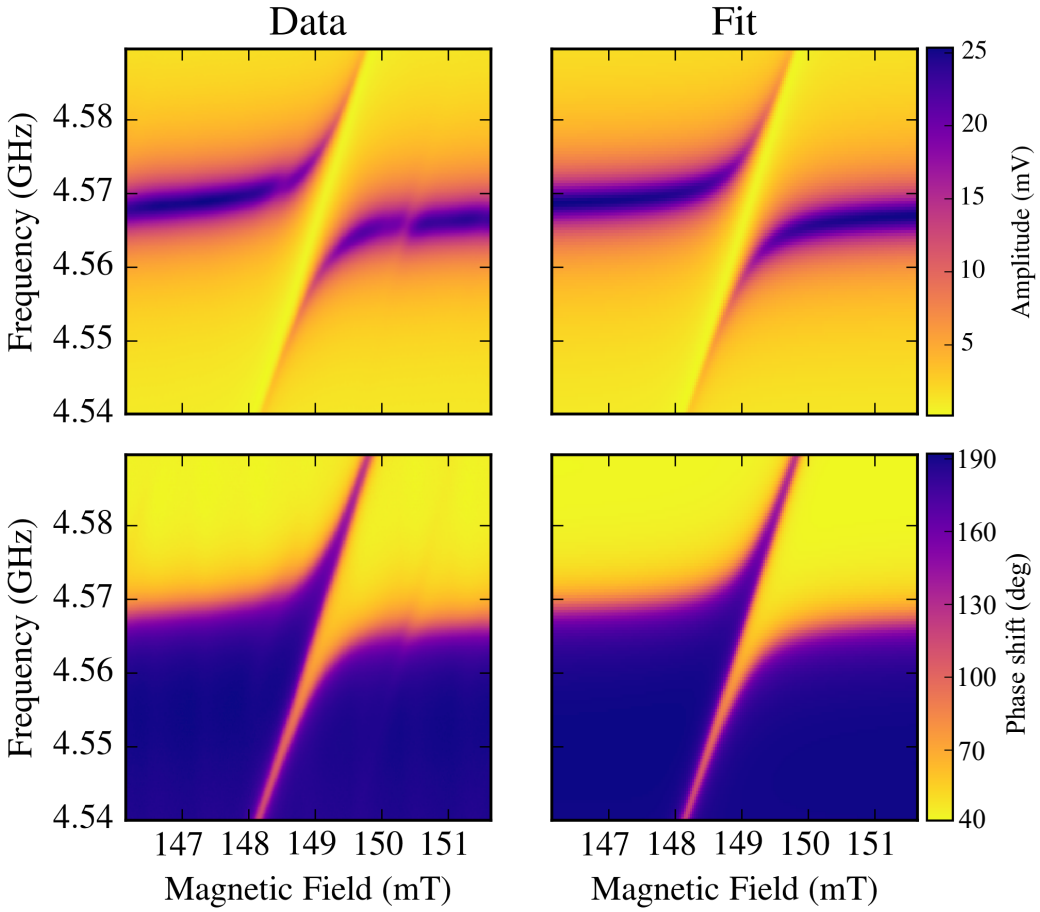

Figure 4.10: High-resolution measurement of the $(1, 1)$ crossing at -100 dBm input power and with 4000 averages. Measured data for amplitude (top) and phase (bottom) of the complex resonator transmission are shown on the left. A fit to the data using Eq. 4.14 is shown on the right. Note the two weak features in the data which appear at higher and lower field than the main resonance. These are due to other modes in the sphere which were not accounted for in the fit. The parameters extracted from the fit are given in Table 4.1.

can potentially be transferred coherently.

Having extracted the parameters of the system, we went on to characterise it in a regime with only a very small number of excitations present in the system. This is a relevant consideration because for future experimental devices incorporating superconducting qubits, the system should be able to interact and provide readout when there are only single photons present. To explore this, we make measurements with the input power to the resonator is reduced to -140 dBm. Under these conditions, since the lifetime of a single excitation in the resonator is approximately $\frac{1}{\kappa_{\mathrm{int}}+\kappa_{\mathrm{ext}}}$, we would expect there to be on average only one photon present in the resonator at any given time. Also, the negligible thermal background at the experimental temperature of 10 mK means that magnons are only created via the CPWR, so there should only be a single excitation in the system. A measurement made under these conditions is shown in Fig. 4.11.

To achieve the result in Fig. 4.11 required a high level of averaging; each point in the field-frequency range shown here was averaged 7.5×10^5 times. The avoided crossing is plainly visible, and the data displays the same characteristics as that collected at higher input powers.

4.4 Conclusions

The measurements in this chapter have established that it is possible to couple a magnonic resonator system to a planar superconducting structure akin to those commonly used in circuit QED. An avoided crossing is observed, which demonstrates that strong coupling has been achieved – an important requirement for building hybrid quantum systems using these structures. In addition, due to

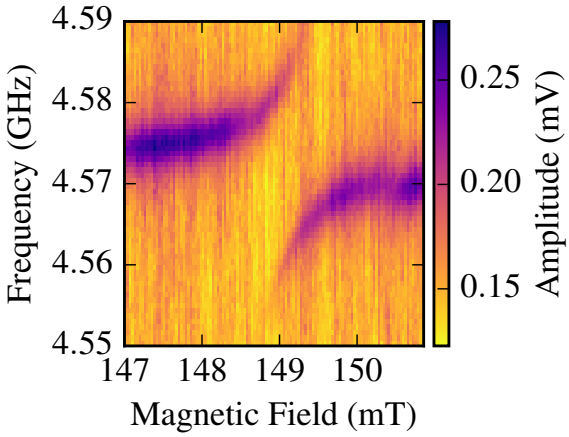

Figure 4.11: Measurement of the $(1, 1)$ crossing when the input power to the resonator is reduced such that, on average, only a single excitation is present in the resonator.

the non-uniform microwave magnetic field of the excitation structure used in the experiment, a wide variety of different modes can be observed coupling to the resonator. We have also shown that our measurement setup is capable of performing measurements even down to the quantum limit, where only single excitations are present in the system.

A theoretical analysis allowed us to identify the strongly coupled modes, although the value of saturation magnetisation determined from the mode spacings is significantly different to the literature value. Further calculations have established that this deviation cannot be explained by the presence of magnetocrystalline anisotropy in the sphere. These observations are an important step towards more complex experiments building hybrid quantum systems incorporating magnonic elements.

Chapter 5

Propagating magnons in thin films

In Chapter 4 we examined the behaviour of resonant magnon modes in a spherical geometry at low temperatures. Resonant modes appear when the size of the sample is small compared to the wavelengths and decay lengths of the excitations being produced. When the sample is large compared to these length scales we can produce propagating excitations such as those examined in Section 2.6.2.

There has been a recent growth in the development of technology based on propagating spin waves for applications in information technology [78]. Using magnons as information carriers avoids the problem of Joule heating in computing devices, and presents the opportunity to build new wave-based logic devices [22]. Interfacing these emerging technologies with existing quantum information systems would potentially create new functionality. Such hybrid systems would also provide an opportunity to study the single-excitation dynamics of magnon systems. Towards this end, this chapter presents an investigation of the behaviour of propagating magnons in yttrium-iron garnet (YIG) films at millikelvin temperatures. Understanding these systems will inform the development of future

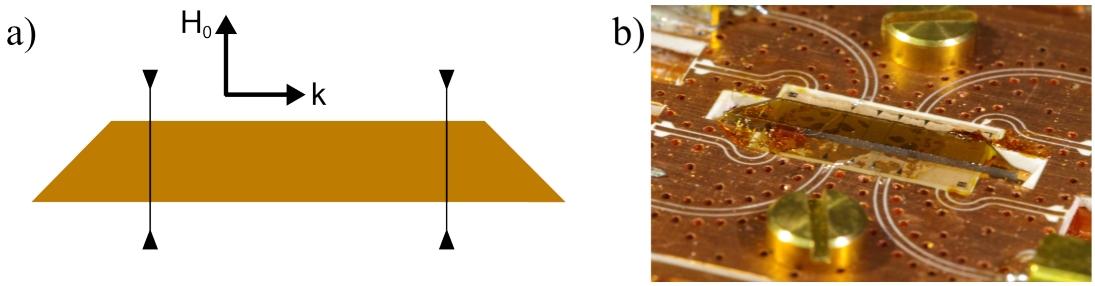

Figure 5.1: a) Schematic of the experimental setup used to study propagating surface-wave magnons. A YIG waveguide is placed on top of two microstrip antennae, one of which is used to excite spin waves and the other to detect them. b) A picture of the waveguide mounted in a sample box. GE varnish is used to glue the waveguide down to the chip which carries the antennae.

hybrid quantum devices.

5.1 Sample design

In this experiment, magnons are excited in a waveguide formed from a thin film of YIG grown on a substrate of gadolinium-gallium garnet (GGG) as described in Section 3.3.1. The waveguide is cut to shape and mounted inside a sample box on top of a sapphire chip patterned with microstrip antennae. The chip is connected to input and output lines, and the sample box mounted in the dilution refrigerator inside the bore of the superconducting magnet as described in Section 3.2.1.

As in Chapter 4, spin waves are excited via oscillating microwave-frequency magnetic fields, in this case generated by the microstrip antennae. The wavenumbers of spin waves that can be excited are limited by the width of the antennae relative to the wavelength of the spin wave. Spin waves will be excited with the wavenumber which, according to the appropriate dispersion relation, corresponds

to the frequency with which the antenna is driven. However, spin waves will be excited simultaneously across the width of the antenna. As the wavenumber of the spin waves increases, this will cause an increasing phase difference between waves excited at one edge of the antenna compared to spin waves excited at the opposite edge. When the wavenumber $k = \frac{2\pi}{w}$, where w is the antenna width, there will be complete cancellation of waves excited from the two halves of the antenna, and so no signal will be transmitted. These factors mean that the wavenumbers of spin waves it is possible to excite with a given antenna will be limited by the width of the antenna. A narrower antenna will allow a greater range of spin waves to be excited. To determine how efficiently the antenna is able to excite waves of a certain wavenumber, we take the Fourier transform of its current distribution. If this is assumed to be uniform across the antenna width, the transform corresponds to a sinc function in wavenumber, so the coupling function is given by

$$C(k) = \frac{\sin\left(\frac{kw}{2}\right)}{\frac{kw}{2}}. \tag{5.1}$$

The effect of this wavenumber-dependent coupling must be accounted for in simulations of the experiment, such as those conducted in later sections. If a signal of a given frequency is applied to the antenna, we must first calculate the corresponding wavenumber using the dispersion relation, and then use Eq. 5.1 to determine the relative efficiency of coupling to that frequency of spin wave. In reality, there will be edge effects due to the finite size of the antenna which lead to a non-uniform current distribution in the antenna. This means the coupling function is more accurately described by a Bessel function, rather than a sinc function [56]. However, in our experiments, we found this difference to have little effect on the accuracy of the simulations, so the simpler sinc function was used.

5.2 Dispersion of spin-wave pulses

In the measurements which follow, spin waves are either excited using a continuous microwave tone, or short pulses of a few tens of nanoseconds duration. In the former case, we effectively measure the transmission of the waveguide as a function of frequency and applied magnetic field. This will allow the dispersion relation of the spin waves to be 'mapped out' as a function of these variables. In the latter case, we can observe the speed with which pulses travel through the waveguide, and how they are affected by the medium as a function of the same parameters. How pulses are affected by the medium depends on the dispersion relation of spin waves.

