Design And construction Of Atomic Comagnetometers

MSC BY RESEARCH THESIS

By

Caleb Arthurs

Supervisor:

Dr Vera Guarrera

Cold Atoms Research Group
School of Physics and Astronomy
College of Engineering and Physical Sciences
University of Birmingham
September 2023

ABSTRACT

This thesis presents an in depth look into the design and construction of a series of new co-magnetometry experiments. Work to design a high-performance magnetic shield that minimises magnetic field noise whilst maintaining a larger internal volume is detailed. A high-performance magnetic coil assembly that optimally fits within a cylindrical environment is designed and implemented.

Preliminary work to characterise one of the co-magnetometry systems utilising novel high pressure ^{87}Rb - ^{21}Ne atomic vapour cells is conducted. This presented possible multiple Larmor precession resonances within the vapour cell, representing new physics. The creation of separated regions within the vapour cell due to high diffusion times is hypothesised and offered for further investigation.

CONTENTS

I. Introduction 6

II. Theoretical Background 8

A. Preliminary Theory 8

1. Zeeman Interaction 8

2. Atomic Complex Refractive Index 9

3. Optical Rotation 10

4. Atomic Transition Polarisations 11

B. Bell-Bloom Magnetometry 13

1. Setup 13

2. Creating circular birefringence 14

3. Quantum description of adding a B-field 14

4. Pump modulation 16

C. Co-magnetometry 17

1. Interaction equations 17

2. Measurement 18

III. First Experiment Magnetic Shielding 20

A. Background Theory 20

B. Custom Shield Design 20

C. Ferrite Design 21

D. MuMetal Design 23

1. Layer Spacing 23

2. Thickness 23

E. Design Verification 25

F. Manufacture 27

IV. Magnetic Coils Design 28

A. Saddle Coils 28

1. Ansys-Python Coil Analysis 29

B. Final Design - First Experiment 30

1. Mounting inside the shield 32

V. First Experiment Optical Setup 35
A. Laser frequency stabilisation 35
B. Polarisation 36

VI. Offset Lock Circuit 39
1. Aside: analogue processing 39
A. Digital processing 39
1. Frequency Division 39
2. Frequency Detection 41
3. PWM conversion 42
4. Signal Processing 43
5. ±4V Clipping 43
B. Design Realisation 44
1. Probe detuning 44

VII. Secondary Experiment Design Evolution 46
A. Magnetic Shield 46
B. Coil Assembly 47
C. Optical Setup 49
1. Probe path changes 51
2. Pump path changes 51

VIII. Vapour Cell 52
A. Secondary Experiment Oven Design 53
1. Final Oven Design 54

IX. Instrumentation Control 56
A. Temperature Control 56
1. Heater circuit 57
2. Temperature sensing circuit 58
B. Coil Control 60

X. Secondary Experiment Characterisation 62
A. Magnetic Field 62
1. Aside: Degaussing 62
2. Setup 62

3. Current - magnetic field relations 63
4. Field Shape 65
B. Magnetometer signal 68
1. Aside: Lock-In Amplifier Theory 68
2. Pump Power - FWHM Relation 73

XI. Cell Regions At High Density 75
A. Hypothesis 75
B. Data analysis 76
1. Discussion 78

XII. Overview and Outlook 81

References 84

APPENDICES 87

A. Polarisation Notation 87

B. Magnetic Shield Drawing 87

C. Coil Type Field Comparison 88
1. Helmholtz 88
2. Saddle 89
3. Maxwell 90
4. Maxwell with holes 91
5. Final Design 92

D. Probing a different state 92

E. Offset Lock Circuit Diagram 93

I. INTRODUCTION

The basic premise of atomic magnetometry is measuring atoms to be able to detect the strength of magnetic fields.

The field of atomic magnetometry developed from Faraday's discovery of the optical rotation of light polarisation as it passes through an optical medium under a magnetic field [1]. As explained further in the thoery section, II, this occurs as the magnetic field causes a splitting of the atomic energy levels within the medium (Zeeman splitting[2]) which creates a circular birefringence. This causes a relative phase shift between the counter-rotating circular polarisations which linearly polarised light decomposes into. Because of this the linear polarisation rotates.

Using Faraday rotation alone and a polarising filter this could be used to detect magnetic field strength based on the optical power transmitted through the filter. There are two drawbacks to this approach. First the Faraday rotation is related to length of material the light passes through via:

$$d\Theta = \mathcal{V}Bd, \tag{1}$$

where $d\Theta$ is the rotation angle, $\mathcal{V}$ is the Verdet constant (material property), B the magnetic field strength, and d distance in the material[3]. Hence, to achieve high enough angle changes for detection of small magnetic fields very long setups of optical medium would be required. Second the resolution of magnetic field strength detectable is limited by both the precision and range of optical intensity that can be measured.

Atomic magnetometers overcome this by shifting the measurement from an intensity change to a frequency change. This was made possible by the advent of optical pumping which led to the development of the Bell-Bloom magnetometer[4]. The Bell-Bloom magnetometer uses laser beams perpendicular to the magnetic field. The split energy levels caused by the Zeeman effect are seen as an oscillation (transition) between those atomic levels when the frame of reference is shifted to an axis perpendicular to the magnetic field (more details in section II B). Hence, the angle of polarisation rotation (of the probe beam) oscillates, which can then be detected very precisely utilising a lock-in amplifier. Optical pumping is needed to shift atoms into one of the two oscillating states such that a circular birefringence occurs.

Due to limited technology (poor lasers etc.) atomic magnetometry was limited in precision as the detected resonances had large widths. This led to drop off in their use. During this time the super-conducting quantum interference device (SQUID) was realised (1987)[5]. This achieved extremely high levels of precision in magnetic field detection and is a staple in ultra-precise

magnetic field detection.

Recent years have seen a revival of interest in the field as the fidelity of these devices improved drastically due to advances such as better lasers, and anti-relaxation coatings[6] which work to reduce the width of detected resonances. The main advantage of atomic magnetometry over the competing SQUID, is the ability to work at room temperature rather than requiring cryogenics [6]. Not only does this make experiments easier to setup and more accessible, it also allows for massive gains[1] in applications limited by the use of cryogenic temperatures such as measuring biotic fields (e.g. human brain magnetic fields). This application is further enhanced by using atomic magnetometers as they recover faster from large fields (often used for stimulation[7]) than SQUIDs.

Co-magnetometry, where 2 species of atoms are used during measurement, allow for differential measurements which improves precision further than a single species magnetometer via common-mode rejection[8]. By introducing a species that will not interact with the probing laser, a noble gas with net nuclear spin, a differential measurement is implicitly made. This is because the noble gas species will too experience the spin precession from magnetic field coupling (see section II A 1) and hence generate its own (opposing) magnetic field from the rotating spin. This field affects total magnetic field experienced by the measured species and therefore the field actually measured is a differential between the two species.

The field generated by the second species will track the external magnetic field to a small degree, allowing for less sensitivity to transverse fields than a single species magnetometer[9]. This insensitivity to transverse fields, gradients and field drift make the co-magnetometer a very sensitive platform for the search of new physics involving spin coupling[9][10]. By introducing a nuclear species the precision of co-magnetometers can be leveraged to search for spin-couplings to nuclear moments, such as the quantum chromodynamics (QCD) axion[11].

This paper reports on the design and construction of two co-magnetometer experiments, underlying theory, and initial characterisation of one of the experiments.

[1]The gain arises as magnetometers can be placed close to the target where as SQUIDs must be kept distant due to the cryogenic temperature they operate at.

II. THEORETICAL BACKGROUND

This section will walk through preliminary theory building up to the theory of co-magnetometry operation.

A. Preliminary Theory

1. Zeeman Interaction

For atoms under a magnetic (B) field the m sublevels of the atoms' hyperfine structure experience an energy shift. The interaction energy is described as follows[2],

$$\mathcal{H}_B = -\mu \cdot \mathbf{B} = -\frac{1}{\hbar}(\mu_N g_I \mathbf{I} - \mu_B g_J \mathbf{J}) \cdot \mathbf{B} \approx \frac{1}{\hbar}\mu_B g_J \mathbf{J} \cdot \mathbf{B}, \tag{2}$$

where $\mathbf{I}$ is the nuclear total angular momentum and $\mathbf{J}$ is the electron total angular momentum; μ_N, μ_B are the nuclear and electronic (Bohr magneton) magnetic moments respectively; μ is the total atomic magnetic moment; and g is the respective Landé factor. The magnetic moment of the atom, μ, opposes the applied magnetic field, $\mathbf{B}$, hence the minus sign convention.

This makes the total Hamiltonian (relevant at this energy scale)[2]:

$$\mathcal{H} = \mathcal{H}_{hf} + \mathcal{H}_B \approx \mathcal{A}\mathbf{I} \cdot \mathbf{J} + \frac{1}{\hbar}\mu_B g_J \mathbf{J} \cdot \mathbf{B}, \tag{3}$$

where, $\mathcal{A}$, is the hyperfine structure constant.

In the low B-field regime (that of sensitive magnetometry) the energy of the hyperfine interaction (from $\mathcal{H}_{hf}$) is larger than that of the magnetic interaction energy. Hence $\mathbf{F} = \mathbf{I} + \mathbf{J}$ is a good basis to work in since the coupling between them is strong[2]. Hence,

$$\mathcal{H} \approx \mathcal{A}\mathbf{I} \cdot \mathbf{J} + \frac{1}{\hbar}\mu_B g_F \mathbf{F} \cdot \mathbf{B}. \tag{4}$$

By aligning the B-field with the z-axis the dot product is simple to deal with and hence the energy shift becomes[2]:

$$\Delta E = \mu_B g_F B \Delta m_{F_z}, \tag{5}$$

noting that the m_{F_z} level is the one parallel to the B-field. Orthoginal levels do not experience an energy shift.

[2] In the high field regime J becomes the basis of choice hence the energy shift uses m_J instead.

Aside: Quantisation Axis

The quantisation axis is the axis you are viewing a quantum system from and does not change the physics of a system. It is equivalent to projecting a state onto a given (quantisation) axis. Since m levels are the projection of an operator onto a given axis already, they may be given without formally stating a quantisation axis. Any implicit quantisation axis lies along one of the main axes defined in a given coordinate system or along a physical feature (e.g. laser propagation direction), and is often noted in the subscript.

Since these m_{F_z} states now have different energy levels a transition can occur between them (following selection rules). The energy of the transition is therefore determined by Δm_{F_z}. A frequency can be assigned to this transition via $E = \hbar\omega$, called the Larmor frequency:

$$\omega_L = \frac{\Delta E(\Delta m_{F_z})}{\hbar} = \frac{\mu_B g_F}{\hbar} B \Delta m_{F_z} \equiv \gamma B \Delta m_{F_z}, \tag{6}$$

where γ is the gyromagnetic ratio.

2. Atomic Complex Refractive Index

Here the atomic refractive index will be shown to illustrate the absorptive drop-off from detuning which is utilised in atomic magnetometers to prevent power broadening from the laser beams (see section II B for magnetometer theory).

For an atomic vapour the refractive index, experienced by a laser beam passing through, is determined by nearby transition lines. For a transition with a frequency ω_0, the complex refractive index, $\tilde{n}$, is as follows [12]:

$$\tilde{n}^2 = \frac{\epsilon}{\epsilon_0} = 1 + \frac{4\pi N f e^2}{m(\omega_0^2 - \omega^2 - i\omega\Gamma)}, \tag{7}$$

where ω is the laser frequency, ω_0 is the transition frequency, N is the atomic state density, f is a 'fudge factor' [12], m is the mass, and Γ is the effective damping (from derivation[3]).

This can be split into a real and imaginary part:

$$\tilde{n} = n_0 + i n_0 \kappa, \tag{8}$$

[3]The derivation is classical and based off an electric driving force of a dipole [12].

where the real part is the normal refractive index, n_0, and the imaginary part, $n_0\kappa$, relates to absorption, where κ is a real coefficient.

$$\begin{aligned} n_0 = \mathrm{Re}(\tilde{n}) &\approx 1 - \frac{2\pi N f e^2(\omega^2 - \omega_0^2)}{m\big[(\omega^2-\omega_0^2)^2 + (\Gamma\omega)^2\big]} \\ &\approx 1 - \frac{\pi \Delta N f e^2}{m\omega_0(\Delta^2 + \Gamma^2/4)}, \end{aligned} \tag{9}$$

$$\begin{aligned} n_0\kappa = \mathrm{Im}(\tilde{n}) &\approx \frac{2\pi N f e^2 \omega}{m\big[(\omega^2-\omega_0^2)^2 + (\Gamma\omega)^2\big]} \\ &\approx \frac{\pi N f e^2 \Gamma}{2m\omega_0(\Delta^2 + \Gamma^2/4)}, \end{aligned} \tag{10}$$

where Δ is the detuning of the laser from resonance (ω_0).

Importantly the refractive index, n_0, is proportional to N and inversely proportional to Δ, for large Δ:

$$\begin{aligned} \delta n_0 &\propto N \\ &\propto \sim \frac{1}{\Delta}. \end{aligned} \tag{11}$$

This means that by varying either the detuning of the (nearest) transition, or the atomic density of the state in transition, the refractive index of the atomic vapour can be varied. An important distinction between the absorptive, $in_0\kappa$, and refractive, n_0, components is that $n_0 \propto \Delta^{-1}$ whereas the absorptive component is proportional to Δ^{-2} and hence drops off faster from detuning:

$$n_0\kappa \propto \sim \frac{1}{\Delta^2}. \tag{12}$$

3. Optical Rotation

Linearly polarised light can be broken down into a superposition of circular polarisation components (please see Appendix A for the notation):

$$\pi_x = \frac{1}{\sqrt{2}}\Big(\sigma_z^+ + \sigma_z^-\Big). \tag{13}$$

If a medium has a different refractive index for the two circular polarisations, n_0^+ and n_0^-, (circular birefringence) then for linear light a relative phase is picked up between the two circular components. This results in the orientation of linear polarisation rotating.

This rotation can be represented easily on a Poincaré sphere (figure 1); further the relation between linear and circular polarisations can be seen using the sphere (similar to Bloch sphere

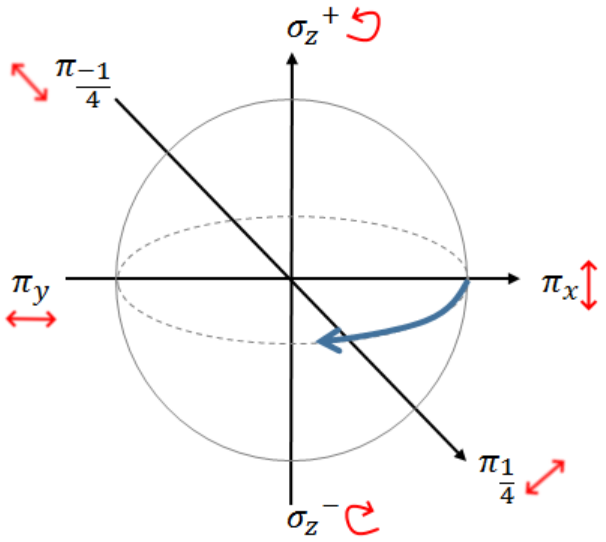

Figure 1: Poincaré sphere showing rotation (blue arrow) of π_x light towards π_y through the linear plane (π_x-$\pi/4$ plane; along the dotted equator) as a result of relative phase pick up between σ^+, σ^-. The rotation is through this plane as a phase is picked up between the circular polarisations but the relative amplitudes remain constant hence the rotation will be equidistant from each σ. The red arrows show a graphical respresentation of the polarisations.

representation). The change in polarisation angle, θ, is therefore proportional to the difference in refractive index for the two circular polarisations. Specifically

$$\Delta\theta = \frac{\pi l(n_0^+ - n_0^-)}{\lambda_0}, \tag{14}$$

where l is the path length, λ_0 is the wavelength in vacuum[13].

4. Atomic Transition Polarisations

A simple atomic system is considered here to allow discussion of the impact of laser polarisation of transition excitation. Figure 2 shows all possible optical transitions ($\Delta m = \{0, \pm 1\}$). The system is quantised along the z axis. It is assumed transitions between the upper and lower states obey electric dipole selection rules, but the states are unlabeled to keep a generic description[4].

For a quantisation axis along z, the polarisations that induce transitions between states of the canonical basis are $\{\sigma_z^+, \sigma_z^-, \pi_z\}$ (see appendix A for notation reference) and hence any other polarisations must be rotated into this basis (equivalent to splitting them into components in this basis). The logic for this[14] is that the light carries angular momentum along its direction

[4]A simple structure is taken where the angular momentum operator (e.g. **F**, **J**...) is 0 for the ground state and 1 for the excited state, since this gives the minimum m levels to describe all the polarisations.

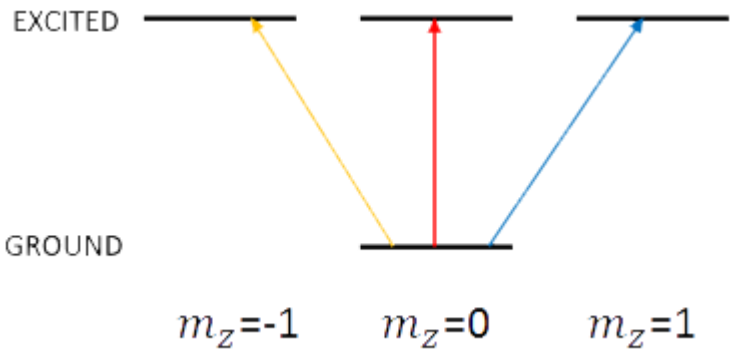

Figure 2: Simple energy level diagram showing transitions for all optical polarisations. The generic diagram is drawn assuming transitions from the lower to upper levels obey selection rules. Here the quantisation axis is parallel to the z axis.

of propagation equal to$\{\hbar_-, -\hbar_+, 0\}$[5], respectively. Hence when viewing the system along zz (quantisation axis along z), only angular momentum carried along the z axis is relevant.

By using equation 13 or the Poincaré sphere (figure 1) all the transitions allowed for each optical polarisation for a given quantisation axis can be found. These are detailed in the table I.

Polarisation	**Transition** Δm_z
σ_z^+	↗
σ_z^-	↖
π_z	↑
π_x , π_y	$\frac{1}{\sqrt{2}}$(↗ + ↖)
$\sigma_{x,y}^+$, $\sigma_{x,y}^-$	$\frac{1}{\sqrt{2}}\left[\uparrow + \frac{1}{\sqrt{2}}(\nearrow + \nwarrow)\right]$

Table I: Table showing all light polarisation (and propagation) options and the allowed Δm_z transitions, quantised along the z axis. Arrows are pointed and coloured to match figure 2.

The breakdown for circular light propagating orthogonal to the quantisation axis (z) is given below:

$$\sigma_{x,y}^{\pm} = \frac{1}{\sqrt{2}}(\pi_z \pm \pi_{y,x}) = \frac{1}{\sqrt{2}}\left[\pi_z \pm \frac{1}{\sqrt{2}}\left(\sigma_z^+ + \sigma_z^-\right)\right], \tag{15}$$

noting the subscripts: $\sigma_{x,y}$ are switched from $\pi_{y,x}$ to indicate which subscript refers to which as the above refers to multiple cases of the equation combined, and not to indicate both simultaneously, i.e. the equation for σ_x has a π_y term in the middle of the equation.