If we consider a plane wave of the form $\sin(kx - \omega t)$, then the speed at which its wavefronts advance is given by $v_\mathrm{p} = \frac{\omega}{k}$. Light in vacuum has a linear dispersion relation given by $\omega = ck$, so the velocity of a plane light wave in vacuum is independent of its frequency. However, the dispersion relations of spin waves, given in Section 2.6.2, are non-linear, and so different wavelengths of spin wave will travel at different speeds. One consequence of this is dispersion of spin wave pulses, an effect analogous to chromatic aberration in lenses. A pulsed signal will consist of a carrier frequency shaped by an envelope function, as shown in Fig. 5.2a. Experimentally this is achieved using microwave mixers, as described in Section 3.2.2. Such a signal can be decomposed into a series of plane waves with different frequencies using a Fourier transform. The overall shape of the pulse arises from the interference pattern created by adding together all of the frequency components of the pulse. The more quickly the envelope function changes, the greater the range of frequency components in the pulse. For instance, for the

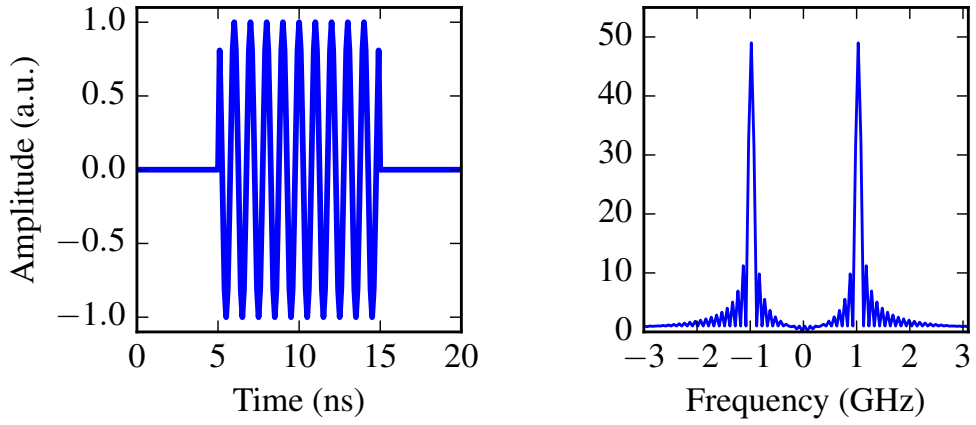

Figure 5.2: a) A 10 ns pulse with carrier frequency of 1 GHz. b) The Fourier transform of the pulse, showing the magnitudes of the frequency components from which it is formed. For a square pulse the distribution of frequencies is a sinc function centred around the carrier frequency.

square pulse shown in Fig. 5.2a, the time dependence is given by

$$F(t) = \begin{cases} \cos(\omega_c t), & t_i < t < t_i + T, \\ 0, & \text{otherwise}, \end{cases} \tag{5.2}$$

where ω_c is the carrier frequency, t_i is the start time of the pulse and T is the duration of the pulse. The Fourier transform of this function is:

$$F(\omega) = \frac{T}{2}\left(e^{i(\omega+\omega_c)\left(t_i+\frac{T}{2}\right)} \operatorname{sinc}\left(\frac{(\omega+\omega_c)T}{2}\right) + e^{i(\omega-\omega_c)\left(t_i+\frac{T}{2}\right)} \operatorname{sinc}\left(\frac{(\omega-\omega_c)T}{2}\right)\right). \tag{5.3}$$

As shown in Fig. 5.2b, this consists of the sum of two sinc functions, with a phase shift affected by the start time, centred at $\omega = \pm\omega_c$. The first minimum of the sinc functions will be at $\omega = \pm\omega_c \pm \frac{2\pi}{T}$, so as the length of the pulse decreases, the width of the frequency spectrum will increase.

When considering how a system responds to such pulses, it becomes necessary to distinguish between the *group velocity* and *phase velocity* of a wave [108]. The

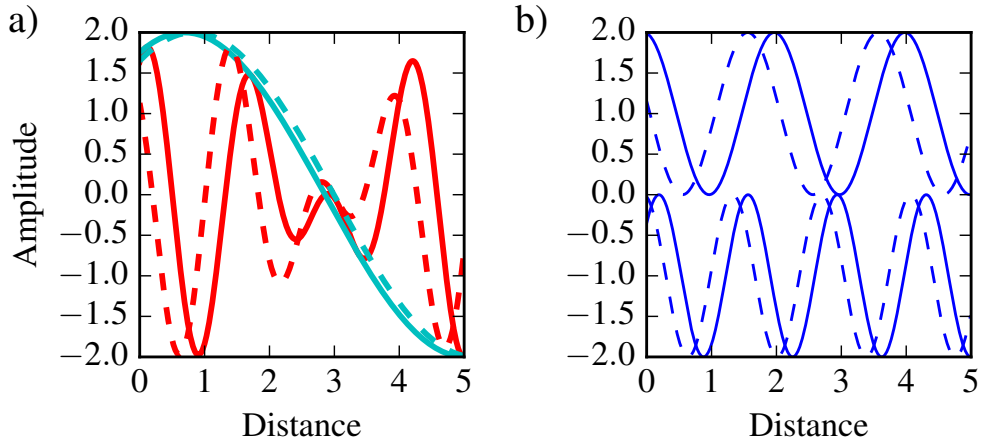

Figure 5.3: An illustration of the difference between phase and group velocity. a) A signal formed from two frequency components, shown in red, with its envelope function in cyan. Between the initial time (solid lines) and a later time (dashed lines) the envelope advances a small distance determined by the group velocity. b) The two frequency components of the signal plotted separately, at the initial time (solid lines) and the later time (dashed lines). Both components advance by their individual phase velocities, which in this case are negative. The phase velocities are different from each other and from the group velocity of the envelope in a.

phase velocity, $v_\mathrm{p} = \frac{\omega}{k}$, is the rate of advancement of the phase of a plane wave of frequency ω. The group velocity, $v_\mathrm{g} = \frac{\partial\omega}{\partial k}$, is the rate of advancement of the envelope of a wavepacket with a carrier frequency ω. The distinction between these two cases is illustrated in Fig. 5.3.

For light in a vacuum, we can see that, since $\omega = ck$, the group and phase velocities are identical, and equal to the constant c. The envelope and all the frequency components will advance at the same rate and the pulse will maintain its shape. However, in a material with a non-linear dispersion relation (also known as a *dispersive* medium), different frequencies of wave travel at different speeds. While the envelope of a pulse with a certain carrier frequency will advance at the group velocity, the individual frequency components determined from the Fourier

transform will each advance at their own phase velocity. This leads to the phases of the frequency components changing with respect to one another as the wave propagates, altering the shape of the envelope. Since YIG is a dispersive medium, we therefore expect the shape of the pulse to be distorted as it travels through the YIG waveguide. The exact nature of this distortion depends on the form of the dispersion relation, which will change as a function of magnetic field.

5.3 Experimental results

Measurements have been conducted on a number of waveguide samples all cut from the same wafer of epitaxially-grown YIG on a GGG substrate. Similar measurements were made on the waveguides with different spacings between the microstrip antennae, ranging from 7.98 mm to 1 mm. Comparing the results of measurements made in these different configurations allows us to understand which factors most strongly affect the observed results. All of the measurements in this chapter were made on waveguides cooled to the base temperature of the dilution refrigerator; approximately 10 mK.

5.3.1 Continuous wave measurements

The results of transmission spectroscopy of a waveguide mounted in surface-wave configuration are shown in Fig. 5.4 as a function of frequency and applied magnetic field. The band of spin-wave excitation is visible as a series of lines across the diagonal of the plot. The variation of signal in the background of the measurement arises from the frequency-dependent transmission of the input and output lines.

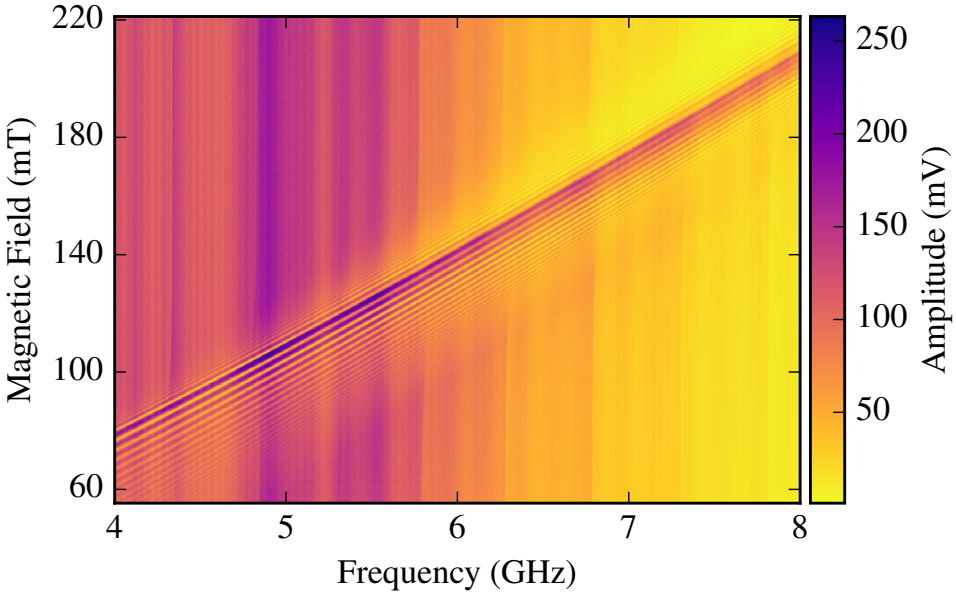

Figure 5.4: Spectroscopy of a YIG waveguide in surface-wave configuration as a function of applied magnetic field. The plot shows the voltage measured at the output antenna for a constant power applied to the input antenna.

As well as exciting spin waves in the waveguide, the oscillating voltage applied to the input antenna causes radiation of electromagnetic waves into the sample box. These electromagnetic waves propagate via the vacuum space in the sample box and are received by the output antenna. The strength of the signal transmitted in this way varies as a function of frequency due to the characteristics of the antennae and the sample box itself. However, these properties are not affected by magnetic fields, so the background varies as a function of frequency but not applied field.

The presence of the direct-coupled signal also results in the oscillations visible in the spin-wave signal. When the magnetic field and frequency are such that spin waves are allowed to propagate, the signal at the output antenna will be the

sum of the direct-coupled signal transmitted via the vacuum, and the spin-wave signal transmitted via the waveguide:

$$S_{\text{out}} = A_{\text{m}} \cos(k_{\text{m}} a - \omega t) + A_{\text{p}} \cos(k_{\text{p}} a - \omega t). \tag{5.4}$$

where A_{m} and A_{p} are the amplitudes of the spin-wave and direct-coupled signals respectively, k_{m} and k_{p} their wavenumbers, and a the separation of the two antennae. The photon wavenumber k_{p} is given by $k_{\text{p}} = \frac{\omega}{c}$ and the spin-wave wavenumber is given by the dispersion relation for surface waves (see Eq. 2.58). The expression for the total signal can be expanded into

$$S_{\text{out}} = (A_{\text{m}} \cos k_{\text{m}} a + A_{\text{p}} \cos k_{\text{p}} a) \cos \omega t + (A_{\text{m}} \sin k_{\text{m}} a + A_{\text{p}} \sin k_{\text{p}} a) \sin \omega t, \tag{5.5}$$

which can be rewritten in the form

$$S_{\text{out}} = A_{\text{t}} \sin(\omega t + \phi), \tag{5.6}$$

where

$$A_{\text{t}} = \sqrt{A_{\text{m}}^2 + A_{\text{p}}^2 + 2A_{\text{m}}A_{\text{p}} \cos(k_{\text{m}} a - k_{\text{p}} a)}. \tag{5.7}$$

The frequency of the signal is unchanged, but its amplitude has become periodic in the difference between the spin-wave and photon wavenumbers. Since the frequency dependence of the two wavenumbers is drastically different, we observe oscillations in the output signal as the frequency is swept across the spin-wave band. The amplitude of the oscillations depends on the relative size of the direct-coupled and spin-wave signals. These oscillations can be used to find an estimate of the thickness of the YIG film. We expect the spacing in wavenumber between successive peaks to be:

$$\left(k_{\text{m}}^n - k_{\text{p}}^n\right) - \left(k_{\text{m}}^{n+1} - k_{\text{p}}^{n+1}\right) = \frac{2\pi}{a} \tag{5.8}$$

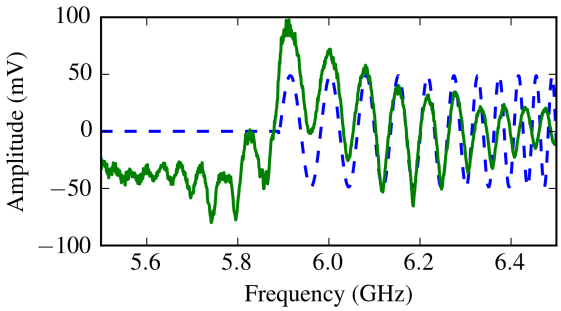

Figure 5.5: Fitting the film thickness using the interference pattern between the direct-coupled and spin-wave signals in continuous wave. The data (solid line) is a cut at 140 mT from Fig. 5.4. The fit to the interference pattern for a thickness of 7.4 µm is plotted as a dashed line.

or:

$$\Delta k_{\mathrm{m}} - \Delta k_{\mathrm{p}} = \frac{2\pi}{a}. \tag{5.9}$$

As noted in Section 2.4.3, we expect $k_{\mathrm{m}} \gg k_{\mathrm{p}}$ and since k_{m} increases much more rapidly as a function of frequency than k_{p}, we also expect $\Delta k_{\mathrm{p}} \gg \Delta k_{\mathrm{m}}$. Therefore, the difference in wavenumber between peaks is approximately

$$\Delta k_{\mathrm{m}} = \frac{2\pi}{a}. \tag{5.10}$$

Our measurement is made as a function of frequency, rather than wavenumber, and to convert between these requires the use of the dispersion relation, which depends on the film thickness. We use the dispersion relation to convert the theoretical wavenumber spacing between peaks Δk_{m} into a frequency spacing $\Delta\omega$. The film thickness used in the calculation is adjusted until the calculated frequency spacing $\Delta\omega$ matches the observed peak spacing in the data. An example of this type of fitting is shown in Fig. 5.5.