[5] π_z light gives 0 angular momentum along z axis as it cannot decompose into $\sigma_z^{\pm}$.

B. Bell-Bloom Magnetometry

1. Setup

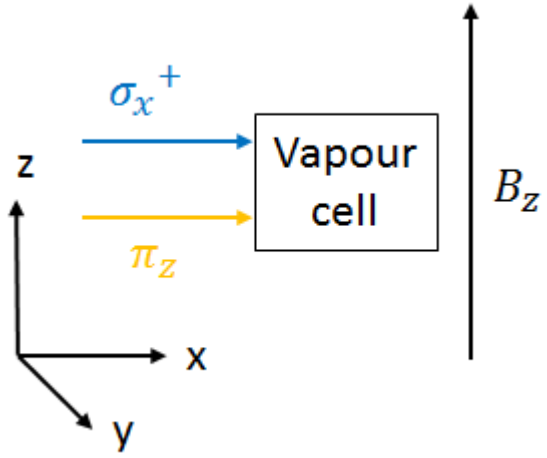

Figure 3: Diagram to explain a Bell-Bloom type magnetometer. The blue arrow represents the pump laser which has σ^+ polarisation and propagates along x. The orange arrow shows the probe laser and also propagates along x but has linear polarisation (whether the polarisation axis is along z or y is unimportant). The B-field is perpendicular to the lasers, shown to align with the z axis here.

The basic setup of a Bell-Bloom type magnetometer is shown in figure 4. In this setup the polarisation of the probe beam needs to be π_z. This is because the polarisation axis cannot be parallel to the pump beam propagation[6]. Often in practical setups the probe beam is perpendicular to the field and pump laser (parallel to y in figure 3).

For explanation purposes a simple atomic hyperfine system is used, as described in figure 4.

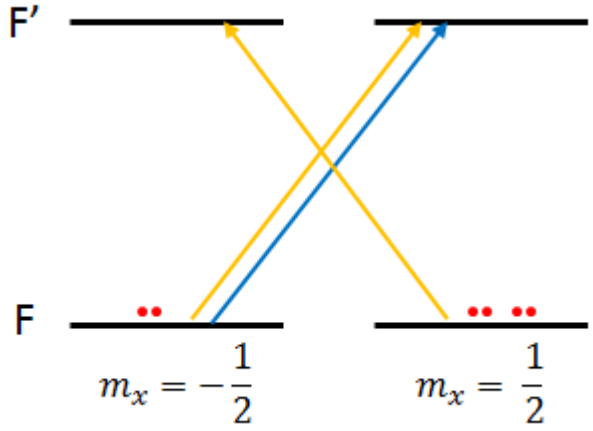

Figure 4: A simple atomic energy level setup to explain the principles of a Bell-Bloom type. The colour coding of the transitions matches the laser colour coding in figure 3. The F levels of the system are unimportant. The quantisation axis is defined parallel to the pump beam (x-axis). The red dots are to illustrate atomic densities in each state due to optical pumping.

[6]The reason for this is seen later. The probe beam needs to decompose its polarisation into σ_x^+ and σ_x^- such that it can excite both the $\Delta m_x = \pm 1$ transitions. If propagating along the pump axis both linear polarisations achieve this. If propagating perpendicular to it then the linear polarisation axis cannot be parallel to the pump axis as this will drive $\Delta m_x = 0$ transitions instead.

2. Creating circular birefringence

In order to create an oscillating signal in the probe beam, which is easier to detect than a static rotation, a circular birefringence is created in the atomic system by the pump laser beam.

The pump beam optically pumps the atoms into the high m_x state due to its polarisation, as shown in figure 4. This occurs when the excited state lifetime is significantly shorter than $\frac{1}{\Omega}$, where Ω is the Rabi frequency of the transition such that spontaneous emission occurs more often than stimulated emission. Since spontaneous emission has a random decay path and the pump excitation path is increasing m_x the net path for atoms is increasing m_x.

This creates a circular birefringence in the vapour since n_0 (see equation 9) is proportional to the atomic state density, N, for the relevant transition. The probe beam can be decomposed into $\sigma_x^{\pm}$ components hence it experiences optical rotation governed by equation 14. It is this optical rotation that is detected experimentally.

3. Quantum description of adding a B-field

Since spontaneous emission occurs quickly in this system (see section II B 2) the atoms tend to populate the ground states (see figure 4), hence the m_F states here implicitly refer to the ground F states. Further, the states will be labelled as in keeping with notation in quantum computing. Hence, $m_z = \left\{-\frac{1}{2}, \frac{1}{2}\right\}$ are labelled $\{|0\rangle, |1\rangle\}$ and $m_x = \left\{-\frac{1}{2}, \frac{1}{2}\right\}$ are labelled $\{|-\rangle, |+\rangle\}$ respectively.

Until now the effects from the lasers in the system described above (section II B 1) have been treated in a non-magnetic field environment. Adding the B-field into the picture has the effect of splitting the energy levels along its axis (z axis for figure 3) via the Zeeman interaction (see section 4).

The physics of interest are easiest to see considering the axis along the pump laser; hence, the rotations between the x-basis and z-basis are given below:

$$|+\rangle = \frac{1}{\sqrt{2}}\big(|0\rangle + |1\rangle\big), \tag{16}$$

$$|-\rangle = \frac{1}{\sqrt{2}}\big(|0\rangle - |1\rangle\big). \tag{17}$$

Just before the magnetic field is turned on, the ground m_F states in the z axis have the same energy, so the distribution of atoms will be equal. Hence after the field is turned on the time

dependant wavefunction for the system is

$$|\psi, t\rangle = \frac{1}{\sqrt{2}} e^{\frac{iE_0 t}{\hbar}} \Big(|0\rangle + e^{i\omega_L t} |1\rangle \Big), \tag{18}$$

where it is quantised/viewed along the z axis since this axis contains the shifted energy levels. A global phase/energy has been taken out where E_0 is the energy of the $|0\rangle$ state. Subsequently a relative phase is picked up on the $|1\rangle$ state, where $\hbar\omega_L$ is the energy difference between the m_{F_z} states.

The time evolution of the population of the $|+\rangle$ state can then be calculated:

$$\begin{aligned} \mathcal{P}\big(|+,t\rangle\big) = \big|\langle +|\psi,t\rangle\big|^2 &= \left| \frac{1}{2} \exp\left(\frac{iE_0 t}{\hbar}\right) \big(\langle 0| + \langle 1|\big) \Big(|0\rangle + e^{i\omega_L t} |1\rangle \Big) \right|^2 \\ &= \frac{1}{4} \Big| 1 + e^{i\omega_L t} \Big|^2 \\ &= \frac{1}{4} \Big(2 + e^{i\omega_L t} + e^{-i\omega_L t} \Big) \equiv \frac{1}{2} \left(1 + \frac{e^{i\omega_L t} + e^{-i\omega_L t}}{2} \right) \\ &= \frac{1}{2} \big(1 + \cos(\omega_L t) \big). \end{aligned} \tag{19}$$

Since this is a simple 2 level system[7] it is clear that the system oscillates between $|+\rangle$ and $|-\rangle$ at the Larmor frequency, ω_L, as shown in figure 5.

Since these states have been optically pumped, such that the refractive index of the two σ_x transitions are different (discussed in section II B 2), when the m_x state populations rotate the optical rotation effect caused by the difference in the refractive indices (see equation 14) also rotates.

The optical rotation oscillates between $\pm\Delta\theta$, where $\Delta\theta$ is defined as the optical rotation under no B-field (as defined in equation 14), at a frequency of ω_L. Hence the B-field strength can be

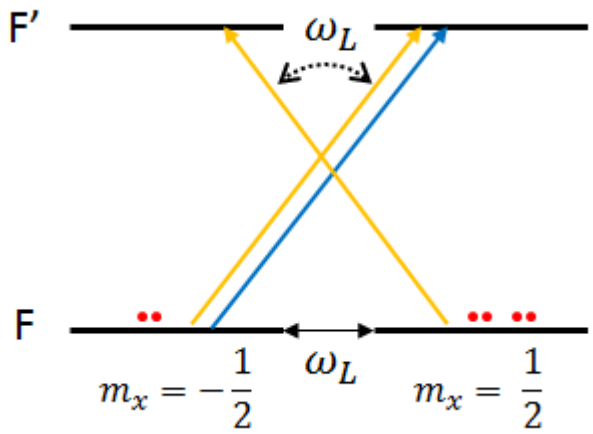

Figure 5: Modified version of figure 4 showing the effect of the B-field. The black arrow represents the effect from the lab frame (quantised along x), where m_x states oscillate/undergo a transition with a frequency of ω_L [8].
The dashed arrow shows the effect from the atoms frame (quantised along μ, the atomic magnetic moment) whereby light polarisation oscillates between σ_x^+ and σ_x^- at the same rate, ω_L, and the states do not.

[7] For more m_F states the oscillations occur between 'magnetic multiplets', which are neighbouring m_F states (separated by ω_L) here.

[8] The upper, F', states also undergo this but it is not an important consideration as they are not significantly populated.

determined experimentally from this oscillation frequency, as below (derived in section II A 1, equation 6; for low fields):

$$B = \frac{\hbar}{\mu_B g_F}\omega_L. \tag{20}$$

4. Pump modulation

The combination of optical pumping (section II B 2) and oscillation from the B-field (section II B 3) causes a complication. As stated the pump beam needs to create a (strong) population difference between the two m_x levels in order for optical rotation to occur (see equation 14). This is achieved by optical pumping.

However, as the B-field causes the m_x states to rotate (when viewed from the, non-rotating, lab frame) the pump will start to excite atoms from the high population state hence the populations will average out between the two states and no optical rotation will occur.

This can be seen easiest by considering the frame of reference rotating with the atomic states (as noted in the caption of figure 5). In this frame the m_{F_x} states do not oscillate, however, the pump beam rotates about the atom instead (at ω_L). As this happens the polarisation of light experienced by the atom changes. The polarisation follows a(n elliptical[9]) path around the Poincaré sphere (figure 1.) from σ_x^+ to σ_x^- and back (seen since $\sigma_{-x}^+ = \sigma_x^-$).

To avoid this problem the pump beam is (amplitude) modulated at ω_L such that its intensity is zero when half a rotation has occurred and therefore doesn't pump population out of the high population state.

[9]Elliptical since at quarter rotation light is σ_y^+ which decomposes as in equation 15

C. Co-magnetometry

The physics of a co-magnetometry experiment derives from a Bell-Bloom type magnetometer. Co-magnetometry adds an additional atomic species (or multiple) that contributes additional, similar, terms to the dynamical equations. These species are noble gases, such that the probe laser only interacts with the alkali atoms (as noble gases have full shells).

The main impact of adding the noble gas is that it also rotates under the B-field but at a different Larmor frequency (see section II A 1 noting J=0). The atoms pick up a net spin, $\langle \mathbf{I} \rangle$, via spin exchange collisions[10] (SECs) with the polarised alkali. As in the Bell-Bloom type they rotate under the B-field (see section II B 3). This rotating spin creates an effective magnetic field that can then interact with the alkali species.

1. *Interaction equations*

The equations for co-magnetometers are derived from Bloch equations[11][9]. The Bloch equations take a macroscopic, classical approach to describe nuclear magnetic resonance (NMR) signals. They consider the system where the (net) spin acts like a classical angular momentum under the effect of a torque which represents the magnetic field[15]:

$$\frac{\mathrm{d}\mathbf{M}}{\mathrm{d}t} = \gamma \mathbf{M} \times \mathbf{B}, \tag{21}$$

where $\mathbf{M}$ is the macroscopic nuclear magnetisation, γ the gyromagnetic ratio, and relaxation terms have been omitted.

To adjust the Bloch equations to the atomic picture the magnetic moment must be converted to spin. This is done by combining equations 2 and 4 to show that

$$\mathbf{m} = \boldsymbol{\mu} = -\frac{\mu_B g_F}{\hbar}\mathbf{F}, \tag{22}$$

where $\mathbf{M}$ in the Bloch equation (21) is the sum over all $\mathbf{m}$ here.

For the alkali atom this leads to[9][16]:

$$\frac{\mathrm{d}\hat{\mathbf{F}}}{\mathrm{d}t} = \frac{1}{q}\left[\gamma_e \underbrace{\left(\mathbf{B}_{\text{ext}} + \lambda M_N \hat{\mathbf{I}} + \mathbf{L} + \mathbf{b_e}\right)}_{\text{effective } \mathbf{B} \text{ terms}} \times \hat{\mathbf{F}} + \overbrace{\mathcal{R}_e}^{\text{relaxation terms}}\right] + \underbrace{\boldsymbol{\Omega} \times \hat{\mathbf{F}}}_{\text{rotating body correction}}, \tag{23}$$

[10] Spin-spin interactions between atoms whereby the total magnetic moment of both atoms is conserved but the moment may be redistributed among $\mathbf{m}_F$ levels

[11] Not to be confused with the quantum state Bloch sphere.

where the cross product is flipped from the Bloch equation (equation 21) to remove the minus sign picked up from equation 22 using the anticommutivity of the vector product. $\hat{\mathbf{F}}$ is the macroscopic unit total atomic spin vector (equivalent to atomic polarisation), explicitly:

$$\hat{\mathbf{F}} = \frac{\langle \mathbf{F} \rangle}{\|\mathbf{F}\|}; \tag{24}$$

similarly for $\hat{\mathbf{I}}$ which refers to the noble gases total nuclear angular momentum, and q is the paramagnetic 'slowing down' coefficient[12]. It is defined as

$$q = \frac{\langle F_z \rangle}{\langle J_z \rangle}. \tag{25}$$

The remaining terms are outlined in table II.

Similarly for the noble gases:

$$\frac{\mathrm{d}\hat{\mathbf{I}}}{\mathrm{d}t} = \gamma_N \underbrace{\left(\mathbf{B}_{\text{ext}} + \lambda M_e \hat{\mathbf{F}} + \mathbf{b_N}\right)}_{\text{effective } \mathbf{B} \text{ terms}} \times \hat{\mathbf{I}} + \overbrace{\mathcal{R}_N}^{\text{relaxation terms}} + \underbrace{\boldsymbol{\Omega} \times \hat{\mathbf{I}}}_{\text{rotating body correction}}.[16] \tag{26}$$

Noting that there is no a.c. stark term, $\mathbf{L}$, since the lasers only interact with the alkali atoms.

A possible anomalous field coupling, $\mathbf{b}$ (explained in table II), is given below[16]:

$$\underbrace{\mathcal{H}_b}_{\text{Anomalous Interaction}} = g_J \mu_B \mathbf{J} \cdot \mathbf{b_e} - \mu_N \hat{\mathbf{I}} \cdot \mathbf{b_N} \equiv g_J \mu_B \mathbf{J} \cdot \mathbf{b_e} - \frac{\mu_N}{\|\mathbf{I}\|} \mathbf{I} \cdot \mathbf{b_N}. \tag{27}$$

2. Measurement

As seen above (equation 23) the dynamics of a co-magnetometer contain many terms. During an experiment one of these terms needs to be isolated to be measured. There are a few different ways to achieve this; one possibility is to use a compensation field[9][16] to mask/cancel-out the fields being ignored.

In most cases you want to prevent a feedback effect occurring where the noble gas interacts with the alkali via the $\lambda M_N \hat{\mathbf{I}}$ term (equation 23) and then the alkali interacts with the noble gas via the $\lambda M_e \hat{\mathbf{F}}$ term (equation 26) and so on. To do this either or both of these terms need to be masked. Usually the impact of the alkali on the noble gas ($\lambda M_e \hat{\mathbf{F}}$) is masked. This is because only the alkali can be measured via the probe laser, in which case the noble gas must influence the alkali to be able to detect the dynamics on the noble gas.

[12] It is a measure of how the total atomic angular momentum, $\mathbf{F}$, is distributed between the electron and nucleus. Equivalently it can be seen as an energy shift due to (hyperfine) interaction between the nucleus and electron.

[a] Since it assumes on the relevant timescales the atoms probed are non-rotating in an inertial frame and the lasers/lab frame are rotating.

Term	Explanation
$\mathbf{B}_{\text{ext}}$	This is the external magnetic field from both the environment and that imposed by us; it is the term identical to that in the Bell-Bloom.
$\lambda M_{N(e)}\hat{\mathbf{I}}(\hat{\mathbf{F}})$	The effective B-field generated by the other atomic species (from their rotating spin under $\hat{\mathbf{B}}_{ext}$). λ is a coefficient with units Kelvin; M_N is the macroscopic magnetic moment: $n_N\mu_N$.
$\mathbf{L}$	The 'light shift' or a.c. Stark Shift caused by the laser(s) [2]. This effect causes the atomic energy levels to shift slightly and hence impacts the system as an external magnetic field would.
$\mathbf{b_{e,N}}$	Represents an anomalous field that acts like a magnetic field. The different subscripts denote that it may couple to electronic and nuclear spins differently as shown in equation 27. This is the term that would be measured to try to detect dark matter presence.
Relaxation terms	These include all decay terms to the precession. The terms within are not detailed here as they are numerous. They consist of collisional terms that decohere the system (precession) including: spin-exchange collisions (SECs), spin destruction collisions (SDC), photon-atom interactions ('collisions'), diffusion (leading to collision with cell walls). These terms should be minimised in an experiment to get the strongest signal.
Rotating body correction	This term is derived classically by considering how the dynamics of a static rigid body can be transformed into a rotating reference frame[a][17]. This is the term measured when using the co-magnetometer as a gyroscope.

Table II: Table explaining the meaning of all important terms in the co-magnetometer equations (23, 26).

III. FIRST EXPERIMENT MAGNETIC SHIELDING

For co-magnetometry it is important to have a very well controlled external magnetic field so that you both know what every reading from the magnetometer means and so that you can apply a very precisely known field if needed (i.e. to mask terms, see section II C 2). To achieve this a magnetic shield is used to prevent outside magnetic fields influencing the experiment (as well as precise coils inside, see section IV).

A. Background Theory

There are 2 main considerations for a magnetic shield design: the shielding factor and the magnetic noise (of the innermost layer).

The shielding factor is given as

$$S_T = S_n \prod_{i=1}^{n-1} S_i \left[1 - \left(\frac{D_i}{D_{i+1}}\right)^k\right], \tag{28}$$

where there are n layers of shielding and layer i has diameter D_i. k depends on the geometry; in this case (cylindrical) k is 2. S_i is the shielding factor for an individual layer defined as

$$S_i = \frac{\mu_i t_i}{D_i}, \tag{29}$$

where t_i is the thickness of the layer, and μ_i its permeability. The derivations can be seen in reference [18].

The noise of a shield layer can be calculated for simple shapes to be

$$\delta B = \frac{\mu_0}{R}\sqrt{\frac{k_B T t}{\rho_\Omega}}\ \underbrace{\sqrt{\frac{2}{3\pi}G}}_{\text{Adjusts for geometry}}, \tag{30}$$

where R is the radius, T the temperature, t the thickness, ρ_Ω the (electrical) resitivity, and G is a scaling factor for which the value changes for different cylindrical geometries[13]. The derivation can be seen in reference [19].

B. Custom Shield Design

Specification: The shield needed to reach a theoretical shielding factor of at least 10^6 with a noise $\sim 1\text{fT}/\sqrt{\text{Hz}}$ or less. This noise level is needed to compete with state of the art magne-

[13] This equation holds for infinite planes and spheres as well by changing the coefficient on the end labeled 'Adjusts for geometry'.

tometry, whilst the shielding factor is to ensure that any typical magnetic field in the lab (nominally the Earth's magnetic field) is suppressed far below world leading magnetically shielded environments[20], such that it is not limiting (in which case the residual magnetisation will be, as discussed in section III C).

Increasing shielding factor is simple as more layers can always be added, as such the noise of the inner layer was considered first. By looking at equation 30 it can be seen that there are only 4 terms[14] that can be adjusted to improve the noise: radius, thickness, resitivity, G factor[15].

The only term that relates to properties of the material is the resistivity, whereby it is desirable to have very high values. This is due to the noise being derived from power losses in the material caused by electromagnetic fields[19]. The main two sources of loss are eddy current losses and hysteresis losses, the former of which is reduced by high resitivity as less current can flow.

MuMetal is the standard material for magnetic shield manufacture[9][21][22] due to its very high permeability[22] which improves its shielding factor (see equation 28). However, its resistivity is very low. The equation for noise (30) has no dependence on permeability implying it can be significantly lower for the inner layer, in order to improve noise, and compensated for by additional outer layers (to achieve desired shielding factor). Based off this alternative materials were looked into with higher resistivities to use for the inner layer[16].

C. Ferrite Design

Newer shield designs incorporate an innermost layer of ferrite to achieve the low noise[6][23]. Ferrite is commonly used in transformers and electrical components to reduce noise (damp magnetic fields).

There is a long list of available ferrites with a wide range of properties. The ferrite for the shield was selected to have the following: high permeability, high resistivity, high Curie temperature. On top of this the ferrite needs to be a soft ferrite to allow it to be easily degaussed[17].

The ferrite used in this experiment was 3C96 by Ferroxcube[24]. Its properties in comparison to muMetal are given in table III. Due to the ferrite's brittleness it has to be much thicker than muMetal would, for structural reasons. Even given this (see equation 30) its noise is orders of magnitude lower than for muMetal.

[14] This assumes it is impractical to cool the shield.

[15] The G factor relates the cylinders diameter to its length; to reduce noise it should be longer than wide[19].

[16] The range over which the other parameters could be altered is much smaller than the range the resistivity can change. In fact, during some testing it was found impractical to reach the $1\text{fT}/\sqrt{\text{Hz}}$ noise desired with muMetal.

[17] Degaussing is a process to reduce the residual magnetic field present in a material. It involves saturating the material with a strong alternating current (AC) magnetic field and then slowly decreasing the field uniformly such that the magnetism is distributed uniformly. A soft ferrite has a low coercive force (field strength to saturate) in comparison to hard ferrites.

Material	**Noise/** $\mathrm{fT}/\sqrt{\mathrm{Hz}}$	**Permeability**	**Resitivity/** $\Omega\ m$
3C96	0.014	2000	5
MuMetal	15	10,000[a]	6×10^{-7}

Table III: Table comparing properties and noise of the ferrite, C396, and muMetal. It is important to note the thicknesses of the materials are different. For muMetal it is 2mm thick, where as the ferrite is 12.7mm thick as it cannot be machined this thin due to its brittleness.

The volume inside the shield for the experiment should be maximised for ease of access. The maximum machine-able size for the ferrite had an outer diameter (OD) of 250 mm, with a half inch (12.7 mm) thick wall (for structural reasons, due to the ferrite brittleness). The length of the piece is also 250 mm for structural (and cost) reasons; this makes the G factor (equation 30) worse than for a longer piece, however, considering the vast improvement ferrite offers this compromise was still beneficial.

Despite some modern shields using ferrite they usually have small interiors or use ferrite sheets to create a box inside[23]. This shield design is the first large cylindrical ferrite layer to be created[18] (to our knowledge), hence machining it proved challenging. Ferrite pieces do not come in blocks large enough to machine the piece out of. To overcome this problem the company manufacturing the piece (Gateway Cando) designed it using segments that would then be held together by a high temperature ferrite glue.

(a) **(b)**

Figure 6: Images showing the manufacture of the ferrite layer by machining individual segments (b) and gluing them together (a).

[18]In fact the company making the ferrite asked to double check the size as they believe it will be the second largest piece of ferrite next to pieces built for the LHC.

[a]The value given by MSL on their website[22] (who are making the shield) is 470,000. However, papers report muMetal permeability to be in the 10,000-100,000 range[21] as well as a contact from the company.

D. MuMetal Design

The geometrical design of the muMetal layers was based around the ferrite piece, and the magnetic design was to bring the shielding factor up to at least the order of 10^6.

Design of the specifics for the muMetal layers was based off testing and analysing the equation for the shielding factor (28).

1. Layer Spacing

By varying the spacing between shield layers, $D_{i+1} - D_i$, in equation 28 whilst keeping the other parameters constant the ideal spacing for a shield with the outer diameter (OD) of the inner layer specified (see section III C) was found, as shown in figure 7. For simplicity this was rounded and a spacing of 50 mm was selected.

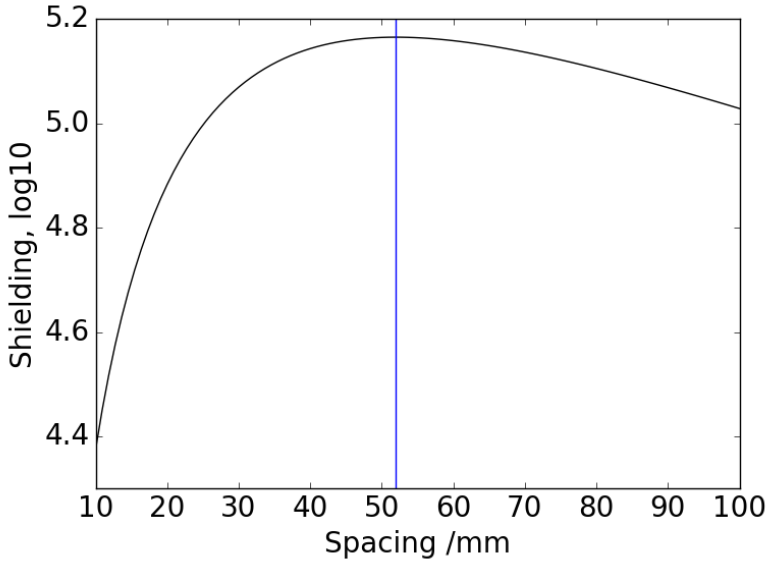

Figure 7: Figure showing the relationship between layer spacing and the shielding factor for a multi-layer magnetic shield, based off equation 28. The blue line shows the detected maximum on the curve.

2. Thickness

The actual permeability of muMetal varies. An analysis was performed using the shielding factor equation (28) to see the impact of the possible range of muMetal permeability[21]. The analysis was based on a shield design with 4 layers (3 muMetal, 1 mm thick, 1 inner ferrite[19], 12.7 mm thick), 50 mm spacing. As seen in figure 8, the analysis showed that the muMetal permeability only needed to be $\sim 20,000$ to reach the desired shielding factor, see figure 8. The contact at Magnetic Shields Ltd. (MSL) stated that 'a value of 10,000 is more expected' for the

[19] At this point a slightly different ferrite was being used but it does not affect the conclusions drawn. The ferrite was changed due to stock.

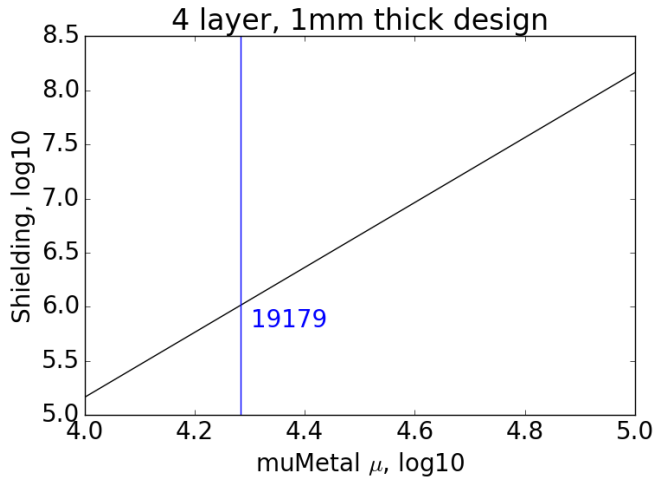

Figure 8: Figure showing the relationship between muMetal permeability, μ, and the shielding factor for a multi-layer magnetic shield, based off equation 28. The blue line shows the point at which the shielding factor reaches 10^6.

permeability of muMetal rather than the stated maximum value of 470,000 on their website[22]. Based off this, an analysis of the impact of muMetal thickness on shielding factor and the shield mass was conducted (based on a permeability of 10,000), the results of which are seen in figure 9.

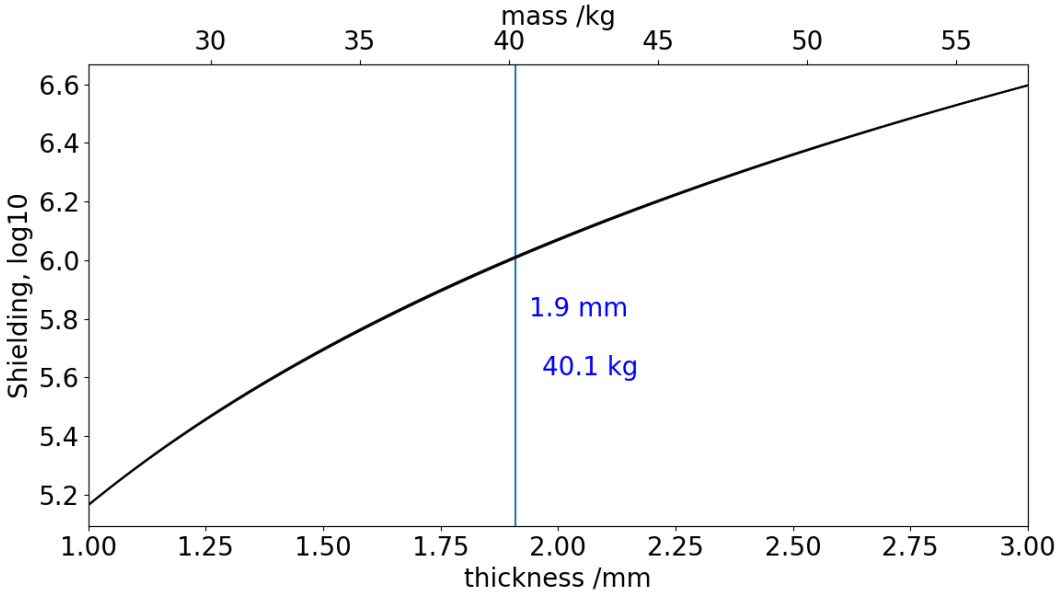

Figure 9: Figure showing the impact of muMetal thickness on the shielding factor and mass of the magnetic shield. The design is still 4 layers, with an inner 12.7 mm thick ferrite piece with 50mm spacing.

Based on this, despite the large mass and cost increase, it was decided to use 2mm thick muMetal to ensure a shielding factor of at least 10^6 was achieved by the shield even if the muMetal permeability was at its minimum of 10,000.

The full design of the shield can be seen in Appendix B, figure 44. Table IV shows the final design specifications and theoretical properties of the shield.

Shielding Factor	0.93×10^6 [a]
Noise	$14\text{aT}/\sqrt{\text{Hz}}$
Thickness (number) of ferrite layers	12.7 mm (1)
Thickness (number) of muMetal layers	2 mm (3)
Spacing between layers	50 mm
Inner diameter	224.6 mm
Length to diameter ratio	1

Table IV: Theoretical properties and specification of shield design.

E. Design Verification

To verify the results given by the shielding factor equation (28)[20] a simple computer-aided design (CAD) design was made for the shield and an electromagnetic simulation of it was ran using the software Ansys AIM. To simulate the Earth's magnetic field[21] a very large Helmholtz coil [25] was made in CAD such that the shield could be placed in the centre most 20% region, where the field is reasonably uniform. The design was simulated with and without holes to see their impact. In the case with holes, a 30mm pair of holes was placed along the central axis of the cylinder (for the pump laser) and another 30mm pair central and perpendicular to this one (for the probe laser), a further 20mm hole was placed to one side of one of the 30mm holes on the cylindrical curved surface (for cable access).

Data was extracted from Ansys AIM using the line chart to csv function. Python was used to analyse the data. The code calculated an external field strength[22] from the data which, due to the curve imperfections caused when adding holes (see figure 10), varies. The reference value is calculated by taking the average value of the data above $10^{-4.5}$ T; an improvement to this method would clip the top data as well so as not to skew as seen in figure 10c. Similarly the value of the bottom of the well was calculated by the average of data below 10^{-9} T. This means the shielding factor calculated from these two are not perfect; instead of using the minimum field value or centre value it takes into account the imperfect sloping, however, this is sufficient for an order of magnitude check.

Figure 10a clearly shows the shielding factor is on the order of 10^6 or above. Adding the holes appears to keep the shielding factor to the same order of magnitude and had minimal impact towards the centre of the shield. Considering the lines taken are directly through the holes this

[a] Slightly below 10^6 spec due to last minute change of ferrite material.

[20] Ansys could not be used to verify the noise.

[21] This assumes the external fields around the shield will be dominated by the Earth's field.

[22] This is done as a quick check mechanism as the shielding factor needs only to be to order of magnitude.

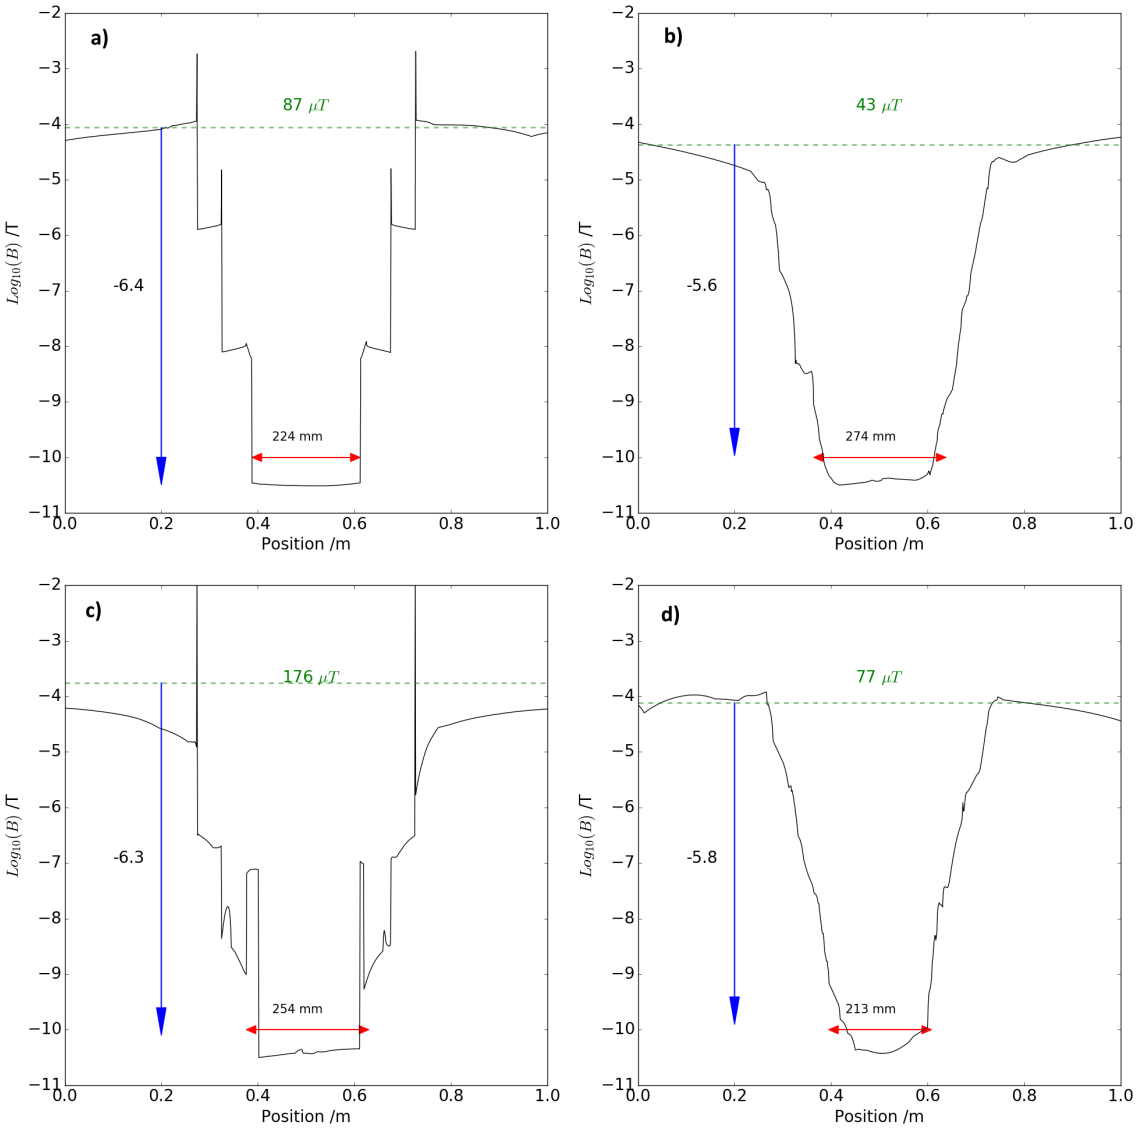

Figure 10: Figure to show the simulation of the magnetic shield in Ansys AIM. Graph 'a' shows the B-field strength along the z-axis (parallel to Helmholtz axis) for the design with no holes. The other 3 graphs show the results for the design with holes. Graph 'b' is through the shield central axis, graph 'd' is through the perpendicular z axis (aligned to Helmholtz), and graph 'c' through the final axis. The shielding factor is shown by the blue arrow, the calculated width inside the shield (at low field) by the red arrow, and the calculated external field (used for calculating the shielding factor) by the green line.

data indicates the effect of the holes is not significant on the overall performance.

The width of the wells (red arrows figure 10) are generated quite inaccurately however it is clear that adding the holes maintains at least 30 mm (size of vapour cell, see section VIII) of low B in the centre. The graphs do show a 'bumpy-ness' in this region. This maybe due to simulation inaccuracies from the software limitations. The field uniformity should be controlled by the coils

inside (section IV) which generates a field strength ($\sim 10^5\,\mathrm{T}$) orders of magnitude above this minimum, hence if this is an effect from the holes it shouldn't impact the experiment.

F. Manufacture

The company manufacturing the magnetic shield (MSL) cannot work with ferrite. As such a separate company that specialise in machining ferrite (Gateway Cando) was asked to manufacture the ferrite layer and then ship it to MSL for installation into the rest of the shield.