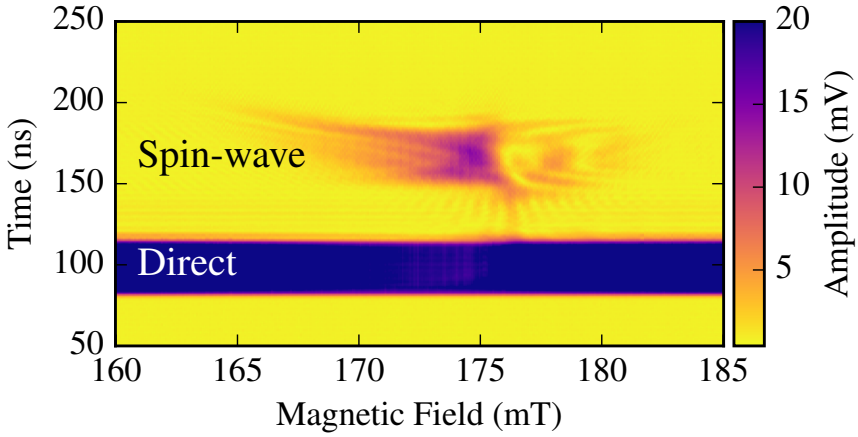

Figure 5.6: An example of a time-resolved measurement of a magnetostatic surface spin-wave pulse of approximately 30 ns. The two antenna are of width 50 µm and are separated by approximately 6.5 mm. A constant pulse carrier frequency of 7 GHz is used while the magnetic field is swept. The directly-coupled signal is seen as a bright stripe around 100 ns, and the spin-wave signal arrives at later times.

5.3.2 Time resolved measurements

In order to better analyse the behaviour of spin waves, it is beneficial to be able to separate the direct-coupled and spin-wave signals received at the output antenna. To do this, we can take advantage of the fact that spin waves in YIG travel much more slowly than light in a vacuum. In the frequency range accessible in our experiments, the maximum group velocity of spin waves is around $3.5 \times 10^6\,\mathrm{ms}^{-1}$, compared to $3 \times 10^8\,\mathrm{ms}^{-1}$ for light.

To perform these measurements, we use pulses typically between 10 and 50 ns in length generated by the AWG in the manner described in Section 3.2.2. An example of such a measurement is shown in Fig. 5.6. A number of features of the measurement are worth discussing.

The directly-coupled signal and the spin-wave signal are clearly separated

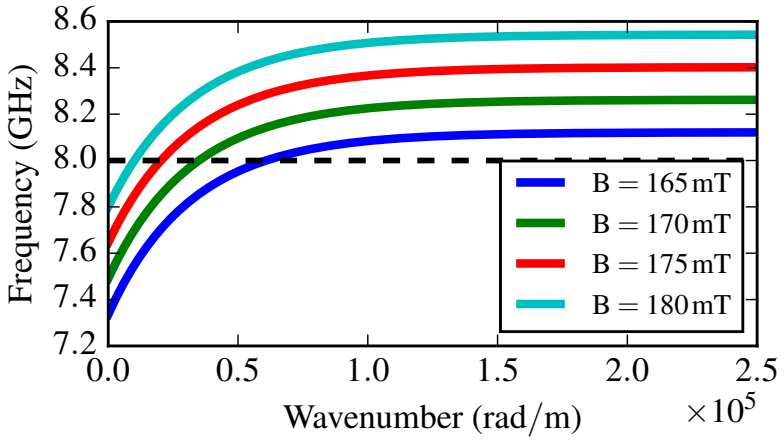

Figure 5.7: The shift in the surface-wave dispersion relation as the bias field is changed. The same frequency corresponds to a lower wavenumber at higher fields.

in time, with the directly-coupled signal appearing as a stripe at early times, and the spin-wave signal arriving at later times, as indicated in the plot. The measurement is made as a function of the magnetic bias field applied to the waveguide, rather than as a function of the carrier frequency of the pulse. The reason for this is the frequency-dependent background signal seen in Fig. 5.4. If the time-resolved measurements were made as a function of frequency, we would observe the complicating factor of a varying input power due to the frequency-dependent transmission of the setup. Making measurements as a function of field ensures a constant background signal. As the field is changed, the dispersion relation also changes, shifting the spin-wave band to higher frequencies at higher fields. This means that the same frequency corresponds to a lower wavenumber as the field increases, as illustrated in Fig. 5.7. Sweeping the field while keeping the input frequency constant will therefore still map out the dispersion relation, just as when sweeping the frequency at constant field.

At high wavenumbers, the group velocity of the spin waves is smaller than at low wavenumbers, reflected in a later arrival time of the pulse at lower fields in Fig. 5.6. The lowest field at which it is possible to observe spin waves is limited by the highest wavenumber to which the antenna can couple, dependent on its width, and by the propagation loss of spin waves in the film. As the field is increased, the band will eventually shift high enough in frequency that the pulse carrier frequency reaches the $k = 0$ limit, or ferromagnetic resonance (FMR) frequency. If the field increases above this point, the input signal can no longer excite spin waves, since its frequency will be below the lower limit of the spin-wave band. This point is seen as the abrupt reduction in signal around 177 mT in Fig. 5.6. The fact that the signal does not disappear entirely above this field will be discussed later.

Using the AWG, we can generate a variety of pulse shapes, or even pulse trains. In Fig. 5.8 we demonstrate the capabilities of the setup with a series of three closely-spaced square pulses, and a pulse with a Gaussian envelope. The former is indicative of the potential for this type of setup to be used for transmission in information technology applications; the three pulses are clearly resolved, and separated from the direct-coupled signal. The latter is an illustration of how the spectral content of a pulse determines its behaviour in a dispersive medium, as discussed in Section 5.2. A Gaussian pulse has a Gaussian frequency distribution, rather than the sinc distribution of a square pulse. Typically this means the range of frequencies contained in a Gaussian pulse is much smaller than that of a square pulse with a similar width, and this is reflected in the much simpler structure of the pulse in Fig. 5.8b compared to that in Fig. 5.6.

As in Chapter 4, it was also possible to make measurements when, on average,

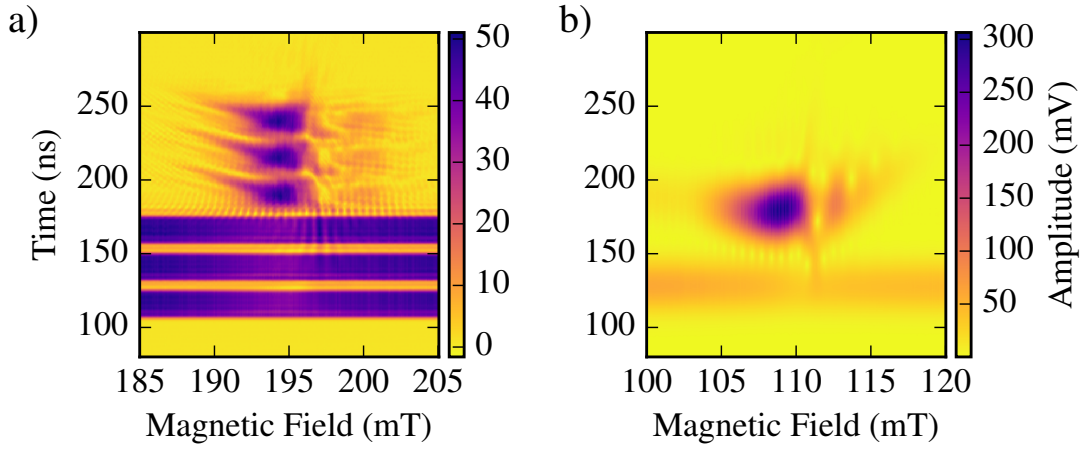

Figure 5.8: a) A train of three 15 ns square pulses separated by 5 ns gaps. b) A pulse with a Gaussian envelope shape, rather than a square envelope.

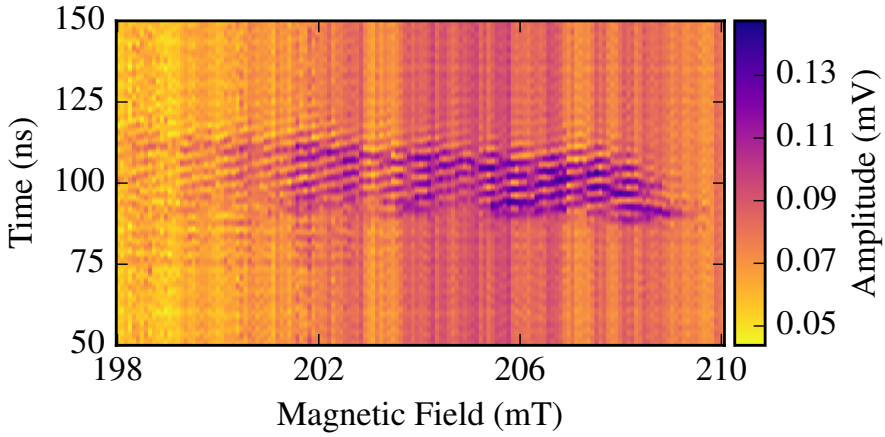

Figure 5.9: A time-resolved measurement with, on average, fewer than one magnon present in the waveguide. A 20 ns pulse was used at an input power of -130 dBm and carrier frequency of 8 GHz, corresponding to approximately 0.4 magnons per pulse.

fewer than one excitation was present in the system. By applying a high level of attenuation to the output of the mixer, the input power was reduced to a level such that each microwave pulse contained, on average, less than one photon when it entered the sample box. The result of such a measurement is shown in Fig. 5.9.

5.4 Simulating time-resolved measurements

In order to test our understanding of the effects of the YIG film on spin-wave propagation, we set out to reproduce the observed experimental results through simulation. To do so, we simulate the effects of the YIG medium, according to the dispersion relation, and the measurement apparatus on the input signal using Fourier decomposition. The simulation is performed in several stages.

Input generation

The input signal is generated to match the experimental input pulse by shaping a continuous wave carrier frequency with an envelope. The sampling frequency of the simulation at this stage is set to 40 GHz, high enough to adequately resolve all the frequencies expected in the measurement. The time-domain signal is then Fourier transformed to split it into its frequency components. For each frequency component, the corresponding spin-wave wavenumber is calculated, and the amplitude of the component modified to take account of the efficiency with which the input antenna can couple to waves of that wavenumber (as described in Section 5.1).

Propagation in the waveguide

Once the signal is 'inside' the YIG film, each frequency component is multiplied by a factor

$$A(\omega) = e^{-i\omega t_\mathrm{d} - \frac{t_\mathrm{d}}{t_\mathrm{l}}}, \tag{5.11}$$

where $t_\mathrm{d} = \frac{a}{v}$ is the time delay for the frequency component with velocity v, and t_l is a characteristic loss time. The first term in the exponential represents

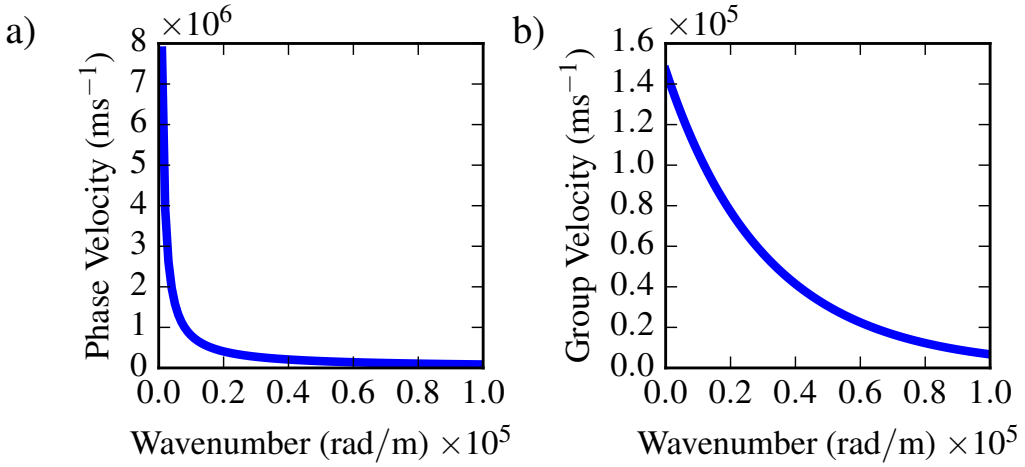

Figure 5.10: a) The phase velocity as a function of wavenumber for spin waves in yttrium-iron garnet in an applied field of 180 mT. b) The group velocity as a function of wavenumber under the same conditions as a.

the phase shift acquired by the wave over the length of the waveguide, while the second term represents the decay of its amplitude over the same distance. It is important to emphasise that, since this loss time is applied to each Fourier component individually based on the phase velocity, it is typically much smaller than a loss time associated with the pulse envelope. This is because, as can be seen in Fig. 5.10, the phase velocity of a Fourier component with a certain wavenumber is much larger than the corresponding group velocity. In other words, a loss time of 1 ns in this simulation does not indicate that we would expect the pulse amplitude to reduce to $\frac{1}{e}$ of its initial value after the pulse has been propagating for 1 ns.

Another important consequence of the difference between group and phase velocity is in calculating the dispersion of the pulse. If we set v equal to the group velocity of the carrier frequency, then the simulation should correctly predict the arrival time of the pulse. However, it will not predict any dispersive effects of

the medium due to the differing speeds of individual frequency components. To account for these effects, we must instead use the phase velocity calculated at the frequency of each individual frequency component. The difference between these two cases is illustrated in Fig. 5.11 for a simulation of the measurement in Fig. 5.6.