IV. MAGNETIC COILS DESIGN

As outlined in sections II C 2, and III a precise, uniform magnetic field needs to be generated over the vapour cell. This prevents inhomogeneous broadening of the magnetic resonance from the non-unfirm field. The precision of the field is limited by the power supply for the magnetic coils and hence this experiment will use an ultra-precise current supply. The field uniformity is determined by the coil geometry used.

The coil designed for the experiment consists of a separate coil geometry for each axis of field control. Along the pump axis (parallel to shield central axis) is a Maxwell coil; this provides a massive improvement in field uniformity over the standard Helmholtz coil[26][27]. The other two axes contain a saddle coil pair.

A. Saddle Coils

Saddle coils are commonly used in magnetic resonance imaging (MRI) machines[28]. This is because they produce good field uniformity and are designed to fit in a cylindrical configuration. It is created by projecting a rectangle over a certain angle of a cylinders curved surface, as seen in figure 11.

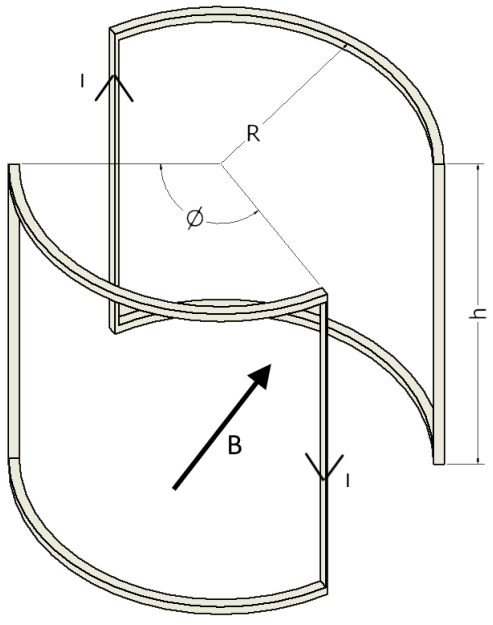

Figure 11: Drawing showing the geometry of a saddle coil pair. The defining parameters are labelled. In bold the current and field directions are labelled.

A paper by Bonetto et al[29] studied the parameter space when defining a saddle coil, labelled in figure 11 to find the optimal geometry for field uniformity. For the angle ϕ they find 120° to be optimal, as such both coils have been designed for this. The optimal ratio of $\frac{h}{R}$, seen in figure 11,

(η) is 4. To maintain space within the shield centre, for the cell and access, this is not practical as the shield has an η of 2, making the radius of the saddle pair (that fits lengthwise) too small. As such the saddle coils have been designed with equal diameter and length, as with the shield, as maximising space is important to generate larger areas of field uniformity.

Due to the efficient space usage they exceed the performance of similar sized Helmholtz coils, see figures 45, 46 in appendix C. Each coil geometry was simulated in Ansys for analysis. Each simulation used the same total current (300 A·turns, arbitrary choice) to first compare the central field strength but second to keep the central field strength comparable, between coils, such that the field strengths as a percentage for each coil are comparable.

1. Ansys-Python Coil Analysis

The region of 1% uniformity was found using Python, for comparison. The vector data was taken from Ansys as a '.csv'; this data has a B_x, B_y, B_z value for every point in the simulation[23]. The Python script sorts the data along one (spatial) axis and then extracts a line along this axis through the centre. It does this by keeping points where the other 2 coordinates are at the specified values for the centre within an error bound[24]. Using this reduced data set it calculates the $\|\mathbf{B}\|$ for the centre point and then uses this to calculate the relative difference for all points on the line (with the central value); further the position data is also made relative to the centre point. The percentage difference is then calculated and the data is clipped for graphing. To calculate the uniformity region the percentage difference subset is interpolated over (to improve precision). A logical OR operator further reduces this subset to values not in this uniformity

		1% uniform region width/mm		
Coil	**Centre Field**/T	**x Axis**	**y Axis**	**z Axis**
Helmholtz	8×10^{-5}	4.0	5.0	0.2
Saddle	2×10^{-4}	49.4	25.1	33.5
Maxwell	1×10^{-4}	99.2	110.2	78.1

Table V: Table showing the field uniformity of different coil types. Each coil geometry used a total of 300 A·turns (301 for Maxwell) and a radius of 100 mm (110 mm for Maxwell). Note the z axis is defined through the coil face, see figure 12a.

[23]This method was used over the previous method of extracting a line from Ansys as it allows greater flexibility. For example it also allowed analysis of the angle of the field and the angle uniformity (important to consider when all 3 coil axes are used in the final design).

[24]Simulation precision means, for example, 2.0 isn't a point, but 1.99 and 2.02 are and both actually mean 2.0.

region. The signs of the relative positional data (for the changing coordinate) for this subset are taken. Hence the boundaries for the region are found by looking for a sign change (since the region must have 0 relative position within it). This is then repeated for the other axes.

The results of simulating each geometry are given in table V. As seen the Maxwell configuration is far superior to either of the other two, in agreement with theory[26][27]. Further the saddle coil geometry is seen to outperform the Helmholtz geometry in field uniformity significantly as well as producing an overall larger field for the same current.

B. Final Design - First Experiment

The final geometry can been seen in figure 12. As stated previously (section IV), the design consists of 1 Maxwell coil and 2 saddle coils. To enable laser access to the shield centre a 30mm hole was added to the shield side (see appendix 44). Since a Maxwell coil has a central coil[30][27], to allow the laser to pass through, a 30 mm hole which the coil goes around was added (see figure 12a. The saddle coils were added to fit between the outer and central Maxwell coils (which require set separations).

	1% uniform region width/mm		
Coil	**x Axis**	**y Axis**	**z Axis**
Maxwell with Holes	126.3	108.2	79.7
Final configuration	54.7	63.7	45.1

Table VI: Table showing the impact on the uniformity region of adding holes to the Maxwell coil and using the saddle coils in conjunction. Note the z axis is defined through the coil face, see figure 12a.

Table VI shows the analysis of the final coil design. One simulation with only the Maxwell coil but with the holes added, the other the full configuration where the saddle coils are set to make the field 45° with respect to each axis[25]. Surprisingly adding the holes improves the performance of the coils. This could be because the hole gives the central coil some extent along z-axis beyond simply its thickness (z-axis runs through the central axis of the (total) cylindrical coil geometry (see figure 12a)). The symmetrical nature of this may lead to better field uniformity in the region close to these holes. This in turn would lead to an observed increase in uniformity along x-axis as this axis goes through the holes. A brief investigation was performed when adding another pair

[25]The relative relationship between the coil sets field strength and current were derived empirically, i.e. the current needed to drive field a certain central field strength for each coil geometry was determined from a regression on simulated data (adding holes to the geometry made analytic calculation more time consuming than was warranted). As such numerically errors exist; not the x-z angle was actually 44°, and y-z 46°, not the desired 45° (equal field strength from all coils, taken as an example case for uniformity).

of opposing holes just above the current one. Initial results showed worse performance, however, further research into the area of adding holes to the coil, or splitting the central coil into two, could lead to improved field uniformity.

As expected using the saddle coils reduces the field uniformity as the Maxwell coil is superior. Notably the y-axis, now has the largest uniformity region (comparing to Maxwell with holes) likely as this axis goes through the outer saddle coil. The saddle coil acting in this axis is larger,

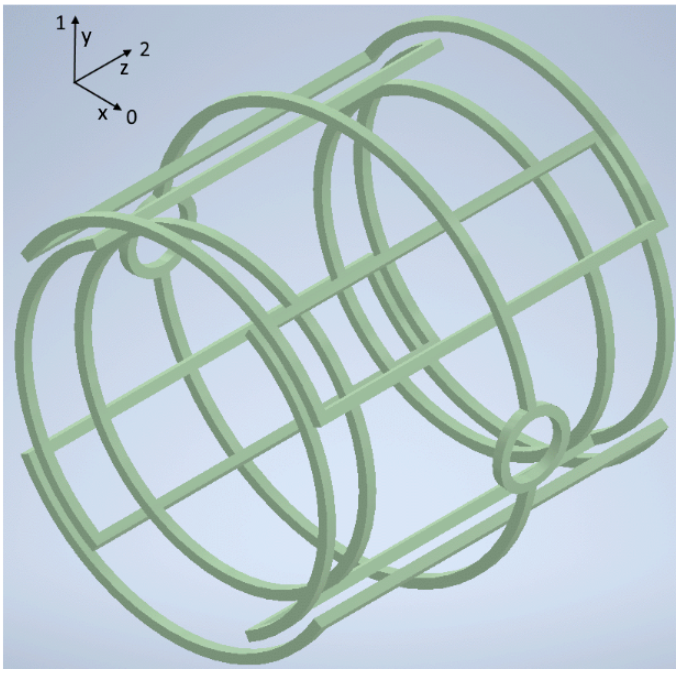

(a)

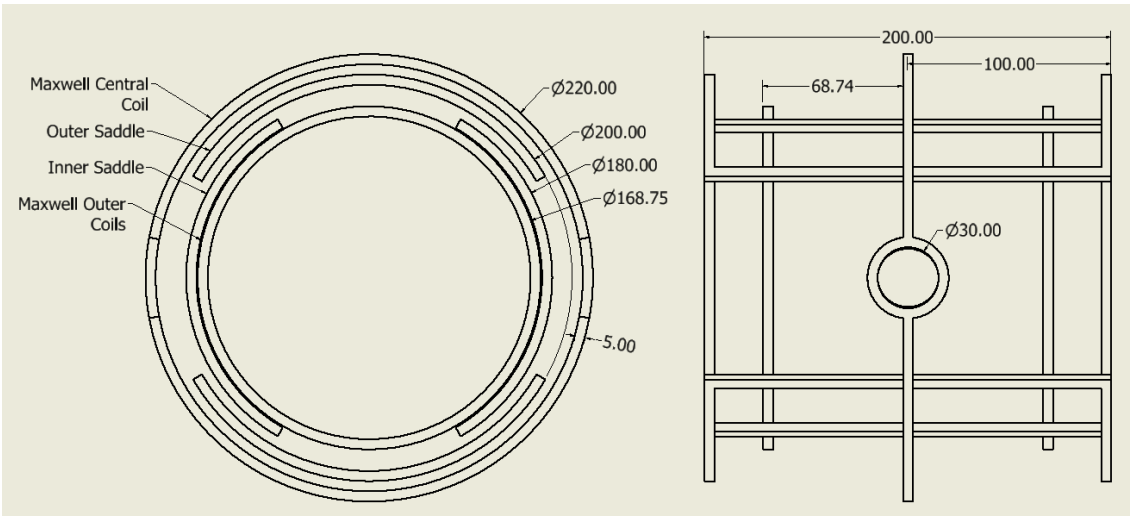

(b)

Figure 12: Figure a shows an image of the complete coil design at an angled perspective to aid visualisation, an axis in the top left is included to line up with table VI. Figure b shows a drawing detailing the specifics of the design, all dimensions in mm.

meaning its uniformity region is better than the other saddle, degrading the Maxwell coil less.

Therefore optimal field uniformity can only be achieved in one axis, however, the saddle coils remain in the case an angle to the field is desired.

1. Mounting inside the shield

The coil design above (section IV B) needs CAD made, which is the inverse of the current CAD, on which the wire can be wound. After creating the inverse[26] excess material was removed to reduce the weight of the design to minimise the load on the brittle ferrite layer (see section III C). Further cable access holes were added to bring the wire in and out of the mount. The coils were broken up into multiple layers: central Maxwell, outer saddle, inner saddle, outer Maxwell (innermost layer). Each layer is designed to slide over each other, as such the layers will be made of Nylon 6 to aid the sliding and for long term wear. A slot and extrusion was added to relevant layers such that the extrusion from one layer slots into the next to lock the layers together in the correct alignment (see figure 13).

Within the coil mounts a mechanism for mounting the vapour cell (see section VIII for specs) and heating it was designed. This is because it needs to be machined from non-metallic materials to prevent it interfering with magnetic fields.

To allow flexibility to add or change components later an optical breadboard was designed. The board attaches via screws to the innermost coil mount. Due to the small space inside the shield the board uses M3 holes.

Most other experiments use a cube enclosure around the vapour cells[9][21] to create an oven to heat the cell. However, here a mount design was made to maximise the surface of exposed cell for the laser to hit, such that the large laser holes in the shield are not wasted. The design has the cell sit on a 'torus'[27] that is mounted onto the breadboard via a tripod design (see figure 14). There is also an adjustable 'piston' mounted on the top of the innermost coil mount to apply pressure on the top of the cell to ensure it is securely held. The 'piston' is adjusted/held by a screw.

To heat the cell an oven is placed underneath the cell. An adjustable (held by screw) hollow tube joins the oven to the cell. This tube creates a seal with the 'torus' mount the cell is on. The air inside the oven is heated[28] to the desired temperature and in turn heats the cell it is in thermal contact with. Unlike the rest of the components the oven will be made from PEEK due

[26] Achieved by extruding the current design into a cylinder.
[27] With an O-ring within a groove to improve grip.
[28] The method for this is currently undecided, though it will likely be hot air flow.

to its higher operating temperature. The oven has foam pads on the tube to improve the seal and on its base to limit heat transfer to the Nylon breadboard.

Figure 13 shows an image of all the interior components to help visualise the above information, especially the coil mounts. Figure 14 shows a drawing of the whole assembly for clearer detail and to visualise components not visible in figure 13.

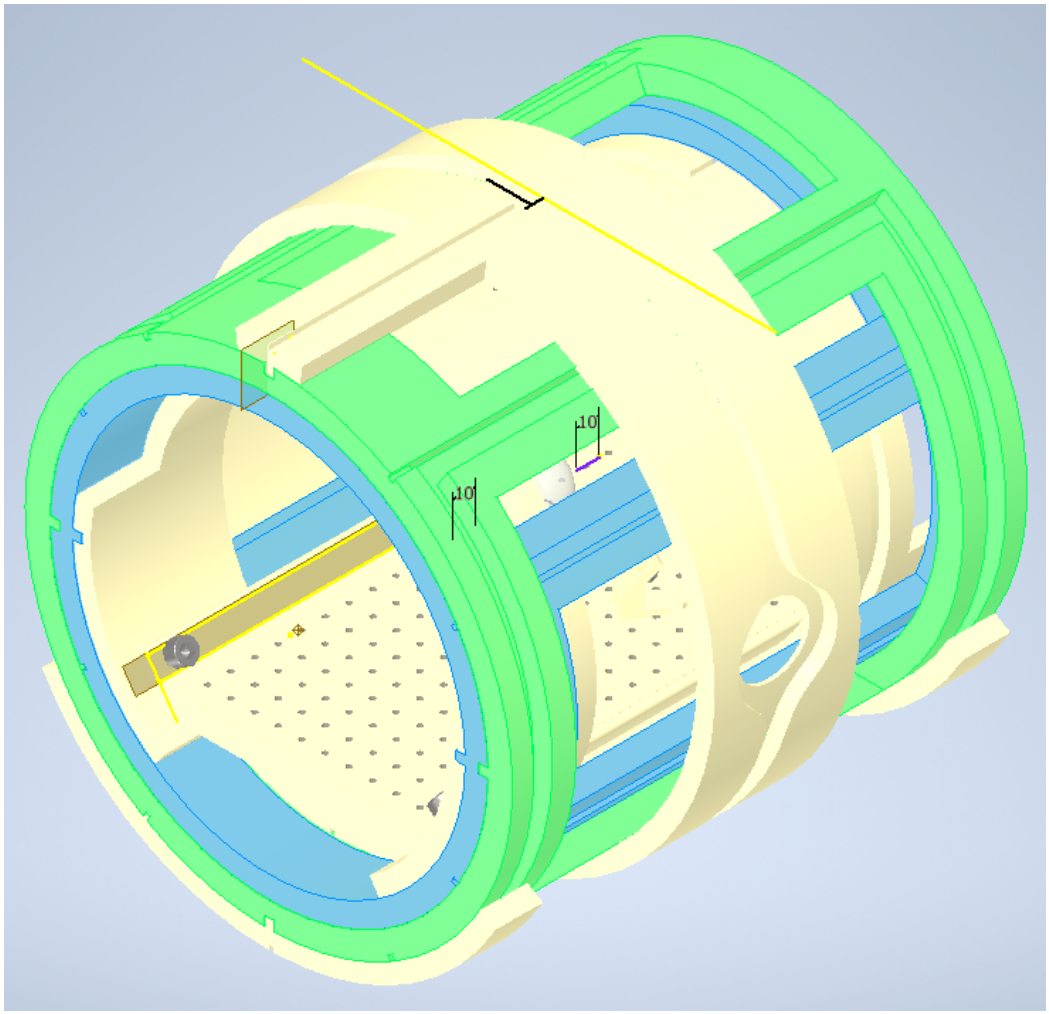

Figure 13: Image showing CAD of all shield interior components at an angled perspective. The 4 coil mount layers have been coloured to visually separate them. Within the coil mounts, the custom optical breadboard can be seen. The custom oven, and cell mount are obscured by the coil mounts.