The phase velocities for surface and backward waves are calculated according to the dispersion relations in Eq. 2.58 and Eq. 2.56 respectively, which are for isotropic, semi-infinite films. However, our films exhibit magnetocrystalline anisotropy, and so the applied field used to calculate the phase velocity must be adjusted to account for this. Assuming that the only significant source of anisotropy is the first-order cubic anisotropy, the FMR frequency for an in-plane magnetised, [111]-oriented YIG film is [109]

$$\omega_{\mathrm{FMR}} = -\gamma\mu_0\sqrt{H_0(H_0 + M_\mathrm{s} - H_{\mathrm{c1}}) + \frac{1}{9}(10\beta H_{\mathrm{c1}}^2 - (18 + 10\beta)H_{\mathrm{c1}}^2\cos 6\phi)}. \tag{5.12}$$

where H_{c1} is the first-order cubic anisotropy field for YIG below 4 K, as given in Table 3.1, $\beta = \frac{H_0}{H_0+M_\mathrm{s}}$ and ϕ is the in-plane angle between the applied field and the crystal $[\bar{1}10]$ axis. The value of ϕ for surface waves will differ from that for backward waves by 90°. In our case, the alignment of the waveguide means that appropriate values are $\phi = 30°$ for backward waves and $\phi = 0°$ for surface waves ($\phi = 0°$ is equivalent to $\phi = 120°$ due to the six-fold symmetry of the expression). To find the field adjustment necessary, we find a ΔH such that

$$\omega_{\mathrm{FMR}}(H_0 + \Delta H, H_{\mathrm{c1}} = 0) = \omega_{\mathrm{FMR}}(H_0, H_{\mathrm{c1}}), \tag{5.13}$$

where H_0 is chosen to be the field at which the isotropic FMR frequency matches

the pulse carrier frequency

$$\omega_{\mathrm{FMR}}(H_0, H_{\mathrm{c1}} = 0) = \omega_{\mathrm{c}}. \tag{5.14}$$

In other words, adding this ΔH to the field used to calculate the dispersion will make the isotropic FMR frequency align with the anisotropic FMR frequency. This scheme is used to account for the anisotropy due to its computational simplicity, which makes performing these simulations possible in a reasonable time. Calculations using more comprehensive considerations of anisotropy on the entire dispersion relation [110, 111] have demonstrated that in our regime, the most significant effect on the dispersion relation is a shift in frequency. While there will be some higher-order effects on the spin waves due to the anisotropy, the good agreement we observe between data and simulation suggests that these are likely to be small. In the case of surface waves, there is also an additional adjustment to the applied field due to the GGG substrate. This will be discussed in more detail in Chapter 6.

Output capture

Once the effects of propagation have been considered, the influence of the measurement apparatus is then simulated. The frequency-domain signal is again multiplied by the antenna coupling function to account for the response of the output antenna, before an inverse Fourier transform returns the signal to the time domain. Firstly, the analogue filtering and down-conversion described in Section 3.2.2 are simulated by multiplication with a continuous wave mixing tone, and the use of digital filters. Then, the signal is down-sampled to the sampling frequency of 2.5 GHz used in experiment. Finally, the digital down-conversion and filtering

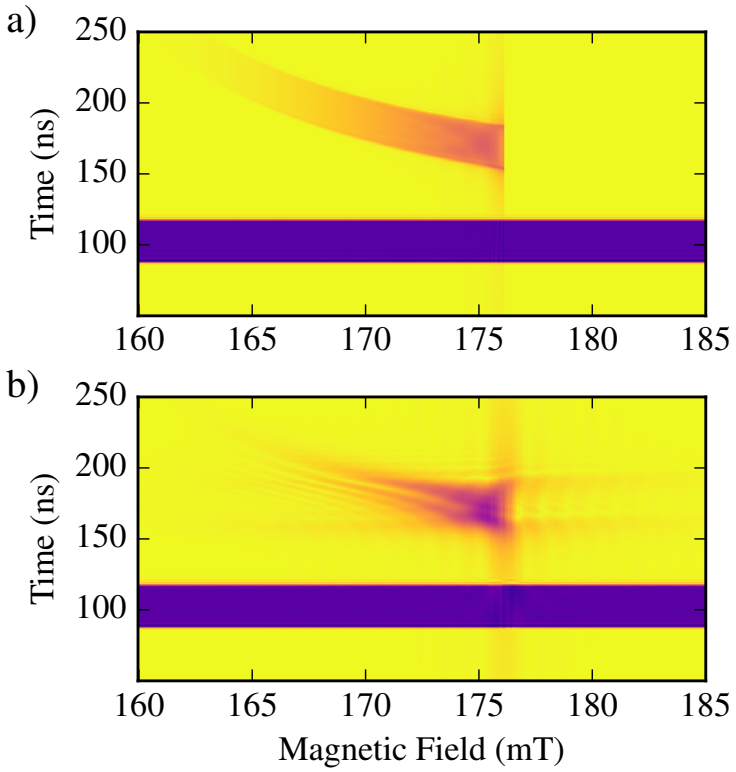

Figure 5.11: a) A simulation of the measurement in Fig. 5.6 using the group velocity of the carrier frequency to advance each frequency component of the input signal. The YIG response is calculated for a saturation magnetisation of $196\,\mathrm{kAm^{-1}}$, and a film thickness of $10\,\mu\mathrm{m}$. The width and spacing of the antennae in the simulation matches that used in the experiment. The arrival time of the pulse is adequately predicted, but not its dispersion. b) A simulation using the same parameters, but phase velocity instead of group velocity. As well as predicting the arrival of the pulse, the shape is much more convincingly reproduced.

described in Section 3.2.2 are performed in the same manner as on the computer used for data capture. All of the filtering steps are reproduced using Butterworth filters, which have a maximally flat response in their passband [112].

5.5 Discussion

In Fig. 5.12 and Fig. 5.13 we plot direct comparisons between measurement and simulation for surface and backward-volume waves respectively. The measurements shown are conducted on the same sample, which was measured in the two different configurations. The parameters used in the simulations are summarised in Table 5.1. In these experiments, the antenna spacing was much smaller than in the measurement shown in Fig. 5.6, and so the direct-coupled and spin-wave pulses overlap at early times. The interference in this region shows the same behaviour seen in the continuous wave measurement in Fig. 5.4. Since the damping is determined by the propagation delay, the reduced propagation distance also allows us to see spin waves over a much wider range of fields than in Fig. 5.6, so a larger proportion of the spin-wave band is sampled. This is because spin waves with higher wavenumbers, existing further from the FMR in field, have smaller velocities, and so experience higher damping than those close to the FMR. The main features of this comparison are typical of simulations made of other measurements with different frequencies, pulse shapes and antenna spacings.

The simulation reproduces most of the major features of the measurement. The arrival time of the pulse and its dispersion close to the FMR are correctly predicted. The interference between the direct-coupled signal and the spin-wave signal is also reproduced where the two pulses overlap. It is worth noting that

the applied field needed to match the carrier frequency to the FMR frequency is different in the two measurements, despite the fact that, according to the expressions in Chapter 2, surface and backward-volume waves should have the same FMR frequency for a given applied field. This is due in part to the effect of magnetocrystalline anisotropy, discussed in the previous section. For surface waves, there is also an additional shift due to the GGG substrate, which is magnetic at low temperatures. This substrate effect is discussed in more detail in Chapter 6.

However, there are several features of the data which are not correctly reproduced. The arrival times of the pulse at high wavenumbers is different in measurement than is predicted by the simulation. In the case of surface waves in Fig. 5.12, the high wavenumber (low field) arrival time is too long in simulation, and the pulse becomes much more steeply curved in the simulation than in the data. This suggests that the true dispersion relation is less steeply curved at high wavenumbers than the theoretical expression (plotted in Fig. 2.11) that is used in the simulation. Conversely, for backward waves in Fig. 5.13 the high wavenumber (high field) arrival time is too short, although the difference is less pronounced

Table 5.1: Summary of parameters used for the simulations of pulsed measurements in surface-wave (Fig. 5.12) and backward-wave (Fig. 5.13) configurations.

Parameter		**Value**
Saturation magnetisation	M_s	$197.4\,\mathrm{kAm^{-1}}$
First-order anisotropy field	H_c1	$-20.13\,\mathrm{kAm^{-1}}$
Film thickness	d	$11\,\mu\mathrm{m}$
Loss time	t_l	$1\,\mathrm{ns}$ (surface wave) $0.4\,\mathrm{ns}$ (backward wave)
Antenna separation	a	$1\,\mathrm{mm}$
Antenna width	w	$25\,\mu\mathrm{m}$
Pulse carrier frequency	ω_c	$6\,\mathrm{GHz}$
Pulse length	T	$30\,\mathrm{ns}$

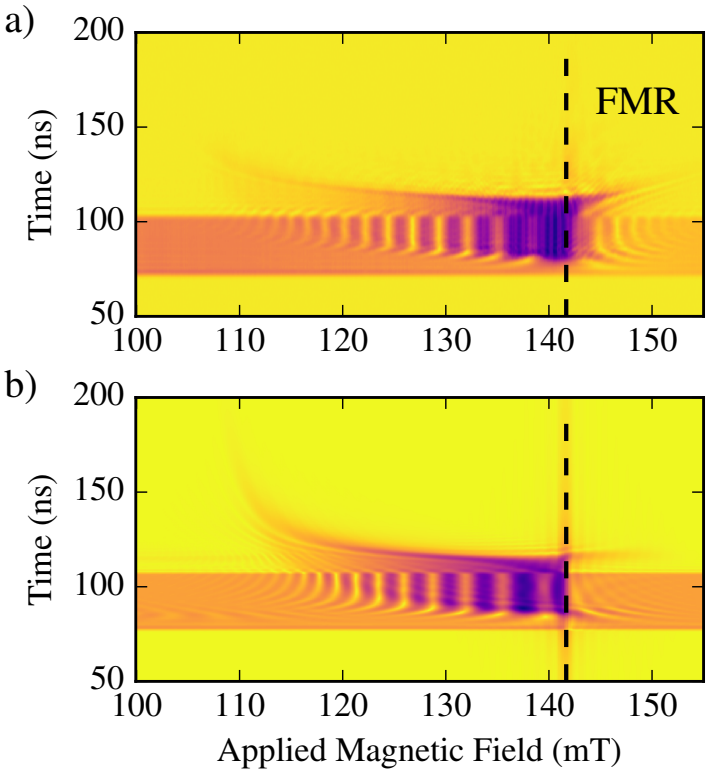

Figure 5.12: a) A time-resolved measurement of the propagation of a surface spin-wave pulse of length 30 ns and carrier frequency 6 GHz. The width of the input and output antennae was 25 µm and their separation was 1 mm. b) A simulation of the measurement in a, using $M_s = 197\,\mathrm{kAm^{-1}}$ and a film thickness of 11 µm. The applied magnetic field used in the simulation has been adjusted to account for the effects of the GGG substrate, which is discussed in more detail in Chapter 6. The characteristic loss time t_l is 1 ns.

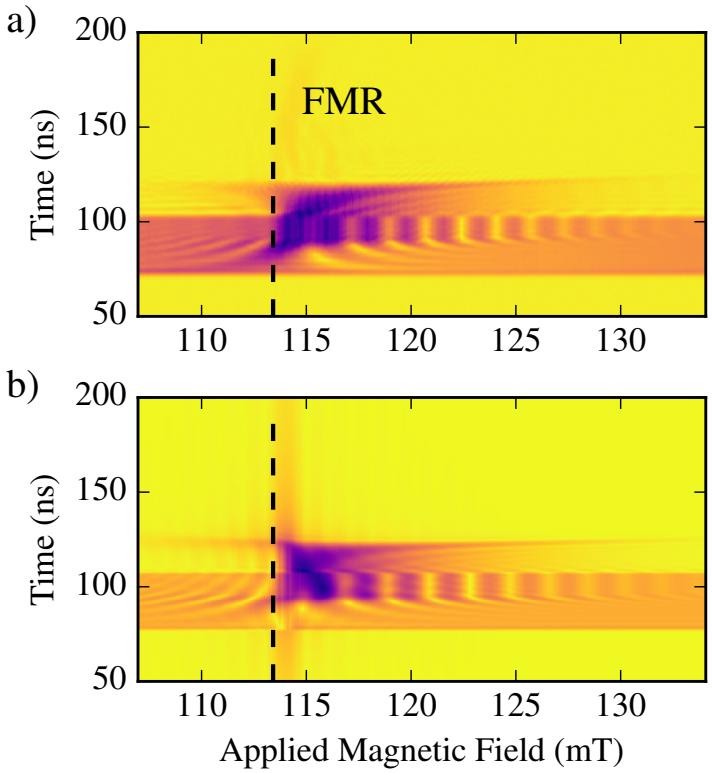

Figure 5.13: a) A time-resolved measurement of the propagation of a backward-volume spin-wave pulse of length 30 ns and carrier frequency 6 GHz. The antenna configuration is the same as in Fig. 5.12. b) A simulation of the measurement in a, using $M_s = 197\,\mathrm{kAm^{-1}}$ and a film thickness of 11 µm. The characteristic loss time t_l is 0.4 ns.

than in the case of the surface waves.

Also apparent is the signal at fields above the FMR in the surface-wave case, which is much stronger in measurement than predicted by simulation. When the field is increased such that the frequency of the input signal is below the FMR frequency, it should not be possible to excite spin waves. Some signal is still seen in the simulation due to the wide frequency spectrum of the square pulse. As seen in Fig. 5.2, the frequency spectrum of a square pulse contains components significantly above and below the carrier frequency of the pulse. If the carrier

frequency is only slightly below the FMR frequency, there will still be frequencies in the input signal above the FMR frequency which can excite spin waves. This effect accounts well for the signal below the FMR in the backward wave case. However, the additional surface-wave signal at lower frequencies (higher fields) than the FMR is also seen in continuous wave measurements, where the input frequency spectrum is extremely narrow (c.f. the oscillations seen in Fig. 5.4 at lower frequencies and higher fields than the brightest FMR line). Such additional signal has been seen in comparable room-temperature measurements in the literature [17, 113] but was not commented on. The presence of this signal in continuous-wave measurements suggests the existence of an additional band of surface spin waves existing at lower frequencies than the theoretical lower limit imposed on surface waves by the FMR frequency.

A number of effects have been considered as potential explanations for these observations, including non-uniformity of the applied magnetic field, magnetocrystalline anisotropy [111], misalignment of the waveguide from the ideal surface-wave geometry [114], proximity of ground planes to the waveguide [71, 72], and finite width of the waveguide [115]. For realistic experimental parameters, none of these effects were able to satisfactorily reproduce the observed phenomena.

5.6 Conclusions

In this chapter we have performed measurements of propagating surface and backward-volume spin waves in a YIG waveguide. A wide variety of continuous-wave and time-resolved measurements have been conducted at different frequen-

cies and magnetic fields with different antenna spacings. Using these measurements, we have been able to determine the parameters of the waveguide and its behaviour. We have built simulations which satisfactorily reproduce the main features of the observed results. We have also been able to rule out a number of different possible explanations for the features of the data which are not predicted by the simulations. Finally, measurements have been made of spin waves in the quantum limit of fewer than one magnon present in the system. This is an important step to building new devices which operate in the limit of single excitations for quantum information applications. The findings in this chapter are reported in [116].

Chapter 6

Temperature-dependent effects of GGG substrate

In Chapter 5 it was mentioned that the applied magnetic field in Fig. 5.12 had to be corrected for the effects of the substrate on which the YIG film is grown. This substrate is gadolinium-gallium garnet ($Gd_3Ga_5O_{12}$, GGG) which, as described in Chapter 3, is almost universally used for this purpose on account of the close matching of its lattice parameter to that of YIG, which enables high-quality, low-strain films of YIG to be grown. The gadolinium ions in GGG are magnetic, having spin $\frac{7}{2}$, and exhibit relatively strong antiferromagnetic coupling with a Curie-Weiss temperature of 2.3 K [117]. However, the system is also frustrated due to a combination of the magnetic coupling between ions and their geometric arrangement in the lattice. The effect of the frustration is so strong that GGG displays no long-range order even significantly below the Curie-Weiss temperature [118]. This has made it an archetypal example of a frustrated magnetic system, and it has been studied for many years [119].

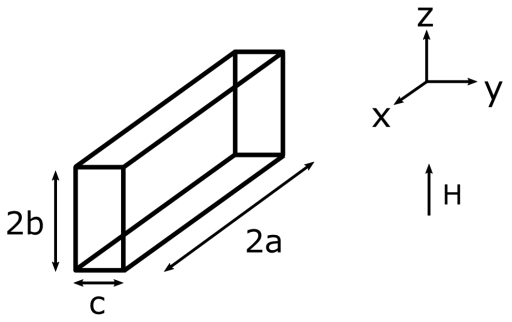

Figure 6.1: The substrate geometry considered to determine the field experienced by the YIG.

Most past investigations into spin-wave propagation have been done at room temperature, where GGG is simply a very weak paramagnet. There has therefore been little need to take account of the properties of the substrate when trying to account for experimental results. However, at the very low temperatures at which our experiments have been conducted, we have observed significant deviations from the theoretical spin-wave behaviour, which are most likely explained by effects arising from the GGG substrate. This chapter describes how the presence of magnetism in the GGG substrate is expected to affect the propagation of spin-waves in YIG. It then presents a thorough characterisation of this behaviour in the temperature range from 10 mK to 10 K.

6.1 Field due to a magnetic substrate

Consider a magnetic material with magnetisation $\mathbf{M}$ and a shape given by the geometry pictured in Fig. 6.1. The origin of the coordinate system is taken to be at the centre of one of the large faces, so the material extends from $-a$ to a in the x-axis, $-b$ to b in the z-axis, and $-c$ to 0 in the y-axis. If a magnetic dipole

$\mathbf{m}$ is placed at the origin, the field it produces at a position $\mathbf{r}$ is given by

$$\mathbf{H_m} = \frac{1}{4\pi}\left(\frac{3(\mathbf{m}\cdot\mathbf{r})\mathbf{r}}{r^5} - \frac{\mathbf{m}}{r^3}\right). \tag{6.1}$$

To find the total field produced by the magnetic material in Fig. 6.1 we must integrate the contributions from all of the moments within the material. The field it produces at the origin is therefore given by

$$\mathbf{\Delta H} = \frac{1}{4\pi}\int_{-a}^{a}\int_{-b}^{b}\int_{-c}^{0}\left(\frac{3(\mathbf{M}\cdot\mathbf{r})\mathbf{r}}{r^5} - \frac{\mathbf{M}}{r^3}\right)\mathrm{d}x\,\mathrm{d}y\,\mathrm{d}z. \tag{6.2}$$

Taking the magnetisation $\mathbf{M}$ of the material to be uniform and aligned with the z-axis, we have $\mathbf{M} = M\hat{\mathbf{z}}$. The only non-zero component of the field at the origin is the z component, given by

$$\Delta H_z = \frac{M}{4\pi}\int_{-a}^{a}\int_{-b}^{b}\int_{-c}^{0}\left(\frac{3z^2 - r^2}{r^5}\right)\mathrm{d}x\,\mathrm{d}y\,\mathrm{d}z, \tag{6.3a}$$

$$= -\frac{M}{\pi}\tan^{-1}\left(\frac{ac}{b\sqrt{a^2+b^2+c^2}}\right). \tag{6.3b}$$

We now consider the implications of this result for measurements of a YIG film grown on top of a GGG substrate. At low temperatures, the GGG substrate will become magnetic when placed into an external field. If the external field $\mathbf{H_0}$ is aligned along the z-axis such that $\mathbf{H_0} = H_0\hat{\mathbf{z}}$, then the effective field experienced by the YIG film will be the sum of the externally applied magnetic field and the field produced by the substrate:

$$H_{\text{eff}} = H_0 + \Delta H_z. \tag{6.4}$$

Since ΔH_z is negative, the film experiences a lower field than would be expected from the current in the superconducting magnet. This means that the spin-wave band will appear at higher applied fields than would be expected if we used the applied field strength directly in our calculations.

In deriving this result, we have assumed that the substrate is uniformly magnetised throughout its whole volume and taken the field at the centre of one of its large faces to be representative of the field experienced by the entire YIG film. In fact, some demagnetising field will exist which changes the magnetisation through the volume. However, calculating this demagnetising field [120], we find that the magnetisation is only significantly non-uniform close to the edges of the substrate. The large majority of the film will experience a uniform field due to the substrate, so we can expect this behaviour to dominate. The YIG film is also very thin on the scale of the substrate dimensions, and so we expect any variation in the field through the thickness of the film to be small.

It is also worth noting the difference in the expected magnitude of this effect for surface and backward-wave configurations of the sample. In Fig. 6.1 we illustrated the surface-wave configuration, where the applied field is in the plane of the YIG film and perpendicular to its long axis, along which the spin-waves will propagate. This means that for the surface-wave case $a \gg b$ in Eq. 6.3b. However for backward-wave measurements, the sample is rotated by 90° so that the long axis of the film is parallel to the applied field. In this case, a and b should be interchanged in Eq. 6.3b, reducing the size of the argument by a factor $\left(\frac{b}{a}\right)^2$. For typical values of a and b in our experiments, this would reduce the size of ΔH_z more than 50-fold.

The difference between these two cases is illustrated in Fig. 6.2 and Fig. 6.3. In the backward-wave configuration the substrate effect is very small and the theoretical FMR frequency is a good match to the observed position of the high-frequency edge of the spin-wave band. However, in the surface-wave configuration there is a substantial reduction in the effective field experienced by the YIG, due

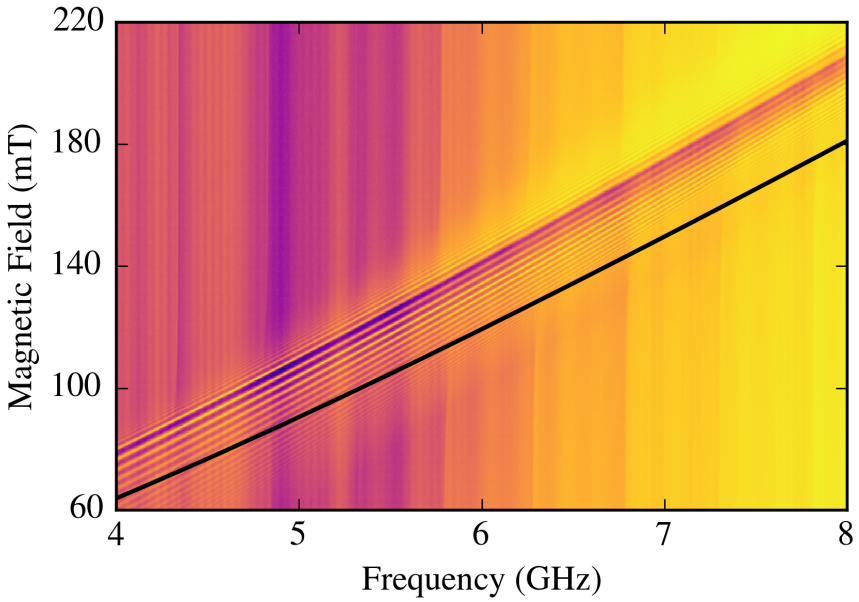

Figure 6.2: Continuous wave measurement at 25 mK in surface-wave alignment from Fig. 5.4. The theoretical FMR frequency (with no substrate corrections, but including the effect of cubic anisotropy according to [109]), is plotted on top and deviates significantly from the observed position of the band edge.

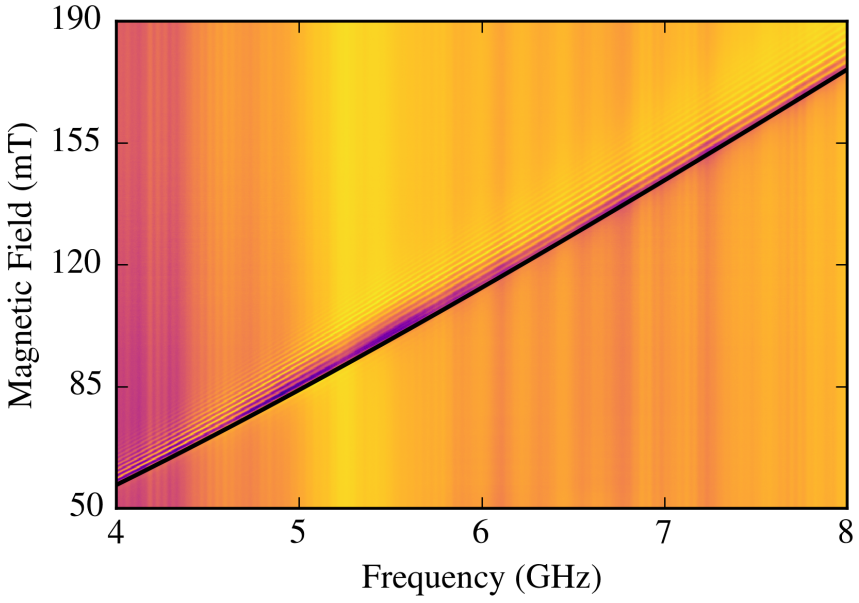

Figure 6.3: Continuous wave measurement at 25 mK in backward-wave alignment. The theoretical FMR frequency (corrected for the cubic anisotropy according to [109]) matches the position of the band edge, since in this geometry the substrate effects are much smaller than for surface waves.

to the substrate. This causes the low-frequency edge of the band to appear at higher applied fields than predicted by the dispersion relation. In both Fig. 6.2 and Fig. 6.3, the theoretical position of the FMR frequency is calculated using Eq. 5.12 and so accounts for the effects of the first-order cubic anisotropy field in YIG.

6.2 Temperature dependence of substrate magnetisation

We have noted that the theoretical effect of a magnetised substrate is expected to be significantly larger for a film in the surface-wave configuration than for a film in the backward-wave configuration. This prediction was supported by the fact that the theoretical FMR frequency, without any correction for the substrate effects, matched the experimental data much more closely for backward waves than for surface waves. We have also observed a clear change in the magnitude of the substrate effect as a function of the temperature of the sample, as demonstrated in Fig. 6.4. This suggests a change in the magnetic properties of the substrate in this temperature range, since the saturation magnetisation and magnetocrystalline anisotropy of YIG are not expected to change significantly in this regime [86, 105]. The fact that the magnitude of the change as a function of temperature is so much smaller for backward waves than for surface waves is also further evidence that the substrate effect is much weaker for backward waves.