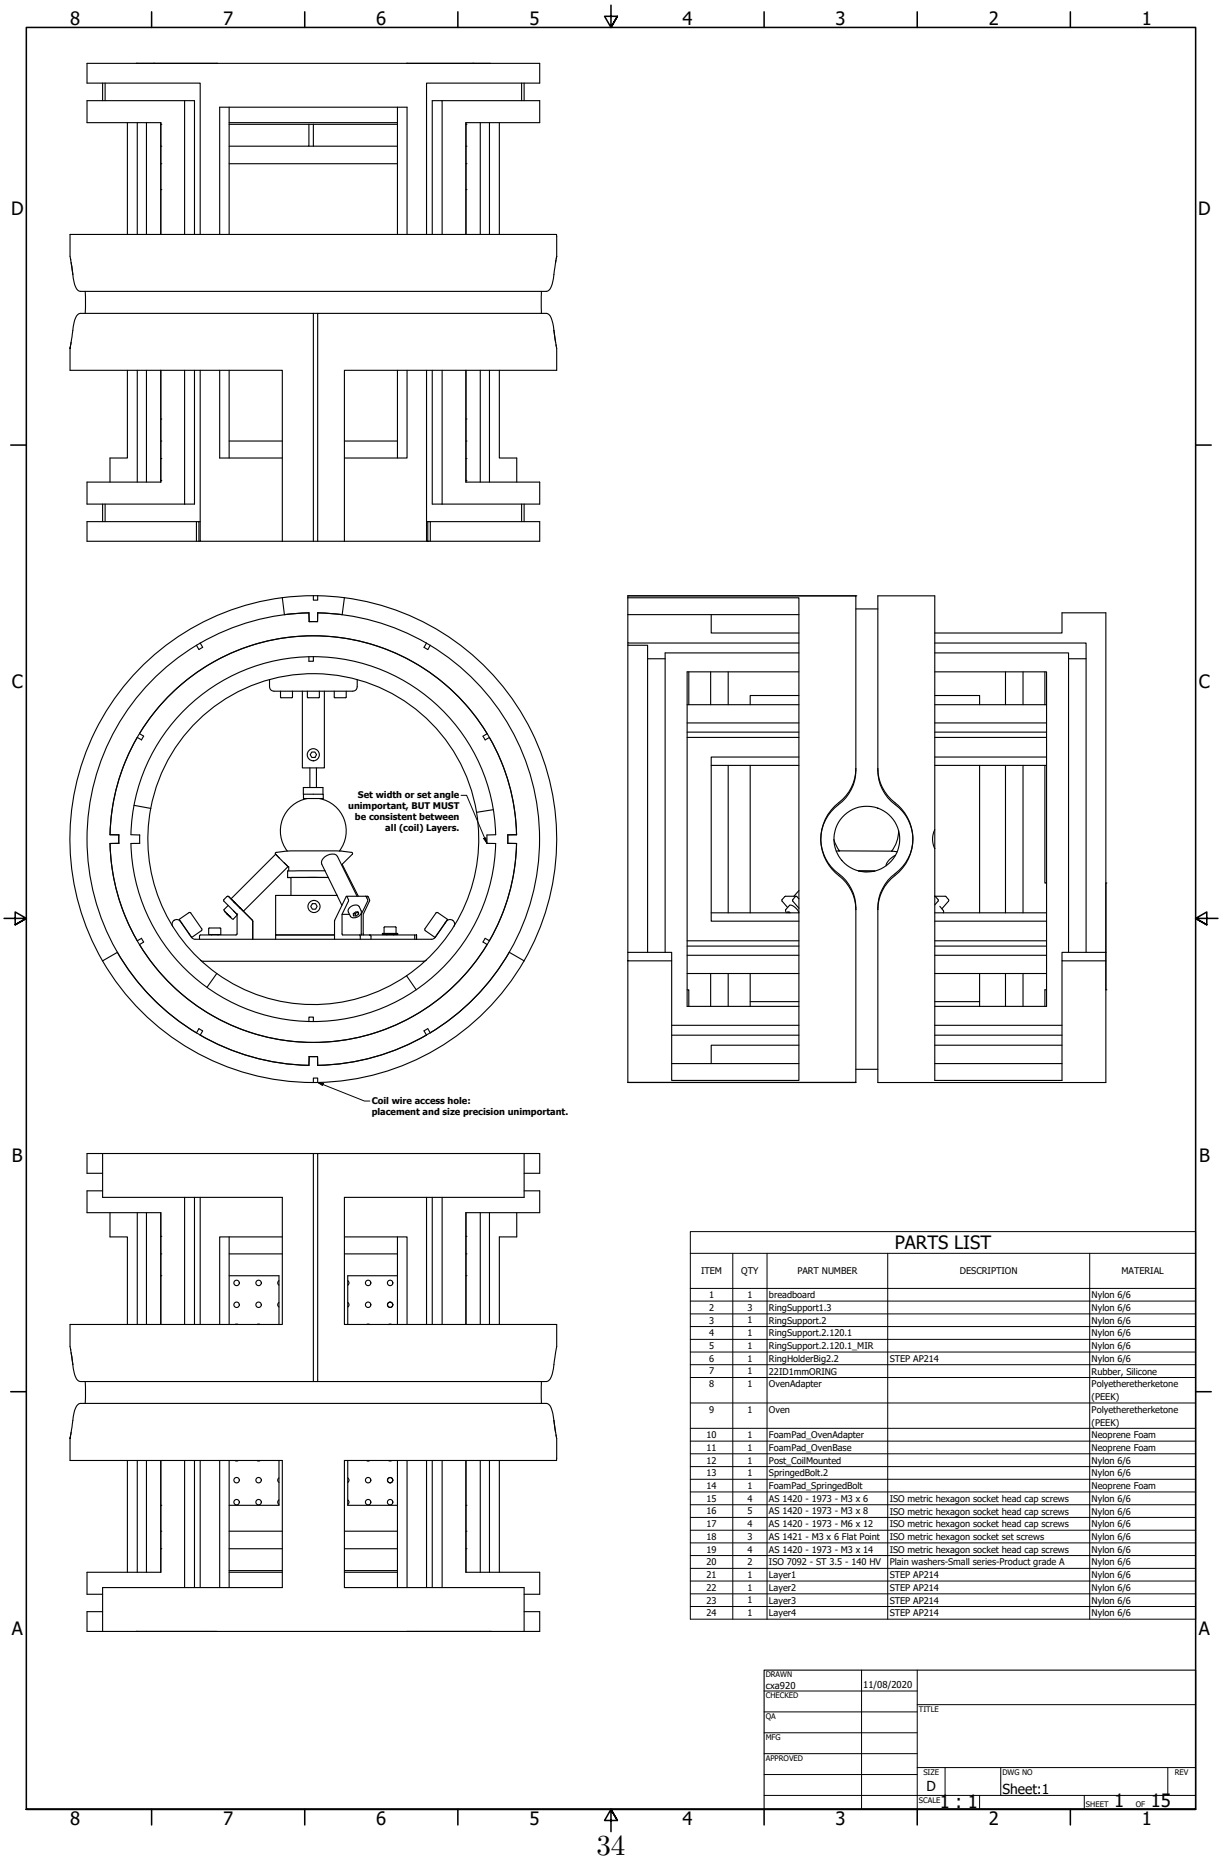

PARTS LIST

ITEM	QTY	PART NUMBER	DESCRIPTION	MATERIAL
1	1	breadboard		Nylon 6/6
2	3	RingSupport1.3		Nylon 6/6
3	1	RingSupport.2		Nylon 6/6
4	1	RingSupport.2.120.1		Nylon 6/6
5	1	RingSupport.2.120.1_MIR		Nylon 6/6
6	1	RingHolderBig2.2	STEP AP214	Nylon 6/6
7	1	22ID1mmORING		Rubber, Silicone
8	1	OvenAdapter		Polyetheretherketone (PEEK)
9	1	Oven		Polyetheretherketone (PEEK)
10	1	FoamPad_OvenAdapter		Neoprene Foam
11	1	FoamPad_OvenBase		Neoprene Foam
12	1	Post_CoilMounted		Nylon 6/6
13	1	SpringedBolt.2		Nylon 6/6
14	1	FoamPad_SpringedBolt		Neoprene Foam
15	4	AS 1420 - 1973 - M3 x 6	ISO metric hexagon socket head cap screws	Nylon 6/6
16	5	AS 1420 - 1973 - M3 x 8	ISO metric hexagon socket head cap screws	Nylon 6/6
17	4	AS 1420 - 1973 - M6 x 12	ISO metric hexagon socket head cap screws	Nylon 6/6
18	3	AS 1421 - M3 x 6 Flat Point	ISO metric hexagon socket set screws	Nylon 6/6
19	4	AS 1420 - 1973 - M3 x 14	ISO metric hexagon socket head cap screws	Nylon 6/6
20	2	ISO 7092 - ST 3.5 - 140 HV	Plain washers-Small series-Product grade A	Nylon 6/6
21	1	Layer1	STEP AP214	Nylon 6/6
22	1	Layer2	STEP AP214	Nylon 6/6
23	1	Layer3	STEP AP214	Nylon 6/6
24	1	Layer4	STEP AP214	Nylon 6/6

DRAWN	cxa920	11/08/2020
CHECKED		
QA		
MFG		
APPROVED		

TITLE

SIZE D | DWG NO Sheet:1 | REV

SCALE 1 : 1 | SHEET 1 OF 15

Figure 14: Drawing of all components to be installed inside the shield: coils, oven, vapour cell and mount.

V. FIRST EXPERIMENT OPTICAL SETUP

The optical setup is relatively simple and consists primarily of locking the lasers and giving them the correct polarisations. Steps to control the amplitude, frequency and modulation of the laser (via acousto-optic modulator (AOM), or laser controller) are also included.

A. Laser frequency stabilisation

The pump laser is locked (the laser frequency is stabilised) using saturated absorption spectroscopy using, the pick-off, PBC1 (polarising beam cube) in figure 16; the half wave plate before PCB1 is rotate-able to control the power used for spectroscopy. The pump is locked to the closed transition, F=2 $\mapsto$ F=3, of ^{87}Rb D_2 (see figure 15), to simplify the setup such that we don't require a re-pumper. Since the F=2 $\mapsto$ F=3 line is weaker than neighbouring lines[31] we will lock to one of the crossover peaks, either the F=2 $\mapsto$ F= $\{1, 3\}$ overlap or the F= $\{2, 3\}$ overlap. Using the acousto-optic modulator (AOM) seen in figure 16a the pump will be shifted from the lock point to the desired F=2 $\mapsto$ F=3 transition.

The probe beam needs to use a transition from the F=1 ground state to avoid the power broadening and stark shift on the F=2 ground state caused by the pump beam. The net polarisation ($\hat{\mathbf{F}}$[29]) is transferred between atoms in the ground states via SECs. Further the beam needs to be detuned from its nearest transition by $\sim$1 GHz. This prevents a reduction in the signal on the photodetector (beam power) caused by absorption. Section II A 2 shows that the coefficient of absorption falls off faster with detuning than the coefficient of refraction; hence, by being far detuned, absorption power losses can be removed whilst still experiencing an optical rotation from refractive interactions, which produces a signal.

Large detuning can be achieved by AOM shifting or offset locking. Offset locking was chosen since AOMs with large detuning are expensive, and have a smaller bandwidth they can work over. Hence figure 15 shows the probe as between the F=1 to F=3 states (with 1 GHz detuning from resonance); the closest transition being F=1 $\mapsto$ F=2[30]. It is shown this way to indicate the actual frequency/energy of the beam since this is more relevant for the highly off resonant beam than the actual transition. Further it helps indicate that the beam is offset locked from the pump beam; the probe beam will be 7.8 GHz detuned from the pump beam, which equates to 1 GHz detuning from the other ground state.

The offset lock signal is formed using the 50:50 beamsplitter in figure 16a, which overlaps the

[29] Also seen in 'microscopic' picture as $\mathbf{m}_F$ transfer that averages to give $\hat{\mathbf{F}}$ in the 'macroscopic' picture.
[30] See appendix D for an explanation as to why the probe can act on another excited state.

pump and probe beam to create a beat signal (difference of frequencies). This signal is RF and detected by the photodiode; the generation of an error signal is managed by the offset lock circuit (see section VI), which is then fed back into the laser controller which locks the laser to this signal. The half wave plates before PBC2 and PBC3 control the power of the pump and probe beam respectively that is fed to the photodiode.

B. Polarisation

Once locked the 2 beams are translated to the magnetic shield in free space. This was chosen over fibres to avoid noise on the beam power.

Near the shield a double Glan-Taylor polariser[31] is used on each beam to create a very clean linear polarisation. For the pump beam a quarter waveplate converts this to a σ^+ polarisation to optically pump the atoms. For the probe beam a rotating half wave plate may be added, if necessary, to rotate the linear polarisation to the correct polarisation, π_z (not parallel to the pump axis[32]).

Both beams are then expanded to cover as much of the cell as practical to improve pumping uniformity and probe signal. The probe beam width is reduced and its two perpendicular linear polarisation powers are split using a Wollaston prism. Each linear polarisation power is detected by a photodiode built into a balanced photodetctor. This allows for a differential measurement allowing for common mode rejection (e.g. total beam power fluctations), which improves precision significantly.

The optical circuit can be seen in figure 16a for clarity.

[31] These have very high extinction ratios.

[32] The probe beam may be rotated off π_z to be placed at an angle such that the effect of its light shift from each perpendicular axis cancel/mitigate each other.

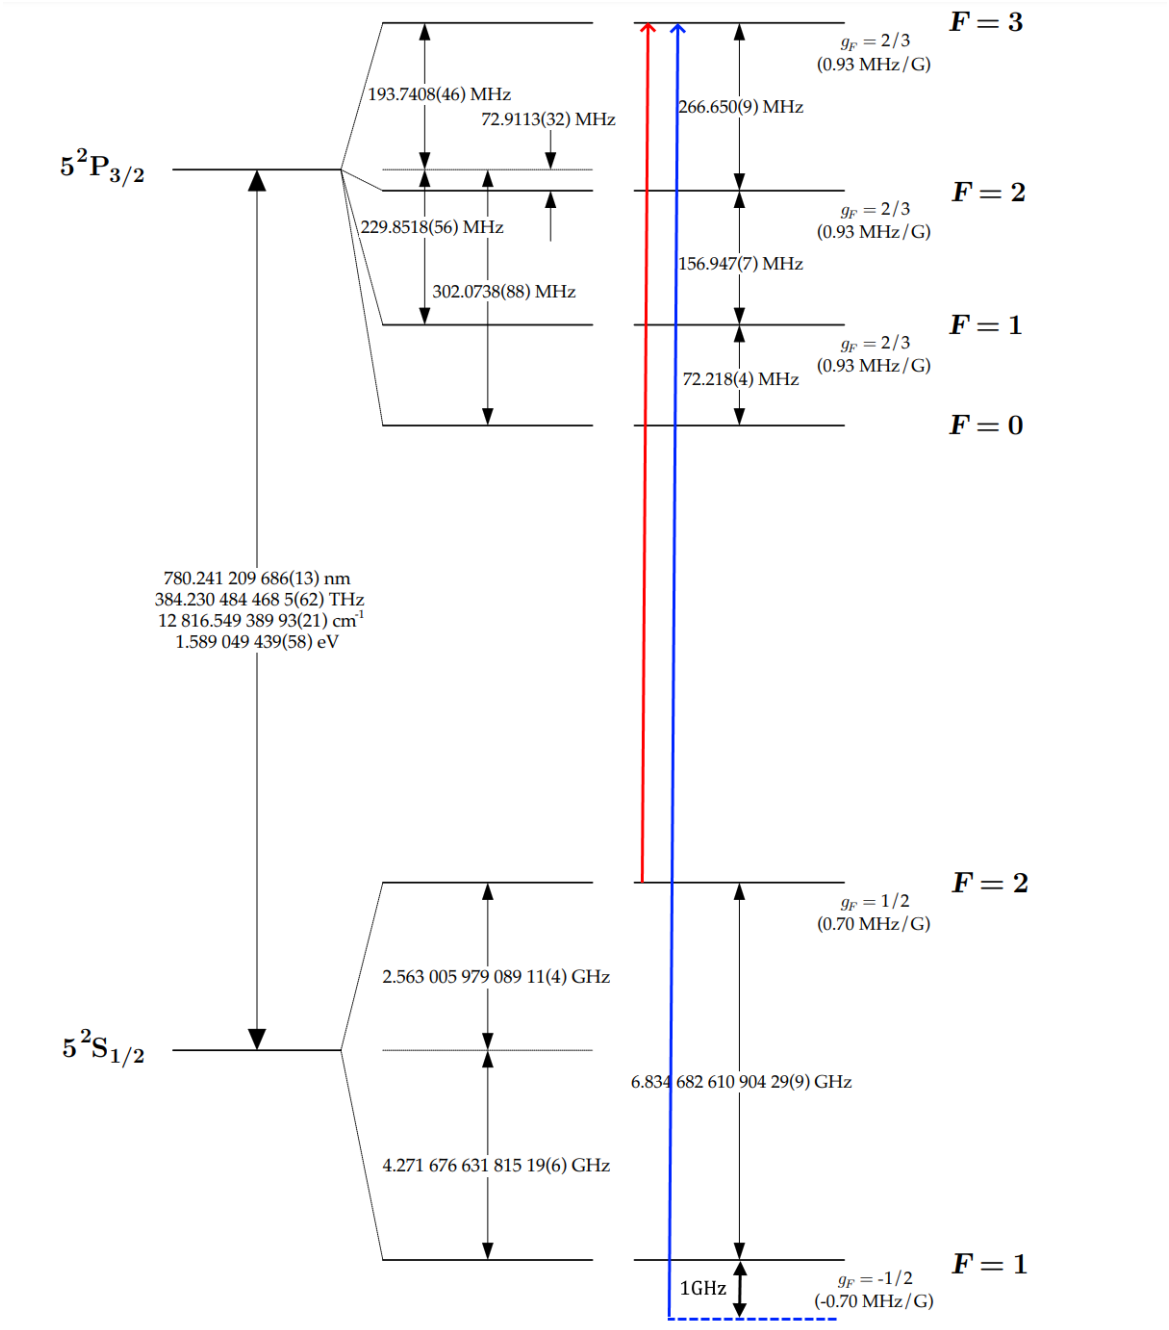

Figure 15: Image showing transitions used by the 2 lasers in the optical setup. The transitions shown are the hyperfine splitting of ^{87}Rb D_2 line. Pump laser in red, probe in blue (as in figure 16a). Modified image from 'Rubidium 87 D Line Data'[31].

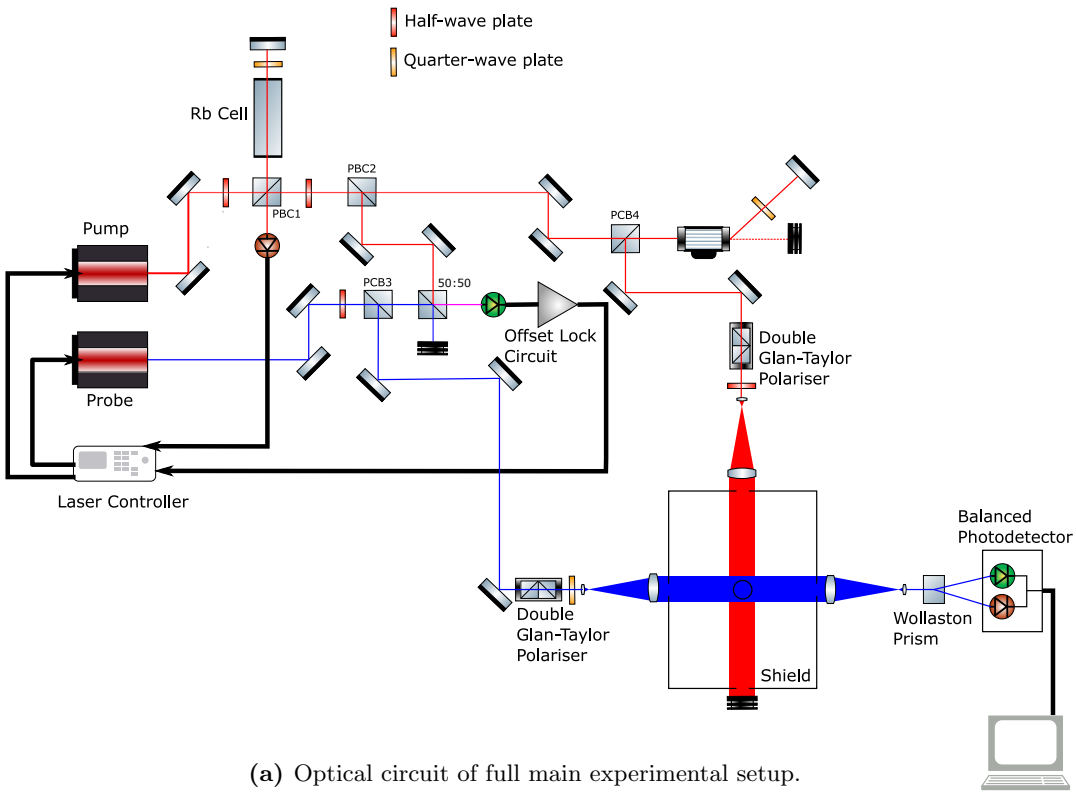

(a) Optical circuit of full main experimental setup.