As mentioned above, the magnetic properties of GGG have been studied for many years [119]. At high temperatures it is paramagnetic [121], displaying Curie-

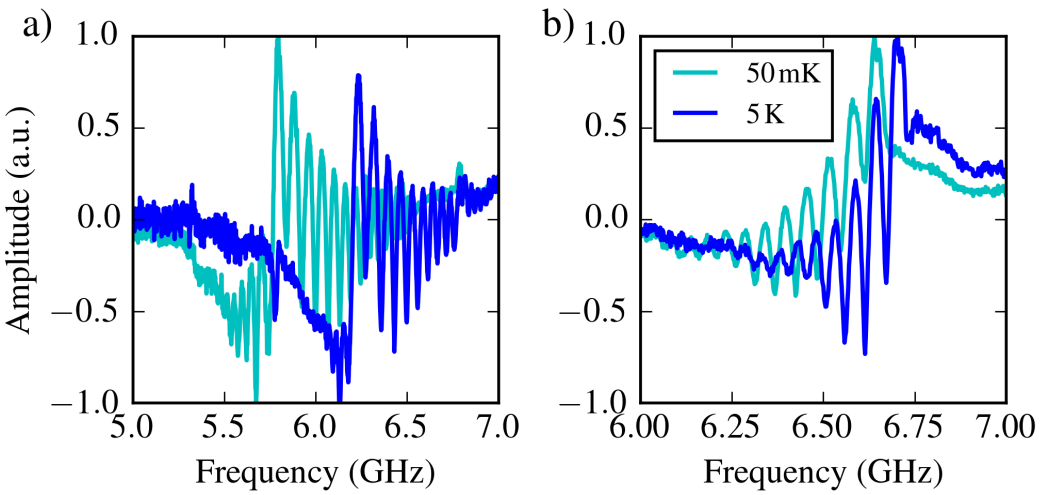

Figure 6.4: A comparison of the transmission of the waveguide as function of frequency at 135 mT at 50 mK and 5 K. a) In the surface-wave configuration, the spin-wave band is shifted up in frequency by approximately 450 MHz between 50 mK and 5 K, indicating a reduction in the field due to the substrate. b) In the backward-wave configuration the shift in frequency between 50 mK and 5 K is only 60 MHz, indicating that the substrate effect is much smaller in this alignment.

Weiss susceptibility above 1.5 K [122]. At low temperatures (below ≈ 0.4 K) it develops long-range antiferromagnetic order under applied fields of around 1 T [123]. In our experiments, we use low magnetic fields (50 to 250 mT) and low temperatures (10 to 50 mK). In this regime the properties of GGG are still not well understood, although possible explanations include a spin glass state [122] or a more exotic state involving locally correlated spin loops [124]. The measurements possible with our apparatus are not sufficient to fully characterise the magnetic state of GGG. However, we can still quantify the effect that it has on the spin waves we observe. To do this, we perform spectroscopy measurements in the surface-wave configuration, similar to those shown in Fig. 6.2, at a range of temperatures between 10 mK and 10 K. At each temperature, the difference

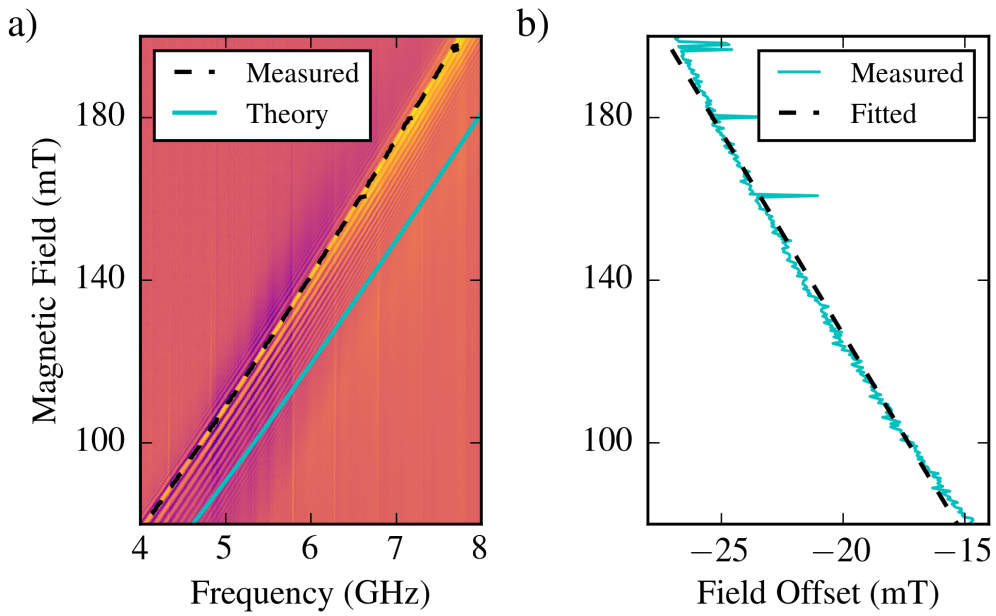

Figure 6.5: a) A spectroscopy measurement on a YIG film in the surface-wave configuration at approximately 40 mK. On top of the data are plotted the measured position of the band edge, as detected by a peak-finding algorithm, and the theoretical position of the band edge. b) The measured difference between the applied field and the field experienced by the YIG film as a function of the applied field. The dashed line is a linear fit to the difference according to Eq. 6.5

between the observed and theoretical frequencies of the surface spin-wave band edge is measured as a function of magnetic field. This is converted into a difference between the applied field and the field experienced by the YIG. We fit the difference using a linear relationship of the form

$$\Delta H_z = A_0 + A_1 H_0. \tag{6.5}$$

An example of this fitting is shown in Fig. 6.5 for a measurement taken at approximately 40 mK. In this case, the offset A_0 is -7.3 mT and the gradient A_1 is -0.1.

The results of the fitting are plotted as a function of temperature in Fig. 6.6

for two different samples. Values of A_0 and A_1 for 50 mK and 5 K are given in Table 6.1. Both were cut from the same wafer of epitaxially-grown YIG on a [111]-oriented GGG substrate, and were measured in the surface-wave configuration. Although the absolute values of the shift are slightly different between the two samples, the trend as a function of temperature is very similar. Below about 300 mKboth the gradient and the offset are nearly constant. Above this temperature, the magnitude of the offset begins to fall, reaching zero at around 10 K. The gradient begins to decrease significantly at a slightly higher temperature of 1 K, and is still non-zero at 10 K.

From Eq. 6.3b we can see that the field shift ΔH_z should be directly proportional to the magnetisation of the substrate. If the substrate were purely paramagnetic, we would expect its magnetisation to be proportional to the applied field, since for a paramagnet $\mathbf{M} = \chi\mathbf{H}$. This would imply $A_0 = 0$ and $A_1 \neq 0$ in Eq. 6.5. This is indeed the situation at the high end of the measured temperature range. However, the fact that at low temperatures ΔH_z is best fitted with a non-zero offset A_0 is intriguing, since it suggests some type of remnant magnetisation present in the substrate. This offset begins to disappear at temperatures above 0.3 K, which is roughly consistent with the temperature range in which GGG is known to become paramagnetic [121].

Table 6.1: Fit parameters for substrate effect

Temperature	**Parameter**	**Sample 1**	**Sample 2**
50 mK	A_0(mT)	-9.3	-7.8
	A_1	-0.084	-0.10
5 K	A_0(mT)	-1.9	-1.0
	A_1	-0.049	-0.052

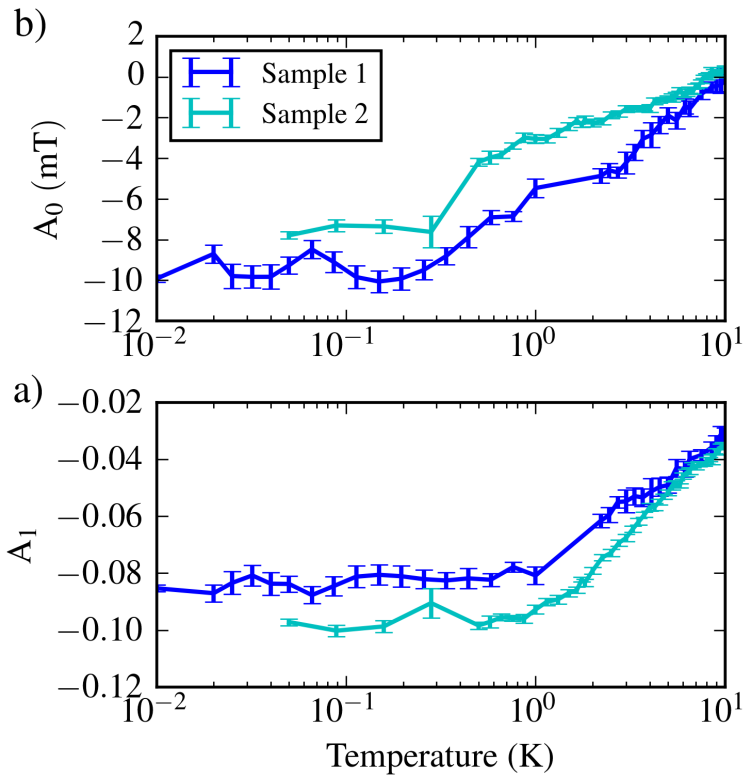

Figure 6.6: The characteristics of the field shift ΔH_z experienced by a YIG film in the surface-wave configuration due to the magnetised GGG substrate. The offset A_0 and gradient A_1, as defined in Eq. 6.5, are plotted as a function of the sample temperature. Lines are guides to the eye. Both offset and gradient are constant at low temperatures, and decrease significantly at temperatures approaching 10 K.

6.3 Conclusions

In this chapter we have established the existence of a shift in the position of the spin-wave band at low temperatures, which disappears as the temperature is increased to around 10 K. Previous studies have identified an effect on the FMR frequency of epitaxially-grown YIG films due to their GGG substrate [125, 126]. However, these studies were largely conducted in a much higher temperature range than that investigated here, and assumed that the substrate was purely

paramagnetic in character. As we have seen, this assumption is inadequate in the temperature range investigated in this chapter.

While the linear fit to ΔH_z used in Eq. 6.5 was adequate at the fields used in these measurements, it may be that measurements at lower fields would reveal a more complex relationship between the substrate magnetisation and the applied field. This would be an interesting subject for future experiments.

Using the results presented here, we are able to determine the field experienced by the YIG film for a given applied external field, and to predict how this will change with temperature. This allows us to correct our simulations of spin-wave propagation and more accurately model how spin-waves in YIG will behave at very low temperatures.

Chapter 7

Outlook

This thesis has described some initial experiments in the emerging field of quantum magnonics. We have examined the theory of spin-waves which underlies the study of magnonics, and described the experimental techniques required to extend this study into the quantum regime. The experiments presented here provide the foundation for future work by cementing our understanding of the physical processes that govern the behaviour of magnons, and demonstrating that this can be extended to the challenging conditions necessary for more complex quantum investigations.

In Chapter 4, we examined the behaviour of resonant magnon modes, and how they interact with the types of structures commonly used to manipulate quantum information devices. We demonstrated that coupling was possible to a wide variety of different resonant modes, and identified a scheme for predicting which of those modes are most strongly coupled. In Chapter 5, propagating spin waves were investigated using high time resolution measurements of short magnon pulses in a waveguide. We showed that propagation of magnon pulses over dis-

tances approaching a centimetre is possible even at these very low temperatures, and demonstrated that we could detect these pulses even when the power levels were reduced such that, on average, only single magnons were present in the system. Computer simulations were developed which successfully model the major characteristics of the magnon dispersion. Finally, in Chapter 6 we demonstrated that the commonly used substrate GGG has a significant effect on the magnetic environment of the waveguide at low temperatures, and provided a thorough characterisation of the nature of this effect as a function of temperature.

The results of these investigations, as well as the techniques for experimental method and computer simulation that have been developed, will inform the design of future experiments in the field. A particular goal is the interface of superconducting quantum bits with magnonic devices. While investigations have been made into the interaction of such devices with resonant magnon modes [44, 45], there have been no corresponding investigations into propagating excitations such as those studied in Chapter 5. Superconducting qubits could be used as single magnon sources and sinks to probe the dynamics of single magnons in these systems. Magnonic elements may also be able to provide new functionality in quantum information systems. Structures such as magnonic crystals [19, 127] could provide interesting switching and filtering functionality in such systems.

Having gained an understanding of how magnonic devices can be brought into the challenging regime of low-power, low-temperature physics, we can move forward with new and exciting investigations which will bring the study of magnonics into the quantum regime.

Acknowledgments

> We can forgive a man for making a useful thing as long as he does not admire it. The only excuse for making a useless thing is that one admires it intensely
>
> Oscar Wilde, *The Picture of Dorian Gray* (1891)

Thanks are due to all of the people who have been a part of my time in the Clarendon Laboratory. Thanks to Alexy for giving me the opportunity to work in her lab. To Arjan, who has enlivened my time here with many coffee breaks and kicking, as well as teaching me lots of stuff about microwaves and fridges (not the kitchen variety). This thesis would have been much less complete without your help. To Sandoko, may you remain 'still alive' for a long time to come, and teach future generations of the extreme-ness of magnonics. To my office-lab mates in 174, Nelson, Calvin, Alistair and Aaron. I hope that the office-lab remains as safe as it has been while I was there, and that you can continue to irradiate whomever takes my place.

I am also extremely grateful to my friends and family, who have listened to my complaints with good humour when things weren't going as planned. In particular, thanks to my parents and my sister for their love and support, and, of course, to Kirstie, who has done an admirable job of keeping me (relatively) sane over the last three years.

Bibliography

[1] R. D. Hicks. *Aristotle De Anima*. Cambridge University Press, 1907.

[2] Plato. *Plato in Twelve Volumes: with an English translation by W.B. Lamb*. Harvard University Press, 1925.

[3] Alexander Wylie. *Chinese Researches*. Chengwen Publishing Company, 1897.

[4] Macklis RM. Magnetic healing, quackery, and the debate about the health effects of electromagnetic fields. *Annals of Internal Medicine*, 118(5):376–383, 1993.

[5] Gerrit L. Verschuur. *Hidden Attraction: The History and Mystery of Magnetism*. Oxford University Press, apr 1996.

[6] Park Benjamin. *A history of electricity (The intellectual rise in electricity) from antiquity to the days of Benjamin Franklin*. J. Wiley, 1898.

[7] F. Bloch. Zur theorie des ferromagnetismus. *Zeitschrift für Physik*, 61(3):206–219, mar 1930.

[8] J. H. E. Griffiths. Anomalous high-frequency resistance of ferromagnetic metals. *Nature*, 158(4019):670–671, nov 1946.

[9] J. R. Eshbach. Spin-wave propagation and the magnetoelastic interaction in yttrium iron garnet. *Phys. Rev. Lett.*, 8:357–359, May 1962.

[10] P. A. Fleury, S. P. S. Porto, L. E. Cheesman, and H. J. Guggenheim. Light scattering by spin waves in FeF_2. *Phys. Rev. Lett.*, 17:84–87, Jul 1966.

[11] J.D. Adam and J.H. Collins. Microwave magnetostatic delay devices based on epitaxial yttrium iron garnet. *Proceedings of the IEEE*, 64(5):794–800, 1976.

[12] J. D. Adam, L. E. Davis, G. F. Dionne, E. F. Schloemann, and S. N. Stitzer. Ferrite devices and materials. *IEEE Transactions on Microwave Theory and Techniques*, 50(3):721–737, Mar 2002.

[13] S. A. Wolf, D. D. Awschalom, R. A. Buhrman, J. M. Daughton, S. von Molnár, M. L. Roukes, A. Y. Chtchelkanova, and D. M. Treger. Spintronics: A spin-based electronics vision for the future. *Science*, 294(5546):1488–1495, 2001.

[14] M. N. Baibich, J. M. Broto, A. Fert, F. Nguyen Van Dau, F. Petroff, P. Etienne, G. Creuzet, A. Friederich, and J. Chazelas. Giant magnetoresistance of (001)Fe/(001)Cr magnetic superlattices. *Phys. Rev. Lett.*, 61:2472–2475, Nov 1988.

[15] G. Binasch, P. Grünberg, F. Saurenbach, and W. Zinn. Enhanced magnetoresistance in layered magnetic structures with antiferromagnetic interlayer exchange. *Phys. Rev. B*, 39:4828–4830, Mar 1989.

[16] V V Kruglyak, S O Demokritov, and D Grundler. Magnonics. *Journal of Physics D: Applied Physics*, 43(26):264001, jun 2010.

[17] A A Serga, A V Chumak, and B Hillebrands. YIG magnonics. *Journal of Physics D: Applied Physics*, 43(26):264002, jun 2010.

[18] C. G. Sykes, J. D. Adam, and J. H. Collins. Magnetostatic wave propagation in a periodic structure. *Applied Physics Letters*, 29(6):388–391, 1976.

[19] A. V. Chumak, A. A. Serga, B. Hillebrands, and M. P. Kostylev. Scattering of backward spin waves in a one-dimensional magnonic crystal. *Applied Physics Letters*, 93(2):022508, jul 2008.

[20] Andrii V Chumak, Alexander A Serga, and Burkard Hillebrands. Magnon transistor for all-magnon data processing. *Nature Communications*, 5:4700, August 2014.

[21] S. Klingler, P. Pirro, T. Brächer, B. Leven, B. Hillebrands, and A. V. Chumak. Design of a spin-wave majority gate employing mode selection. *Applied Physics Letters*, 105(15):152410, oct 2014.

[22] T. Schneider, A. A. Serga, B. Leven, B. Hillebrands, R. L. Stamps, and M. P. Kostylev. Realization of spin-wave logic gates. *Applied Physics Letters*, 92(2):022505, 2008.

[23] Yutaka Tabuchi, Seiichiro Ishino, Atsushi Noguchi, Toyofumi Ishikawa, Rekishu Yamazaki, Koji Usami, and Yasunobu Nakamura. Quantum magnonics: The magnon meets the superconducting qubit. *Comptes Rendus Physique*, 17(7):729–739, aug 2016.

[24] Atac Imamoğlu. Cavity QED based on collective magnetic dipole coupling: Spin ensembles as hybrid two-level systems. *Phys. Rev. Lett.*, 102:083602, Feb 2009.

[25] J. H. Wesenberg, A. Ardavan, G. A. D. Briggs, J. J. L. Morton, R. J. Schoelkopf, D. I. Schuster, and K. Mølmer. Quantum computing with an electron spin ensemble. *Physical Review Letters*, 103(7), aug 2009.

[26] Y. Kubo, C. Grezes, A. Dewes, T. Umeda, J. Isoya, H. Sumiya, N. Morishita, H. Abe, S. Onoda, T. Ohshima, V. Jacques, A. Dréau, J.-F. Roch, I. Diniz, A. Auffeves, D. Vion, D. Esteve, and P. Bertet. Hybrid quantum circuit with a superconducting qubit coupled to a spin ensemble. *Phys. Rev. Lett.*, 107:220501, Nov 2011.

[27] Hans Huebl, Christoph W. Zollitsch, Johannes Lotze, Fredrik Hocke, Moritz Greifenstein, Achim Marx, Rudolf Gross, and Sebastian T. B. Goennenwein. High cooperativity in coupled microwave resonator ferrimagnetic insulator hybrids. *Physical Review Letters*, 111(12):127003, sep 2013.

[28] Yutaka Tabuchi, Seiichiro Ishino, Toyofumi Ishikawa, Rekishu Yamazaki, Koji Usami, and Yasunobu Nakamura. Hybridizing ferromagnetic magnons and microwave photons in the quantum limit. *Phys. Rev. Lett.*, 113:083603, August 2014.

[29] J. Bourhill, N. Kostylev, M. Goryachev, D. L. Creedon, and M. E. Tobar. Ultrahigh cooperativity interactions between magnons and resonant photons in a YIG sphere. *Physical Review B*, 93(14):144420, apr 2016.

[30] Yunshan Cao, Peng Yan, Hans Huebl, Sebastian T. B. Goennenwein, and Gerrit E. W. Bauer. Exchange magnon-polaritons in microwave cavities. *Physical Review B*, 91(9), mar 2015.

[31] Maxim Goryachev, Warrick G. Farr, Daniel L. Creedon, Yaohui Fan, Mikhail Kostylev, and Michael E. Tobar. High-cooperativity cavity QED with magnons at microwave frequencies. *Physical Review Applied*, 2(5):054002, nov 2014.

[32] Nikita Kostylev, Maxim Goryachev, and Michael E. Tobar. Superstrong coupling of a microwave cavity to yttrium iron garnet magnons. *Applied Physics Letters*, 108(6):062402, feb 2016.

[33] N. J. Lambert, J. A. Haigh, and A. J. Ferguson. Identification of spin wave modes in yttrium iron garnet strongly coupled to a co-axial cavity. *Journal of Applied Physics*, 117(5):053910, feb 2015.

[34] N. J. Lambert, J. A. Haigh, S. Langenfeld, A. C. Doherty, and A. J. Ferguson. Cavity-mediated coherent coupling of magnetic moments. *Physical Review A*, 93(2), feb 2016.

[35] Dengke Zhang, Xin-Ming Wang, Tie-Fu Li, Xiao-Qing Luo, Weidong Wu, Franco Nori, and J. Q. You. Cavity quantum electrodynamics with ferromagnetic magnons in a small yttrium-iron-garnet sphere. *Npj Quantum Information*, 1:15014, November 2015.

[36] Xufeng Zhang, Changling Zou, Liang Jiang, and Hong X. Tang. Superstrong coupling of thin film magnetostatic waves with microwave cavity. *Journal of Applied Physics*, 119(2), 2016.

[37] Xufeng Zhang, Chang-Ling Zou, Na Zhu, Florian Marquardt, Liang Jiang, and Hong X. Tang. Magnon dark modes and gradient memory. *Nature Communications*, 6:8914, nov 2015.

[38] Xufeng Zhang, Chang-Ling Zou, Liang Jiang, and Hong X. Tang. Strongly coupled magnons and cavity microwave photons. *Phys. Rev. Lett.*, 113(15), October 2014.

[39] Lihui Bai, M. Harder, Y. P. Chen, X. Fan, J. Q. Xiao, and C.-M. Hu. Spin pumping in electrodynamically coupled magnon-photon systems. *Physical Review Letters*, 114(22), jun 2015.

[40] J. A. Haigh, N. J. Lambert, A. C. Doherty, and A. J. Ferguson. Dispersive readout of ferromagnetic resonance for strongly coupled magnons and microwave photons. *Physical Review B*, 91(10), mar 2015.

[41] Michael Harder, LiHui Bai, Christophe Match, Jesko Sirker, and CanMing Hu. Study of the cavity-magnon-polariton transmission line shape. *Science China Physics, Mechanics & Astronomy*, 59(11), sep 2016.

[42] H. Maier-Flaig, M. Harder, R. Gross, H. Huebl, and S. T. B. Goennenwein. Spin pumping in strongly coupled magnon-photon systems. *Physical Review B*, 94(5), aug 2016.

[43] Yi-Pu Wang, Guo-Qiang Zhang, Dengke Zhang, Xiao-Qing Luo, Wei Xiong, Shuai-Peng Wang, Tie-Fu Li, C.-M. Hu, and J. Q. You. Magnon Kerr effect in a strongly coupled cavity-magnon system. *Phys. Rev. B*, 94:224410, Dec 2016.

[44] Yutaka Tabuchi, Seiichiro Ishino, Atsushi Noguchi, Toyofumi Ishikawa, Rekishu Yamazaki, Koji Usami, and Yasunobu Nakamura. Coherent coupling between a ferromagnetic magnon and a superconducting qubit. *Science*, 349(6246):405–408, July 2015.

[45] Dany Lachance-Quirion, Yutaka Tabuchi, Seiichiro Ishino, Atsushi Noguchi, Toyofumi Ishikawa, Rekishu Yamazaki, and Yasunobu Nakamura. Resolving quanta of collective spin excitations in a millimeter-sized ferromagnet. *Science Advances*, 3(7), 2017.

[46] H. J. Kimble. The quantum internet. *Nature*, 453:1023–1030, June 2008.

[47] R. W. Andrews, R. W. Peterson, T. P. Purdy, K. Cicak, R. W. Simmonds, C. A. Regal, and K. W. Lehnert. Bidirectional and effcient conversion between microwave and optical light. *Nature Physics*, 10:321–326, apr 2014.

[48] R. Hisatomi, A. Osada, Y. Tabuchi, T. Ishikawa, A. Noguchi, R. Yamazaki, K. Usami, and Y. Nakamura. Bidirectional conversion between microwave and light via ferromagnetic magnons. *Physical Review B*, 93(17), may 2016.

[49] S. Klingler, H. Maier-Flaig, R. Gross, C.-M. Hu, H. Huebl, S. T. B. Goennenwein, and M. Weiler. Combined brillouin light scattering and microwave absorption study of magnon-photon coupling in a split-ring resonator/YIG film system. *Applied Physics Letters*, 109(7):072402, aug 2016.

[50] A. Osada, R. Hisatomi, A. Noguchi, Y. Tabuchi, R. Yamazaki, K. Usami, M. Sadgrove, R. Yalla, M. Nomura, and Y. Nakamura. Cavity optomagnonics with spin-orbit coupled photons. *Physical Review Letters*, 116(22), jun 2016.

[51] Xufeng Zhang, Na Zhu, Chang-Ling Zou, and Hong X. Tang. Optomagnonic whispering gallery microresonators. *Phys. Rev. Lett.*, 117:123605, Sep 2016.

[52] B. I. Bleaney and B. Bleaney. *Electricity and Magnetism*. Oxford University Press, 3rd edition, 1976.

[53] Claudine Lacroix, Philippe Mendels, and Frédéric Mila, editors. *Introduction to Frustrated Magnetism*. Springer Berlin Heidelberg, 2011.

[54] James Binney and David Skinner. *The Physics of Quantum Mechanics*. Cappella Archive, 3rd edition, 2010.

[55] A. Karenowska. *Some Magnetic Reflections on Wave Dynamics*. DPhil Thesis, 2011.

[56] Daniel D. Stancil and Anil Prabhakar. *Spin Waves*. Springer Nature, 2009.

[57] B A Kalinikos and A N Slavin. Theory of dipole-exchange spin wave spectrum for ferromagnetic films with mixed exchange boundary conditions. *Journal of Physics C: Solid State Physics*, 19(35):7013–7033, dec 1986.

[58] L. R. Walker. Resonant modes of ferromagnetic spheroids. *J. Appl. Phys.*, 29(3):318, 1958.

[59] L. R. Walker. Magnetostatic modes in ferromagnetic resonance. *Phys. Rev.*, 105(2):390–399, January 1957.

[60] J. A. Osborn. Demagnetizing factors of the general ellipsoid. *Physical Review*, 67(11-12):351–357, jun 1945.