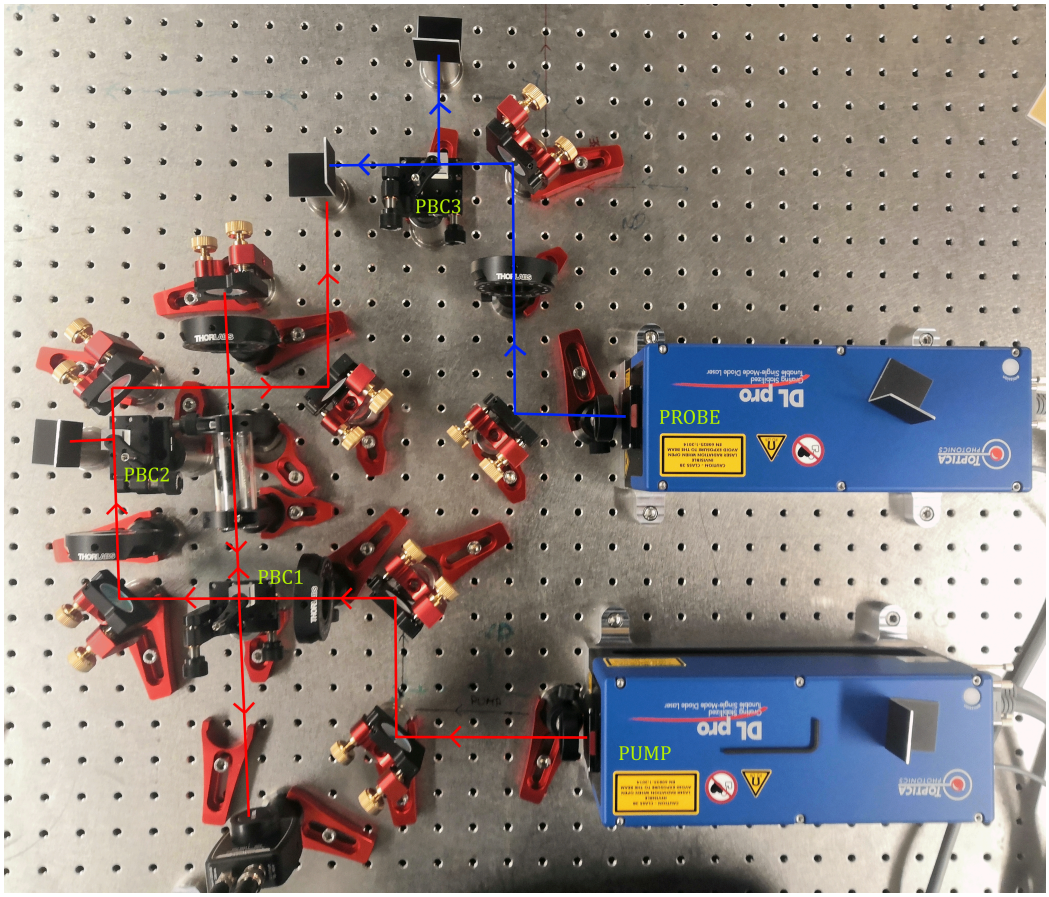

(b) Image of current optical setup. Pump and probe path annotated, colour coordinated with figure 16a.

Figure 16: Circuit diagram (a) of full experimental setup and image of current optical setup (b)

VI. OFFSET LOCK CIRCUIT

To achieve a generic laser lock an error signal needs to be fed into a feedback system that has control over the laser. Since the laser controller in our experiment (DLC pro) has a feedback and control system already, the error signal will be fed back into it.

For an offset lock the error signal needs to be proportional to the frequency difference of the two lasers. The difference frequency (beat frequency) is extracted by mixing[33] the lasers in free space (see figure 16a).

There are 2 main ways to process this frequency to produce an error signal: analogue, and digital processing. Due to my familiarity with digital electronics versus analogue the design for this experiment will be digital.

1. *Aside: analogue processing*

A common analogue method involves mixing the beat frequency with a phase delayed version of itself. The explicit mixing is shown below,

$$\{f_b\} \bigotimes \{f_b + \phi\} = \{f_b\} + \{f_b + \phi\} + \underbrace{\{\phi\}}_{\text{Diff}} + \underbrace{\{2f_b + \phi\}}_{\text{Sum}}, \tag{31}$$

where $\{\}$ represents a frequency/Fourier term[34], f_b is the beat frequency, and ϕ is the phase.

A low pass filter then removes all terms containing the beat frequency (note the phase difference term is static/no-frequency for a static phase change). The phase delay is proportional to the frequency of the signal being delayed, hence the resultant signal varies as the beat frequency varies.

A. Digital processing

The block/logic diagram for the digital circuit we will use can be seen in figure 17; the circuit diagram can be seen in figure 50 in appendix E.

1. *Frequency Division*

For our current setup (see figure 15) the probe beam is detuned by 7.8 GHz from the pump beam. Most digital electronics will not work at these frequencies, hence the first functional block

[33] Frequency mixing produces the sum and difference frequencies. As in equation 31.

[34] Note that ϕ is not a frequency even when represented here. Explicitly terms would be $ft + \phi$ and hence the phase mixes but is not a frequency.

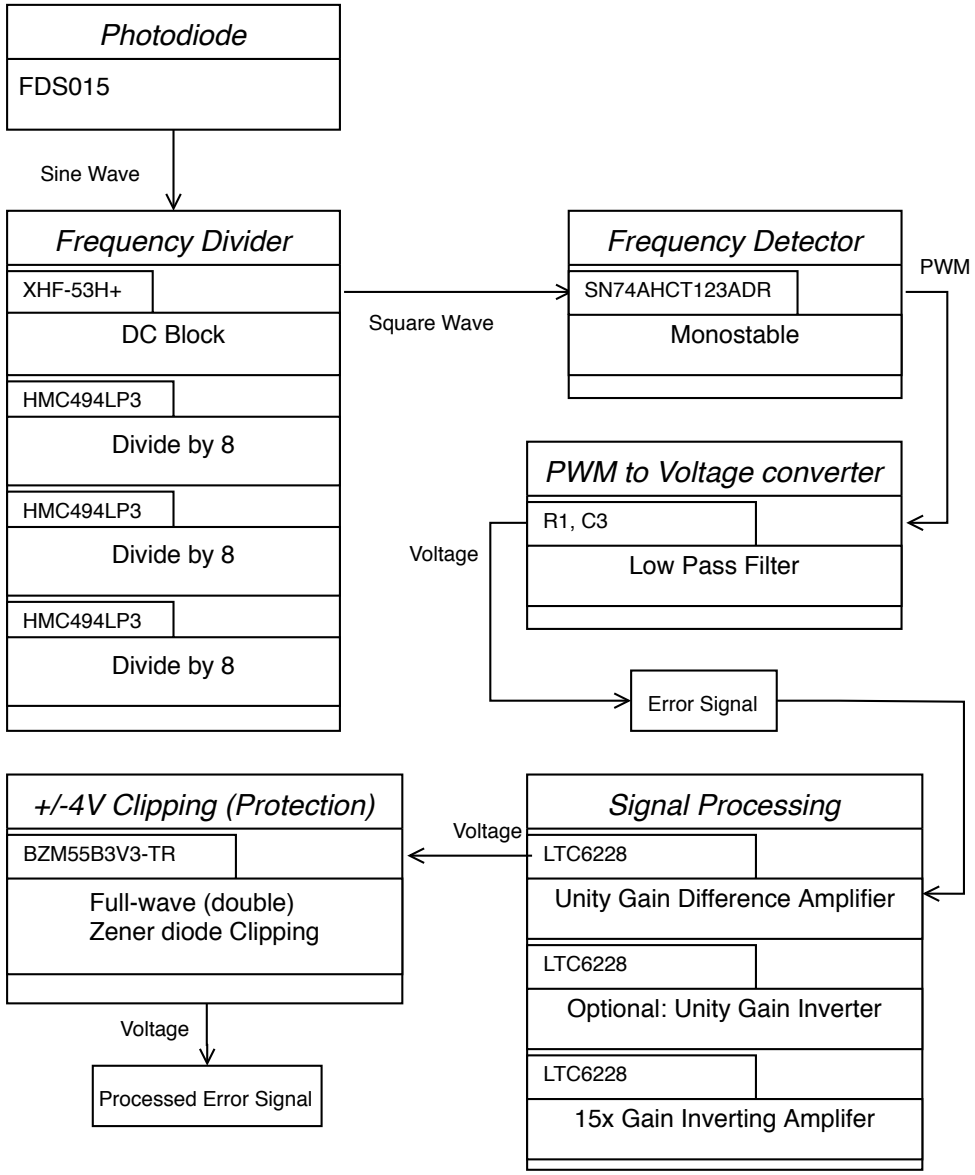

Figure 17: Flowchart showing the block diagram/logic of the offset lock circuit. Each block is entitled (italics) with its functional description. Within are folders giving more detail as to how the function is achieved. These folders are ordered, top to bottom, as in the circuit. Each folders body has a functional description of the step, and its tab has the component name in the circuit (seen in figure 50, appendix E). On the link from each function block the type of signal[a] being sent is stated to help understand the processes.

(see figure 17) involves frequency dividing the signal down to a use-able frequency ($\sim 10\,\mathrm{MHz}$).

The direct current (DC) block component seen in this functional block in figure 17 is required for the high frequency dividers (HMC494LP3) as they need to have their inputs blocked from DC signals. The DC block functional step is therefore a high pass filter.

[a]PWM stands for pulse width modulated.

At this point the signal is also at a low enough frequency that it can be read by an oscilloscope, which can be connected here via BNC. This allows for an absolute reading of the beat signal frequency whenever needed.

2. Frequency Detection

Offset locks normally use phase-locked loop (PLL) integrated circuits (IC) to directly control the laser. The PLL incorporates a feedback system that will allow locking of the laser on its own. Since our laser controller already has a sophisticated locking system the full PLL is unnecessary. To simplify this, a basic frequency detector was implemented instead. This uses a monostable[35] triggered by the input signal. The pulse width of the monostable is set up to be half the time period of the desired frequency, $f0$. If the input frequency is at $f0$ then the monostable outputs a signal at the same frequency with a duty cycle of 50% (as seen in figure 18). If the input frequency varies from $f0$ then the time period of the monostable varies to match the input[36]. However, the pulse width remains the same, hence, the duty cycle changes as the ON time is constant but the OFF time varies with the frequency of the input. An example of this, where the input frequency is below $f0$, is illustrated in figure 18.

Therefore this stage outputs a 'Pulse Width Modulated' (PWM)[37] signal where the duty cycle varies with the input frequency, centred at $f0$. It can be seen that this is sufficient for the application as this signal is an (very) unprocessed error signal, i.e. a signal that varies based off frequency deviation from $f0$.

The value of $f0$ is set by an RC timing circuit (components R4, C2 in figure 50), hence to adjust the centre frequency these components need to be changed. For our experiment this won't

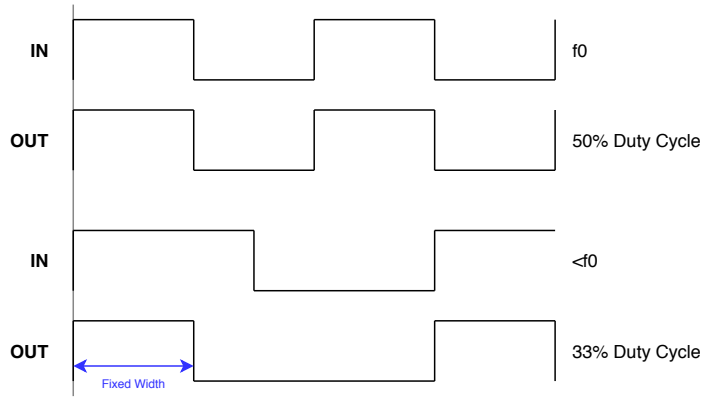

Figure 18: Timing diagram illustrating how a monostable can be used as a frequency to voltage converter. The top two signals show the input signal at the desired frequency, $f0$, and its resultant PWM signal; the bottom signals show the case for a frequency below $f0$ ($\frac{f0}{2}$).

[35] A monostable circuit produces a square pulse of set width whenever its clock pin is triggered.
[36] Providing the input time period > monostable pulse width.
[37] True PWM signals have constant frequency and varying pulse width (to interface with other digital chips on the same clock), this system is the opposite.

need to be continuously varied hence it is easy enough to re-solder new components. A variable resistor was not included as it will introduce noise into the RC circuit making the time period of the pulse less precise/more noisy.

3. PWM conversion

To use an error signal the signal needs to be a voltage level that varies with frequency not a PWM signal. Converting between the two is easy and involves using a low pass filter. The cut-off frequency of the filter is set well below the PWM frequency, such that the filter acts to remove the oscillations from the signal (as an integrator). The voltage level of the signal is the duty cycle multiplied by the supple voltage[38].

The choice of cut-off point for the low pass is a trade off. The integrator circuit has two characteristics: response time, and ripple voltage. The ripple voltage is the sinusoidal signal that survives on top of the voltage level and hence gives the precision of the system. The response time is the time it takes the signal to settle into this ripple after a step (up) voltage change. The more aggressive the cut-off point (lower frequency) the lower the ripple voltage (as it filters better), however, the response time increases as the time constant of the RC circuit increases[32].

16 kΩ and 1 nF were selected for the RC values (see figure 50). This gives a ripple voltage of 5 mV and a response time of 37 μs[32]. Since the scan speed of the laser controller is currently 11 Hz, or 90 ms, the response speed is fine since it is orders of magnitude below this rate.

By combining the 2 definitions[39] of duty cycle a relation between the frequency and output voltage can be found:

$$V_{\text{out}} = V_s RCf, \tag{32}$$

where V_s is the supply voltage.

Using the values in the circuit (32.4 Ω, 1 nF; see figure 50) the frequency for 2.5 V (half supply/centre frequency) is 7.90 GHz (after un-dividing the frequency). Using the ripple voltage this equates to a frequency resolution of 16 MHz (un-divided frequency) (0.2%). Since this detuning is only rough and the system experiences broadening higher than this (room temperatures atoms etc.) this is fine for our purposes.

[38] This can be seen by viewing the low pass as an integrator or by viewing it as averaging the voltages as it removes the oscillations.

[39] One being time on divided by total time; the other being voltage after low pass divided by supply voltage.

4. *Signal Processing*

After being converted to a voltage in the step above the signal is a use-able error signal. However, it is not useful for the laser controller. The current signal is centred at 2.5 V and can range from 0 to 5 V. The DLC pro accepts voltages from -4 to 4 V (centred 0 V). The remainder of the circuit is therefore processing the error signal for use with the DLC pro.

The op-amps in the circuit are full-swing meaning they can output a voltage all the way from the $-$V to $+$V supplied. However, they have a 1.2V input common mode voltage, meaning the input voltage can only go up to $+\text{V} - 1.2\,\text{V}$ meaning it cannot utilise a full 4 V signal off a 5 V supply. To avoid using a separate power supply for the op-amps, and having to clip the output, the signal is not amplified until the final step.

The first step is a unity gain difference amplifier. This re-centres the signal to 0 V by difference 'amplifying' it with a 2.5 V reference signal[40]. The signal now ranges from -2.5 to 2.5 V. Most of this range is useless since the frequency range is from 0 Hz to $2f0$ which is excessively broad. Hence an inverting amplifier is used to amplify the signal[41] such that only a certain bandwidth exists within the -4 to 4 V range. Using equation 32, and clipping the gained voltage to $\pm 4V$, the relationship between gain and bandwidth is given by

$$B = \frac{8\,\text{V}}{V_s RC\ G}, \tag{33}$$

where B is bandwidth (divided by 512), G is gain, and the '8 V' comes from the max and min voltages of $\pm$4V.

Larger bandwidths make it easier to find the non-clipped region on a scan however it makes the slope less steep which makes it harder for the laser controller to lock to. For the current circuit the gain is 15 which equates to a bandwidth of 1.69 GHz. The gain can be changed by replacing the op-amp feedback resistor, Rf4, in the circuit, see figure 50.

5. *$\pm 4V$ Clipping*

Since the op-amps can swing to $\pm$5 V and this could damage the laser controller a protection circuit is added to ensure the output clips to $\pm$4 V. Two 3.3 V Zeners are placed opposing each other (see figure 50) to from a full-wave zener clipper. Since a Zener diode drops 0.7 V over itself in forward bias (as with normal diodes) the combination drops 4V in either direction. Any

[40] A stable IC is used, not a potentiometer. A zener diode was considered but 2.5 V ones don't commonly exist.

[41] An optional inverter is included in case the slope direction of the error signal is important for the lock in the laser controller.

voltage exceeding this is dropped across the resistor between the op-amp and zeners; the value of which is selected to limit the current through the zeners (when max voltage is dropped over R2).

A 50 Ω is included after this to impedance match with the BNC cable.

B. Design Realisation

After design review the circuit was made as a printed circuit board (PCB). Unfortunately, the high frequency sections of the design (Frequency Divider block in figure 17) did no meet requirements for the experiment. The prototype/initial PCB design needed refining for high frequency lines: the tracks needed to be kept straight (no turns) and grounding tracks needed to run alongside them[42].

Due to time constraints and the secondary experiment starting up further design iterations on the PCB were not possible. However, the rest of the circuit (frequency detector onwards, see figure 17) was later shown to work within another application within the department.

Due to this, a commercial offset lock was bought and used for the fisrt experiment.

1. Probe detuning

During use of the commercial offset lock on the fisrt experiment, it was found that the detuning needed on the probe laser was in excess of 10 GHz, which agreed with literature[33]. At these detunings any drift in probe frequency (found to be ∼MHz over hours[43]) will have minimal impact on the atomic vapour and magnetometer signal, therefore to reduce cost and complexity for the secondary experiment the probe laser will not be locked.

[42]Design advice credit to Farzad Hayati: an electronics expert at University of Birmingham.
[43]Frequency tracked via ultra precise laser frequency spectrometer.

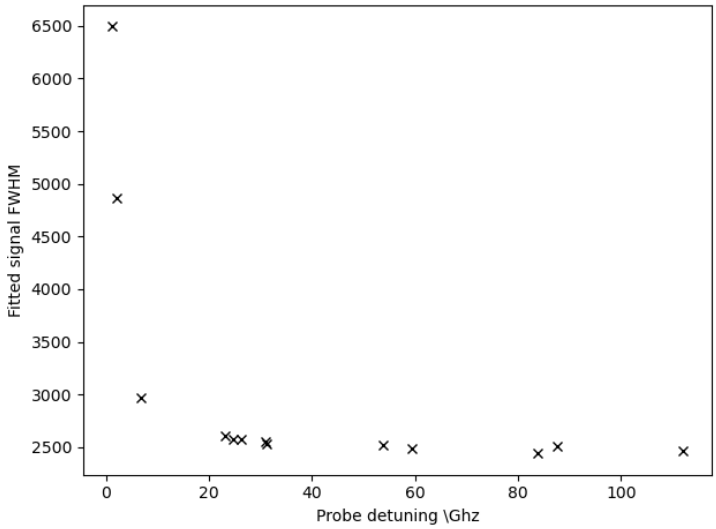

Figure 19: Figure showing the relation between probe detuning (from the pump frequency, see 15) and the full-width half-maximum (fwhm) of a magnetometer signal captured by the secondary experimental setup. Due to the small drift of the laser frequency and good fit of signal to extract fwhm (see X) no errors have been shown on the graph.

An investigation into the effect of probe detuning on magnetometer signal full-width half-maximum (fwhm) was conducted on the secondary experiment. Figure 19, above, shows that the signal improvement from probe detuning drops off after $\sim$25 GHz. As such the secondary experimental setup used a detuning of $\sim$30 GHz after this investigation to ensure the laser is in this optimal region.

VII. SECONDARY EXPERIMENT DESIGN EVOLUTION

A. Magnetic Shield

Due to time and budget restrictions an existing magnetic shield within the department was utilised for this experimental setup. This shield consists of 2 layers of 1.5 mm mu-metal that are cylindrical with the bottom end closed off with mu-metal and an open top end. A single cap piece was provided for the larger of the 2 layers that had a very large central hole for cable and laser access. This hole is 110 mm diameter which will impact shield performance overall and introduce a strong gradient field from the opening. Though not ideal, this can be counteracted by strong gradient field coils in the design.

A second cap piece was made for the inner layer to improve front shielding performance. The piece was designed in CAD and 3D printed via laser sintering within the Metallurgy and Materials department at University of Birmingham. The process uses a laser to sinter fine mu-metal powder random sections (to avoid structural weakness) until a layer is complete and it advances to the next. An image of the process can be seen below in figure 20.

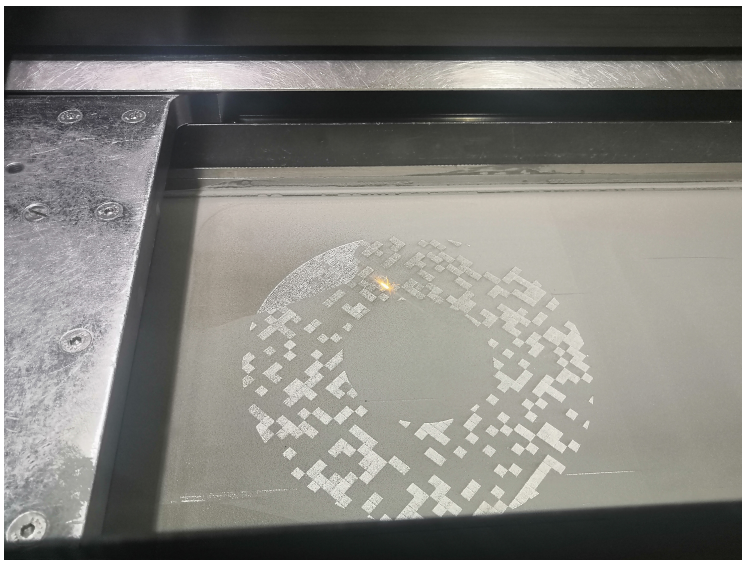

Figure 20: Image of mu-metal shield cap piece being 3D printed via a laser sintering machine.

The 2 layers were presented separate and as such mounting pieces were designed and were machined via computer numerical control (cnc) processes. Four Nylon ring sections were designed to be bolted together around the inner shield layer to act as a separator. The assembly utilised 2 separators. Nylon supports were made to mount to the optics table and support the cylindrical shield on its side. The supports were machined in 2 pieces and bolted together for strength, the shield sat on inside their joined curved edge.

Adhesive foam was added to both sides of the separators and on the support curved edge. The

foam helped to damp vibrations from the sensitive optical setup. For the separators it also helped to make contact between the layers eating up tolerance from the machining, for the supports it added friction sufficient to keep the shield in place.

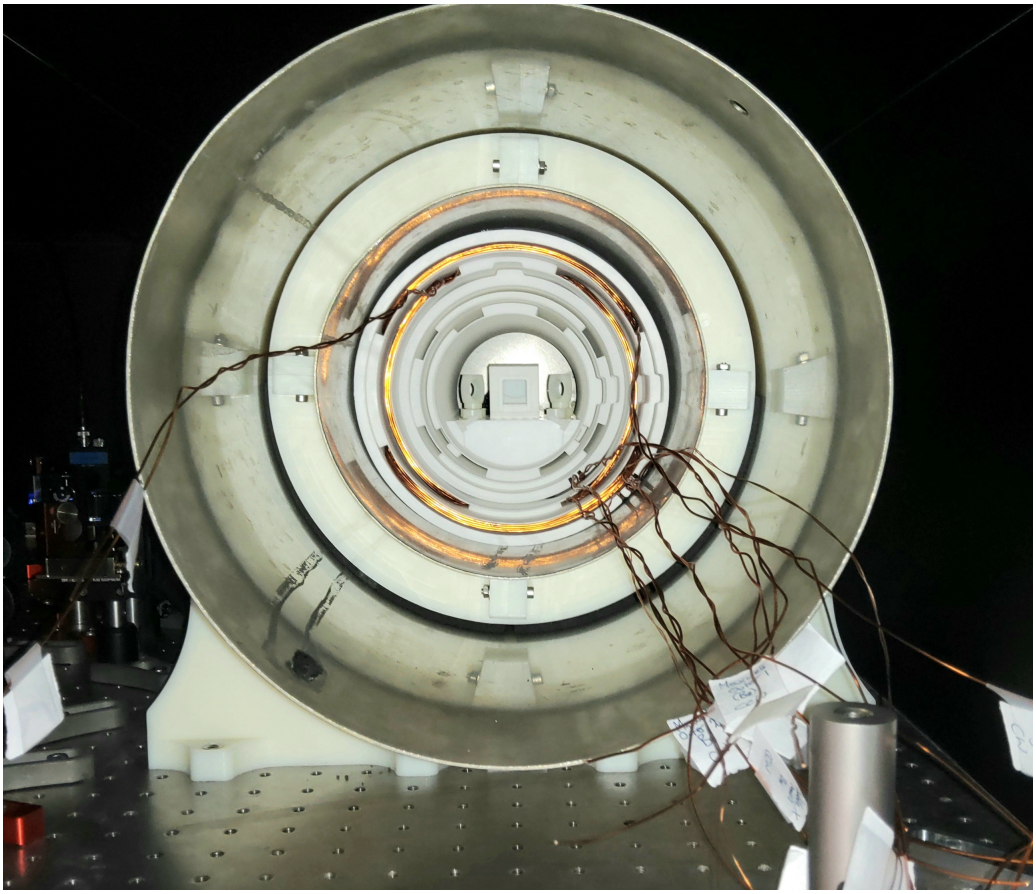

Figure 21: Image of shielding for secondary experiment. Image shows complete setup from front view showing the 2 mu-metal layers (without lids), mounting hardware (Nylon parts) with internal coil geometry (wound) and optical setup (complete).

B. Coil Assembly

The same coil geometries were used in this setup: Maxwell on main axis and Saddles on off-axes. However, to mitigate the gradient field issue presented by the shield and generally improve performance 3-axes gradient coils were added to the design. Saddle coils (with current flow opposing on each sub-coil) were used for off-axes gradient coils[44] whilst a Maxwell gradient coil[34] was used on main axis.

As seen in figure 21 the internal space of the shield was far smaller than on the first experiment. The internal diameter of the shield is 178 mm. In order to leave room to fit optics (discussed in

[44]Golay coils were not used due to size limitations.

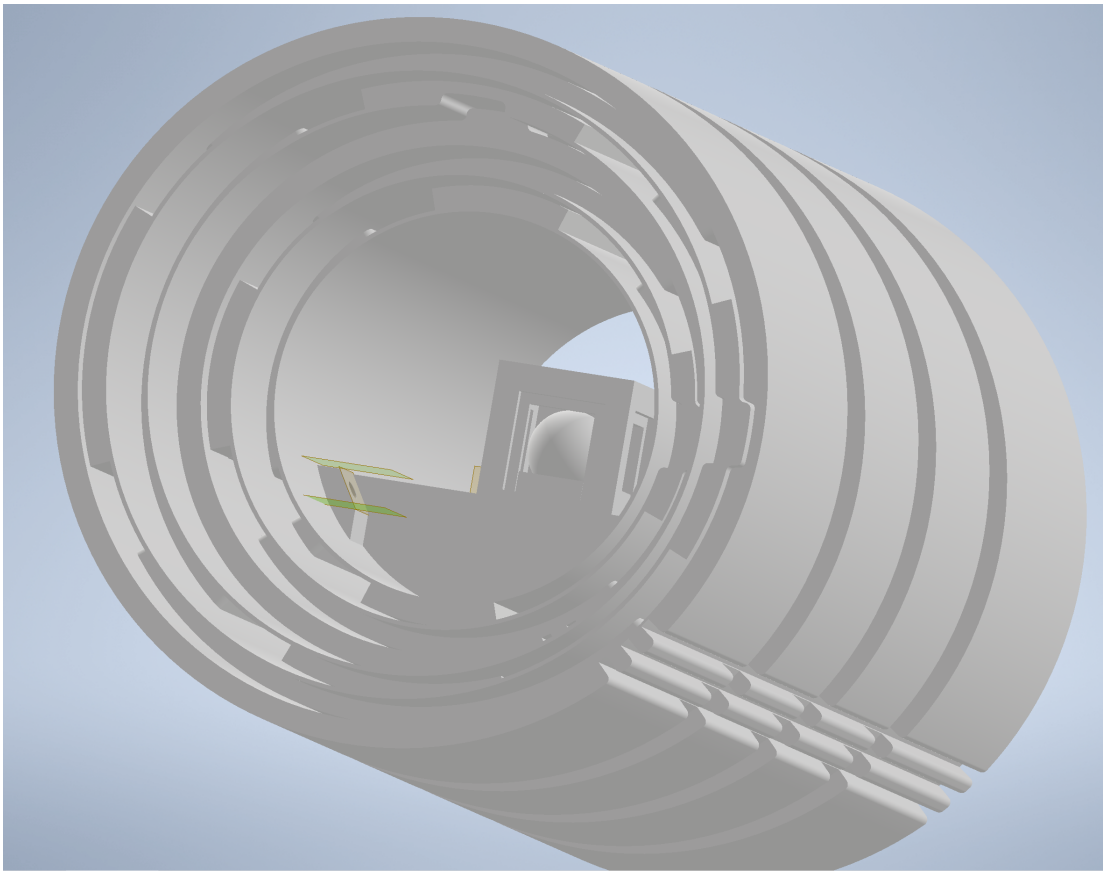

Figure 22: Image showing computer-aided design (CAD) of new coil geometry. 6 layer design can be seen with slot system for fixing layers in place. Bottom-right the outer channels to pass coil wires outside the design and shield can be seen. Inside the layers the optical board with oven design can be seen.

next sub-section) each coil groove (for winding) was made thinner than in the previous design for the first experiment and wider (to compensate). Each layer is 5.8 mm thick with a gap of 1 mm. Figure 22 shows the final design of the coil holders.

The large gap between layers was needed as these designs were small enough to allow 3D printing. 3D printing the coil assembly was far cheaper and, importantly, quicker than having them machined. The first design took over 3 months to get machined[45] whereas 3D printing took 2 weeks. Due to using 3D printing for manufacture, a large tolerance was needed as the design was likely to warp as it cooled; the 3D print company, 3DPRINTUK, advised a tolerance of 1mm as the layers needed to slide past each other. The prints worked well as an alternative to machining, though the fit was looser: in retrospect the tolerance could be lowered.

The coil assembly by layer is given below in table VII. After these 6 layers there is an internal diameter of 90mm to fit the optics.

[45]Noting there was extra complexity due to the removal of material that does not exist in this design.

Layer	Coils
1 (outer)	Maxwell centre coil, Maxwell Gradient
2	Saddle Y (vertical)
3	Saddle X
4	Maxwell outer coils
5	Saddle Gradient Y
6	Saddle Gradient X

Table VII: Table to show the coil assembly for the secondary experiment. Each layer is an equal 5.8mm thick. The main axis maxwell coil is split across layers 1 and 4 due to its geometry requirements.

C. Optical Setup

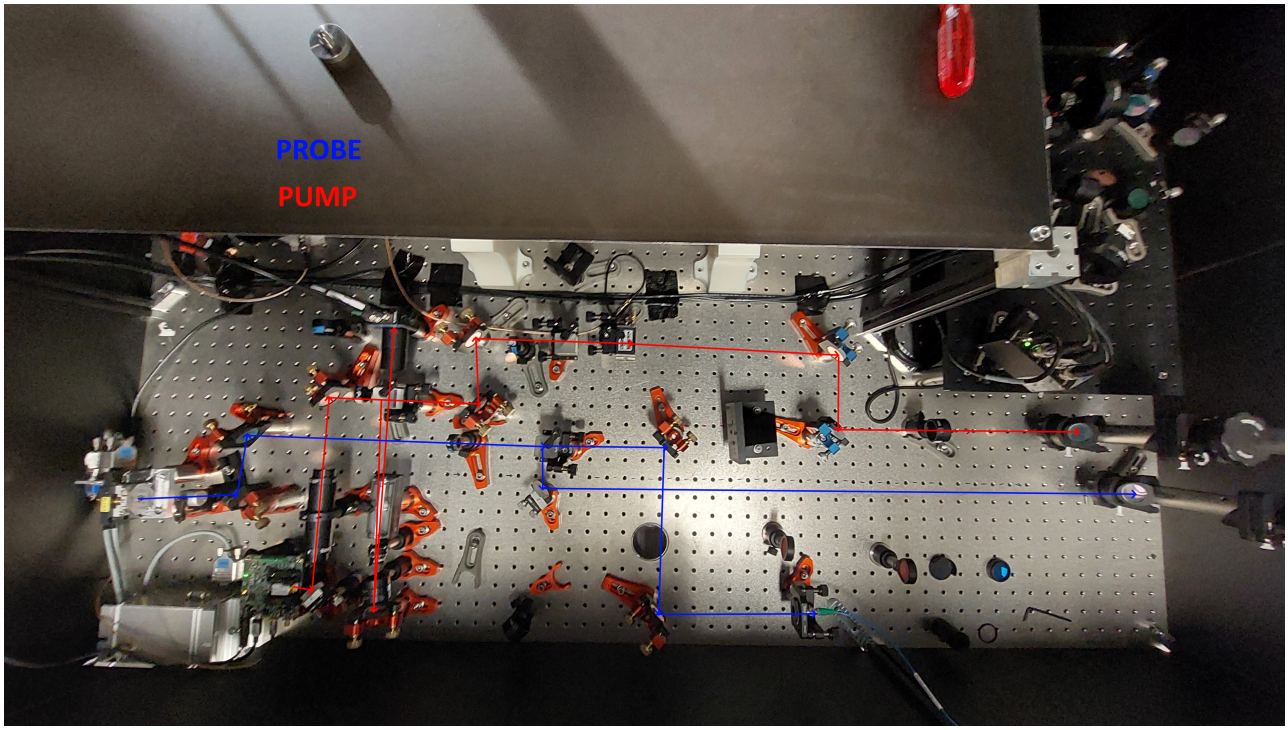

Figure 23: Image of full setup of the secondary experiment. Optical paths added for lasers up to periscopes: pump path in red, probe path in blue. See figure 24 for paths beyond periscope.

The general principles of the optical setup remain the same as the first experiment (see section V and figure 16a), figures 23, 24 show the optical setup for the secondary experiment. As noted in section VI B 1 no offset lock was used on the probe laser as it is not needed for stability, at the detunings required.

Due to the shield only having access on the front face both lasers had to enter from this side.

The polarisations of the pump and probe were kept as designed (see section V) meaning the probe beam needed to remain perpendicular to the pump beam. As mentioned earlier (section VII B), the internal diameter for optics was only 90mm creating a challenge to fit the vapour cell, oven and 2 mirrors to bounce the probe beam perpendicular in and out the cell. Figure 24 gives a view of the intended laser paths: the pump laser is fired direct through the vapour cell (oven centre: centre of figure) and is dissipated on the back of the shield interior; the probe laser is sent towards the left mirror, passes through the cell to the right mirror then back out.

Due to the size requirements the mirrors are half inch and mounted in a custom designed Nylon mirror holder (Thorlabs mounts are too large and made of metal which will affect the magnetic field generated inside). The nylon mounts can still be rotated after securing to the optics board since Nylon is self lubricating, this ended up being helpful for alignment as it was important to be able to align the mirrors once inserted into the shield and they don't have any form of adjustment themselves. The optical path required to achieve this is shown below in figure 24. As seen in the probe optical path in figure 24 the small space within the shield mapped out

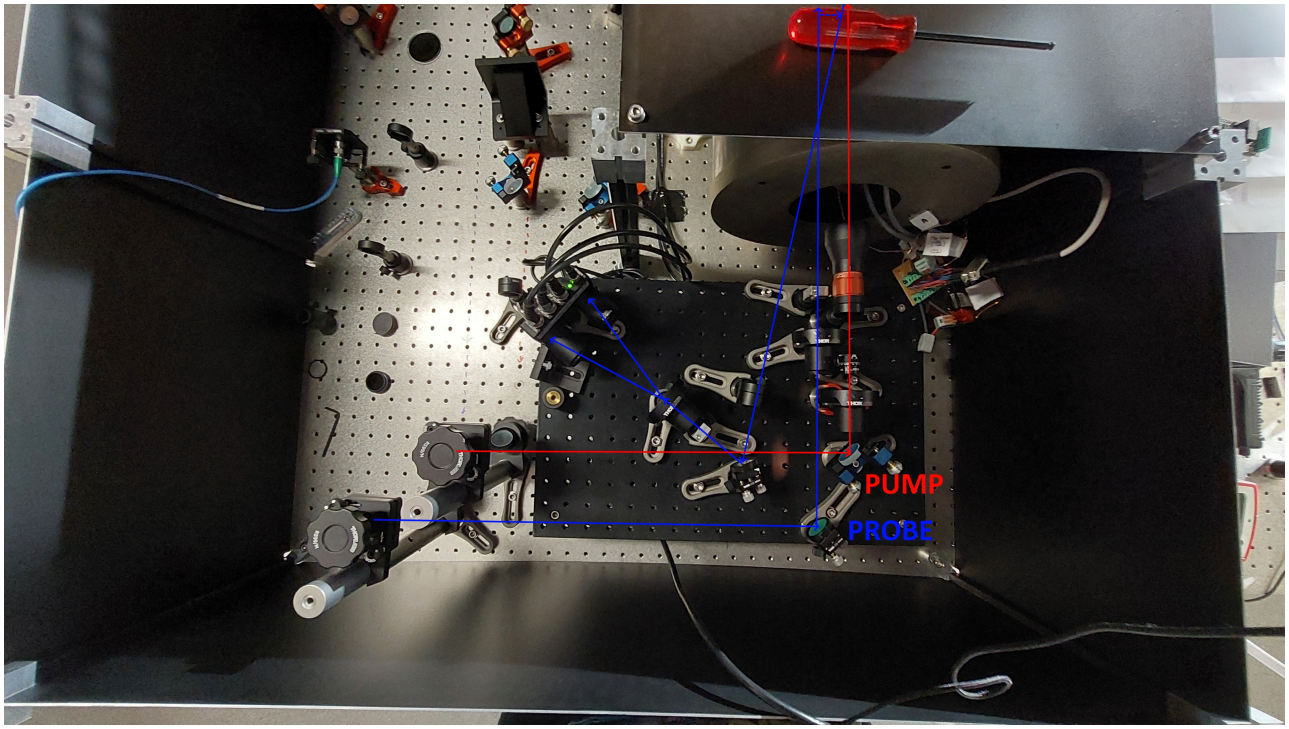

Figure 24: Image of raised optics on the secondary experiment. Optical paths added for lasers from to periscopes: pump path in red, probe path in blue.

to a small space for optical components outside the shield too (to achieve parallel beams going in). Aligning the internal mirrors was extremely difficult as the external optics partially blocked access to the shield internals and mirror alignment only involved manual rotation of the mounts.