[61] K. F. Riley, M.P. Hobson, and S. J. Bence. *Mathematical Methods for Physics and Engineering*. Cambridge University Press, third edition, 2006.

[62] P. C. Fletcher and R. O. Bell. Ferrimagnetic resonance modes in spheres. *Journal of Applied Physics*, 30(5):687–698, may 1959.

[63] J. E. Mercereau. Ferromagnetic resonance g factor to order $(kR_0)^2$. *Journal of Applied Physics*, 30(4):S184–S185, 1959.

[64] Robert L. White. Use of magnetostatic modes as a research tool. *J. Appl. Phys.*, 31(5):S86, 1960.

[65] P. Röschmann and H. Dötsch. Properties of magnetostatic modes in ferrimagnetic spheroids. *Physica Status Solidi (b)*, 82(1):11–57, jul 1977.

[66] I. H. Solt and P. C. Fletcher. Magnetic anisotropy from the magnetostatic modes. *Journal of Applied Physics*, 31(5):S100–S102, 1960.

[67] V. L. Launets, M. M. Novitskas, and V. K. Shugurov. Magnetostatic oscillations of a ferrite ellipsoid with allowance for crystallographic anisotropy. *Radiophysics and Quantum Electronics*, 14(6):734–738, jun 1971.

[68] R. W. Damon and J. R. Eshbach. Magnetostatic modes of a ferromagnetic slab. *Journal of Applied Physics*, 31(5):S104–S105, may 1960.

[69] B.A. Kalinikos. Excitation of propagating spin waves in ferromagnetic films. *IEE Proceedings H - Microwaves, Optics and Antennas*, 127(1):4, feb 1980.

[70] Daniel D. Stancil, Benjamin E. Henty, Ahmet G. Cepni, and J. P. Van't Hof. Observation of an inverse Doppler shift from left-handed dipolar spin waves. *Physical Review B*, 74(6), aug 2006.

[71] Toshinobu Yukawa, Jun ichi Yamada, Kenji Abe, and Jun ichi Ikenoue. Effects of metal on the dispersion relation of magnetostatic surface waves. *Japanese Journal of Applied Physics*, 16(12):2187, 1977.

[72] W. L. Bongianni. Magnetostatic propagation in a dielectric layered structure. *Journal of Applied Physics*, 43(6):2541–2548, jun 1972.

[73] Frank Pobell. *Matter and Methods at Low Temperatures*. Springer-Verlag Berlin Heidelberg, 2007.

[74] A. T. A. M. de Waele. Basic operation of cryocoolers and related thermal machines. *Journal of Low Temperature Physics*, 164(5):179, 2011.

[75] David Meeker. Finite element method magnetics, nov 2013. Version 4.2.

[76] David M. Pozar. *Microwave Engineering*. John Wiley & Sons, Inc, second edition, 1998.

[77] C.E. Shannon. Communication in the presence of noise. *Proceedings of the IRE*, 37(1):10–21, jan 1949.

[78] A. V. Chumak, V. I. Vasyuchka, A. A. Serga, and B. Hillebrands. Magnon spintronics. *Nature Physics*, 11(6):453–461, jun 2015.

[79] H.D. Arnold and G.W. Elmen. Permalloy, an alloy of remarkable magnetic properties. *Journal of the Franklin Institute*, 195(5):621–632, may 1923.

[80] Simon Trudel, Oksana Gaier, Jaroslav Hamrle, and Burkard Hillebrands. Magnetic anisotropy, exchange and damping in cobalt-based full-Heusler compounds: an experimental review. *Journal of Physics D: Applied Physics*, 43(19):193001, apr 2010.

[81] T. Shinjo. Magnetic vortex core observation in circular dots of permalloy. *Science*, 289(5481):930–932, aug 2000.

[82] S. Urazhdin, Norman O. Birge, W. P. Pratt, and J. Bass. Current-driven magnetic excitations in permalloy-based multilayer nanopillars. *Physical Review Letters*, 91(14), oct 2003.

[83] O. Mosendz, V. Vlaminck, J. E. Pearson, F. Y. Fradin, G. E. W. Bauer, S. D. Bader, and A. Hoffmann. Detection and quantification of inverse spin Hall effect from spin pumping in permalloy/normal metal bilayers. *Physical Review B*, 82(21), dec 2010.

[84] Vladimir Cherepanov, Igor Kolokolov, and Victor L'vov. The saga of YIG: Spectra, thermodynamics, interaction and relaxation of magnons in a complex magnet. *Physics Reports*, 229(3):81–144, jul 1993.

[85] N.S. Chang, T. Hayamizu, and Y. Matsuo. YIG-tuned Gunn effect oscillator. *Proceedings of the IEEE*, 55(9):1621–1621, 1967.

[86] P. Hansen, P. Röschmann, and W. Tolksdorf. Saturation magnetization of gallium-substituted yttrium iron garnet. *Journal of Applied Physics*, 45(6):2728–2732, jun 1974.

[87] Elmer E. Anderson. Molecular field model and the magnetization of YIG. *Physical Review*, 134(6A):A1581–A1585, jun 1964.

[88] A. Wallraff, D. I. Schuster, A. Blais, L. Frunzio, R.-S. Huang, J. Majer, S. Kumar, S. M. Girvin, and R. J. Schoelkopf. Strong coupling of a single photon to a superconducting qubit using circuit quantum electrodynamics. *Nature*, 431(7005):162–167, sep 2004.

[89] M. Göppl, A. Fragner, M. Baur, R. Bianchetti, S. Filipp, J. M. Fink, P. J. Leek, G. Puebla, L. Steffen, and A. Wallraff. Coplanar waveguide resonators for circuit quantum electrodynamics. *Journal of Applied Physics*, 104(11):113904, dec 2008.

[90] Sandor Caplan and Gerald Chanin. Critical-field study of superconducting aluminum. *Physical Review*, 138(5A):A1428–A1433, may 1965.

[91] M. Tinkham. Effect of fluxoid quantization on transitions of superconducting films. *Physical Review*, 129(6):2413–2422, mar 1963.

[92] T. E. Faber and A. B. Pippard. The penetration depth and high-frequency resistance of superconducting aluminium. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 231(1186):336–353, sep 1955.

[93] D. K. Finnemore, T. F. Stromberg, and C. A. Swenson. Superconducting properties of high-purity niobium. *Physical Review*, 149(1):231–243, sep 1966.

[94] Alexandre Blais, Ren-Shou Huang, Andreas Wallraff, S. M. Girvin, and R. J. Schoelkopf. Cavity quantum electrodynamics for superconducting electrical circuits: An architecture for quantum computation. *Physical Review A*, 69(6), jun 2004.

[95] L. DiCarlo, J. M. Chow, J. M. Gambetta, Lev S. Bishop, B. R. Johnson, D. I. Schuster, J. Majer, A. Blais, L. Frunzio, S. M. Girvin, and R. J. Schoelkopf. Demonstration of two-qubit algorithms with a superconducting quantum processor. *Nature*, 460(7252):240–244, jun 2009.

[96] R. G. E. Morris, A. F. van Loo, S. Kosen, and A. D. Karenowska. Strong coupling of magnons in a YIG sphere to photons in a planar superconducting resonator in the quantum limit. *Scientific Reports*, 7:11511, sep 2017.

[97] http://www.ferrisphere.com/. Accessed 28-04-2017.

[98] Lauren Peters. CMP20: Rotatable superconducting coplanar waveguide resonator in magnetic fields. Master's thesis, Oxford University, 2015.

[99] S. Ulmer, H. Kracke, K. Blaum, S. Kreim, A. Mooser, W. Quint, C. C. Rodegheri, and J. Walz. The quality factor of a superconducting rf resonator in a magnetic field. *Review of Scientific Instruments*, 80(12):123302, 2009.

[100] ANSYS. Electromagnetics suite, 2014. release 16.0.0, build 2014-11-21 00:25:49.

[101] http://www.ansys.com/. Accessed 28-04-2017.

[102] E.T. Jaynes and F.W. Cummings. Comparison of quantum and semiclassical radiation theories with application to the beam maser. *Proceedings of the IEEE*, 51(1):89–109, 1963.

[103] T. Holstein and H. Primakoff. Field dependence of the intrinsic domain magnetization of a ferromagnet. *Physical Review*, 58(12):1098–1113, dec 1940.

[104] E. M. Purcell, H. C. Torrey, and R. V. Pound. Resonance absorption by nuclear magnetic moments in a solid. *Phys. Rev.*, 69:37–38, Jan 1946.

[105] P. Hansen. Anisotropy and magnetostriction of gallium-substituted yttrium iron garnet. *Journal of Applied Physics*, 45(8):3638–3642, aug 1974.

[106] A.G. Gurevich and G.A. Melkov. *Magnetization Oscillations and Waves*. CRC Press, Inc, 1996.

[107] Irvin H. Solt. Temperature dependence of YIG magnetization. *Journal of Applied Physics*, 33(3):1189–1191, mar 1962.

[108] William R. Hamilton, Mallet, Mac Cullagh, and H. Lloyd. First series of researches respecting vibration. *Proceedings of the Royal Irish Academy (1836-1869)*, 1:341–349, 1839.

[109] M. Maryško. Ferromagnetic resonance relations in {111}-oriented garnet films. *Czechoslovak Journal of Physics*, 30(11):1269–1278, nov 1980.

[110] B A Kalinikos, M P Kostylev, N V Kozhus, and A N Slavin. The dipole-exchange spin wave spectrum for anisotropic ferromagnetic films with mixed exchange boundary conditions. *Journal of Physics: Condensed Matter*, 2(49):9861–9877, dec 1990.

[111] M.J. Hurben and C.E. Patton. Theory of magnetostatic waves for in-plane magnetized anisotropic films. *Journal of Magnetism and Magnetic Materials*, 163(1-2):39–69, oct 1996.

[112] Stephen Butterworth. On the theory of filter amplifiers. *Experimental Wireless*, 7:536–541, October 1930.

[113] A. V. Chumak, A. A. Serga, S. Wolff, B. Hillebrands, and M. P. Kostylev. Scattering of surface and volume spin waves in a magnonic crystal. *Applied Physics Letters*, 94(17):172511, apr 2009.

[114] S. N. Bajpai and N. C. Srivastava. Magnetostatic bulk waves in an arbitrarily magnetized YIG–dielectric layered structure. *Physica Status Solidi (a)*, 57(1):307–315, jan 1980.

[115] T. W. O'Keeffe and R. W. Patterson. Magnetostatic surface-wave propagation in finite samples. *Journal of Applied Physics*, 49(9):4886–4895, sep 1978.

[116] A. F. van Loo, R. G. E. Morris, and A. D. Karenowska. Time-resolved measurements of surface spin-wave pulses at millikelvin temperatures. *arXiv:1610.08402*, 2016.

[117] W. I. Kinney and W. P. Wolf. Magnetic interactions and short range order in gadolinium gallium garnet. *Journal of Applied Physics*, 50(B3):2115–2117, 1979.

[118] S. Hov, H. Bratsberg, and A.T. Skjeltorp. Magnetic phase diagram of gadolinium gallium garnet. *Journal of Magnetism and Magnetic Materials*, 15-18:455–456, jan 1980.

[119] P. P. Deen, O. Florea, E. Lhotel, and H. Jacobsen. Updating the phase diagram of the archetypal frustrated magnet $Gd_3Ga_5O_{12}$. *Physical Review B*, 91(1), jan 2015.

[120] R. I. Joseph and E. Schlömann. Demagnetizing field in nonellipsoidal bodies. *Journal of Applied Physics*, 36(5):1579–1593, may 1965.

[121] P. Schiffer, A. P. Ramirez, D. A. Huse, and A. J. Valentino. Investigation of the field induced antiferromagnetic phase transition in the frustrated magnet: Gadolinium gallium garnet. *Physical Review Letters*, 73(18):2500–2503, oct 1994.

[122] P. Schiffer, A. P. Ramirez, D. A. Huse, P. L. Gammel, U. Yaron, D. J. Bishop, and A. J. Valentino. Frustration induced spin freezing in a site-ordered magnet: Gadolinium gallium garnet. *Physical Review Letters*, 74(12):2379–2382, mar 1995.

[123] O.A Petrenko, D.McK Paul, C Ritter, T Zeiske, and M Yethiraj. Magnetic frustration and order in gadolinium gallium garnet. *Physica B: Condensed Matter*, 266(1-2):41–48, may 1999.

[124] J. A. M. Paddison, H. Jacobsen, O. A. Petrenko, M. T. Fernandez-Diaz, P. P. Deen, and A. L. Goodwin. Hidden order in spin-liquid $Gd_3Ga_5O_{12}$. *Science*, 350(6257):179–181, oct 2015.

[125] V. V. Danilov, D. L. Lyfar', Yu. V. Lyubon'ko, A. Yu. Nechiporuk, and S. M. Ryabchenko. Low-temperature ferromagnetic resonance in epitaxial garnet films on paramagnetic substrates. *Soviet Physics Journal*, 32(4):276–280, apr 1989.

[126] M. Maryško. Influence of GGG substrate on FMR and magnetostatic wave propagation. *Journal of Magnetism and Magnetic Materials*, 101(1-3):159–161, oct 1991.

[127] A V Chumak, T Neumann, A A Serga, B Hillebrands, and M P Kostylev. A current-controlled, dynamic magnonic crystal. *Journal of Physics D: Applied Physics*, 42(20):205005, sep 2009.