As such the beam coming out is not at an ideal angle. Further it was found that the spherical vapour cell (as detailed in section VIII) distorted the beam shape (worsened by the imperfect spherical nature towards the stem), including expanding it, meaning not all the probe power was returned. A lens was added before the Wollaston prism to focus the expanded beams onto the balanced photodetector.

1. Probe path changes

The probe path was most significantly changed from the first experiment design (figure 16a). Due to the high detuning the doppler-free spectroscopy section was completely removed (as well as offset lock), in replacement a beam pick off is coupled into a fibre (lower path, figure 23 leading to an ultra-precise laser spectrometer located centrally in the department. This spectrometer uses a reference stabilised laser to measure the frequency of multiple other lasers to sub kHz precision. This enables far easier tracking of probe detuning and still operates in real time.

2. Pump path changes

Since the atomic vapour is at room temperature the atomic line widths are so broad the F-levels (see figure 15) overlap causing there to be no distinction if the pump beam is at the frequency for the F=2 $\mapsto$ F=3 or F=2 $\mapsto$ F=2,3 crossover. As such the frequency shift from the AOM is irrelevant, leaving its only purpose being to modulate the pump beam. Therefore, to simplify the setup, the AOM is now setup as single pass not double pass. A pinhole was added before the periscope to spatially select the first order diffration from the AOM.

VIII. VAPOUR CELL

A variety of spherical cells were used between the first and second experiments. On the second experiment the focus was on using cells with high pressure of Neon-21. Work was done in collaboration with the national physics laboratory (NPL) to research into the effects of high pressure Neon cells. NPL's work involved testing Cs-Ne cells whilst Birmingham worked with Rb-Ne.

The Birmingham cells used enriched ^{87}Rb with a varying amount of ^{21}Ne as detailed in the table below, VIII.

	Diameter/mm	^{21}**Ne pressure**/Torr
Cell 1	20	300
Cell 2a	20	500
Cell 2b	20	1400

Table VIII: Table showing the relevant specifications of the primary vapour cells used in the secondary experiment. Both cells contain ^{87}Rb enriched droplets.

The cell diameters were chosen to be small to maximise polarisation of the vapour (since atoms have to travel less distance to become polarised by the laser or other polarised atoms).

Due to the high pressure of Neon this acted as the buffer gas in the cells. This works to reduce the SDCs between the atoms and the cell wall (diffusion term, see table II). The disadvantage of this high pressure buffer method is that it causes large broadening of the atomic energy levels.

The alternative to buffer gases are anti-relaxation coatings on the cell interior. These collide elastically with the alkali atoms inside the cell preventing spin destruction during wall collision. Unfortunately, it was not possible to get a cell made with an anti-relaxation coating at the high pressures desired as the cell creation process was too technically challenging for current techniques. However, these coatings have seen common use in modern atomic magnetometry[6] as these present a major advantage over buffer gases: it eliminates the pressure broadening in buffer gas.

The standard coating used is paraffin[21]. This creates 2 new complications with its use. First is imperfection in the coating of the cell whereby it is not uniform, leading to SDCs occurring where the coating is lacking. The second limitation is that the paraffin melts at 60°C limiting the temperature the cell can be heated to. This limits the maximum vapour density of the alkali and hence limits probe signal strength.

A new type of coating overcomes the temperature problem. Octadecyltrichlorosilane (OTS)

also works as an anti-relaxation coating for potassium and rubidium and only starts to breakdown at $\sim 170°$C[6]. However, it still has the same problems from imperfect coating[46].

A. Secondary Experiment Oven Design

In order to heat the vapour cell to increase vapour density and therefore signal strength an oven was designed to house the cell. This design was challenging due to the space constraints imposed by the shield as outlined in the previous section, VII B.

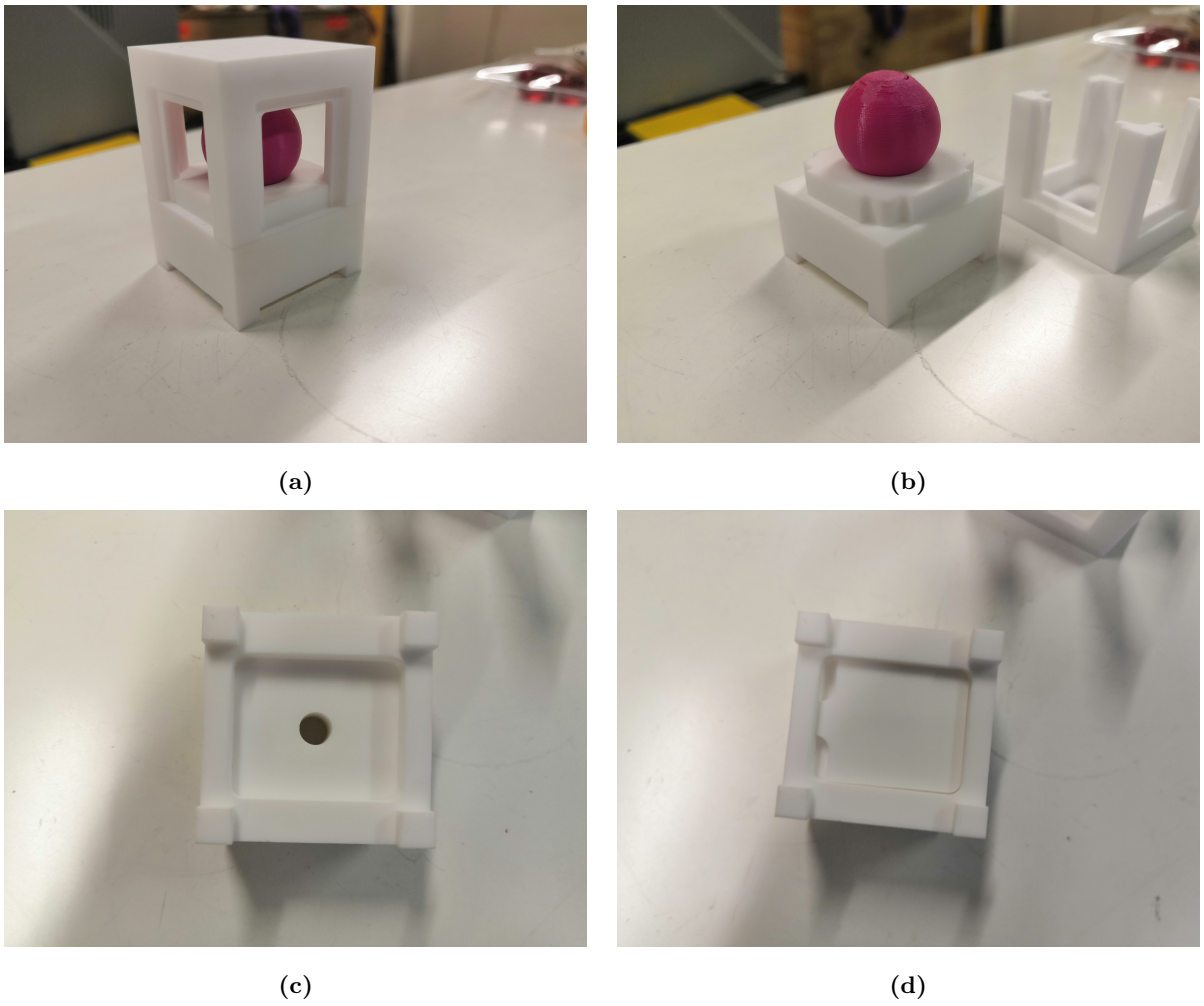

(a) **(b)**

(c) **(d)**

Figure 25: Images showing the initial oven design. Parts are machined from Macor. There is a 3D printed cell (pink) for reference. 20mm square optical glass panels not shown in images. Image (a) shows the full oven setup; image (b) shows the oven with the lid piece removed; (c) shows the bottom of the oven with the cap removed showing the hole for the cell stem; (d) shows the cover over the oven bottom, note the cut outs in the cap piece to allow cable access.

[46]These come from difficulties involved in wet chemical processes.

To meet temperature requirements whilst being non-metallic the Oven needed to be made of a ceramic, as such the machinable ceramic Macor was chosen. To fit within the shield and allow for space for 2 mirror mounts the dimensions had to be kept as small as possible, the initial design had a square footprint of 33 mm by 33 mm, which gave an internal space 24.8 mm by 24.8 mm: making cell placement and assembly delicate.

To meet these requirements the oven design involved a base with a curved section for the cell to rest securely in (see figure 25b), the cell stem was designed to run below it through the base piece to a indent underneath the base (see figure 25c) where a ceramic heating element could be added and sealed in via a ceramic cap (see figure 25d). If needed thermal paste could be added to this channel to increase heat flow to the cell. The lid piece would include the top and sides (surrounding the cell) of the oven. This piece would have 4 open sides to fit 20 mm square optical glass, this as well as all oven joins would be sealed together[47]. Since the internal space around the cell was very tight this design with the sides attached to the removable lid (see figure 25b) allowed the cell to be mounted in the oven without the sides, giving easy maneuvering during mounting, and then the lid would simply slot over the top into cut guiding grooves.

1. Final Oven Design

Figure 26: Image showing high pressure vapour cell as delivered. Note the stem is significantly larger than the 20 mm diameter sphere (designed to be ∼4 mm), also note the imperfect sphere shape causing alignment issues as noted in section VII C

Due to the high pressure of the vapour the cell stem ended up being far longer than designed (see figure 26). Instead of the expected 4 mm stem the length was closer to 40 mm. This required a major revision to the oven design. The initial design is presented above as example for other experiments needing small oven designs as most cells will not have these stem issues.

As the main oven was machined already an adapter piece was added to the design (see figure 27). This piece slots into one of the 4 windows on the oven the the stem now passes through. The

[47]This means to remove the cell is an ordeal where the sealant has to be peeled off gently.

original stem hole was sealed with sealant. The adapter piece is a simple design incorporating a base piece with a rest piece for the cell stem. The lid piece included a side access section to fit the 2 ceramic heaters ultimately used[48], and the PT100 sensor for the temperature controller (see section IX A).

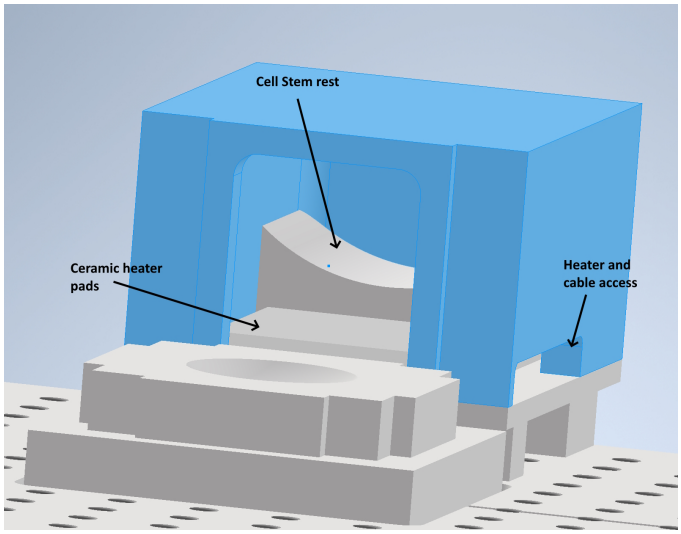

Figure 27: Image showing CAD design of oven adapter piece. Cell and oven lid removed to show adapter piece clearly. Adapter lid shown in blue. Adapter and oven shown on internal breadboard to highlight height difference in oven and adapter feet (since adapter sits on the breadboard, the oven within it).

The addition of the adapter block made oven assembly more difficult. To aid assembly 2 mounts where 3D printed to hold the optical breadboard as the oven would need to be assembled on this now. The wires for the heater units and PT100 sensor were quite rigid in the end as the heater wires needed to be shielded twin pair to reduce interference with the sensing wires. This caused a lot of torque to be applied to the oven and overcame its weight requiring the oven to be clamped whilst it was sealed, see figure 28.

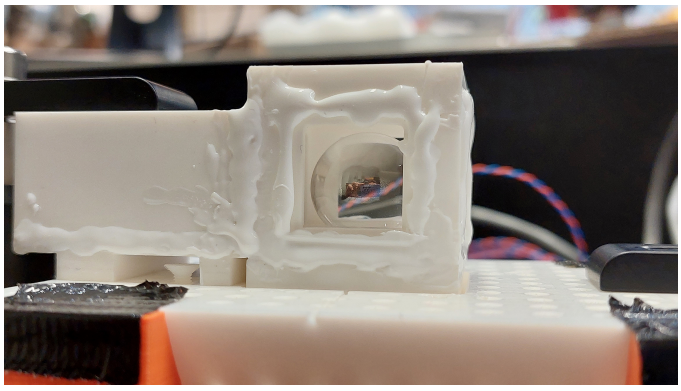

Figure 28: Image showing the oven waiting for sealant to cure after assembly with the vapour cell.

[48]Single heater performance was tested on the original oven and found to just be sufficient so a second was added to be safe.

IX. INSTRUMENTATION CONTROL

Due to limited funds for the secondary experimental setup a lot of instrumentation control required custom control hardware to be made. A general architecture was used for the control design: a Raspberry Pi Pico was used. This could be PC controlled or used via the connected 16-key keypad and 16x2 LCD (liquid crystal display), see figure 29. The Pico could control a series of instrumentation via its PWM pins which would connect to a PCB custom made for the instrumentation.

Figure 29: Image setup for a Raspberry Pi Pico controlled system. This image shows the control piece for the temperature control (see section IX A). The Pi Pico can be seen on a 3D printed mounting board alongside a 16x2 LCD display and 16-key keypad. In the background the shielded casing for the temperature sensing board can be seen.

A. Temperature Control

A custom control system was used to control the vapour cell temperature in the secondary experiment.

The temperature control circuitry consisted of 2 connected circuits: the sensing circuit and the heating circuit. Keeping the circuits separate allowed for modification of each separately and helped reduce the effect of interference from the heating circuit on the sensitive sensing circuit.

The Pi Pico was used as a control interface between these 2 circuits. The analogue output of the temperature sensing circuit was read by the Pi and then calculated the heater power needed to reach a user set target temperature. A software PID loop[49] was programmed into the Pi to handle this conversion. The software allowed setting of the PID components, setting a target temperature, or manually setting heater PWM duty cycle percentage.

[49]Form of feedback control loop utilising a proportional (P) term, integral (I) term, and derivative (D) term in the feedback.

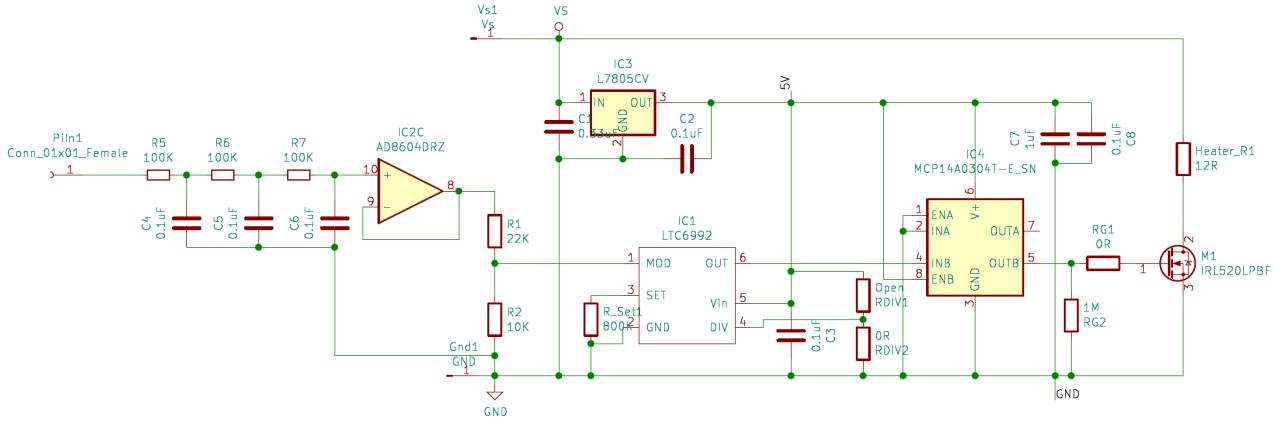

Figure 30: Figure showing the circuit diagram for the heater circuit used to control the vapour cell temperature. Pi Pico control signal comes in via 'PiIn1' on the left.

1. Heater circuit

The heater circuit needed to use AC (alternating current) heating on the vapour cell to keep the frequency above the bandwidth of the magnetometer to prevent the heating current creating a magnetic field that would interfere with the magnetometer signal.

Due to time restrictions the heater circuit design was kept simple. Inspired by robotic motor control techniques the design uses PWM to control the power supplied to the heaters. The circuit, as seen in figure 30, simply up-converts the Pi Pico PWM to a higher frequency that drives a MOSFET[50] that drives current through the heaters. The up-conversion involves converting the Pico PWM to analogue via low pass filtering and then using a potential divider to down-convert the voltage range to $[0, 1]$ V; a buffer is included after analogue conversion. The LTC6992[35] is a chip capable of driving high frequency PWM. To reduce MOSFET price the chip generates a lower frequency 62.5kHz PWM (still well above magnetometer bandwidth). The MCP14A0304T is a MOSFET driver to ensure sufficient current is supplied to drive the MOSFET.

Unfortunately, this design was seen to have a couple major issues. Firstly, the PWM signals contain a DC component[36] since the circuit is single-ended, centering the PWM about ground would alleviate this issue. However, the second issue creates a lot more problems.

Since the PWM signal is square in nature and it is being used to drive a MOSFET at high power this creates very large $\frac{dI}{dt}$ spikes which means a lot of electrical noise is created. This interfered with the sensing circuit greatly. To reduce the interference a few steps were taken. The 2 circuits were put in separate grounded boxes; the heater cables were made of shielded twisted pair, the

[50] Metal-oxide-semiconductor field-effect transducer

sensor cable from twisted pair[51]; 100 mH of inductance was added on the ground connection to the heater circuit. All these steps reduced the interference to the point of creating a temperature reading error of less than 1°C on a laboratory bench setup, this level of error was acceptable for the temperature stability needed. However, once inserted into the shield system the noise had greater impact again. This was likely because the heater and sensing cables now ran alongside each other for an extended period of time. Given the difficulty, and sensitivity (see discussion on cell oven assembly, section VIII A 1) of adjusting these cables and the time restrictions on the project initial experimentation was done without heating (but with temperature sensing) to allow testing of other crucial components.

In the future the PWM circuit should be replaced with a sinusoidal current driver to remove these noise issues. Due to the separation of the heater circuit, once designed, this will be a simple swap out.

2. *Temperature sensing circuit*

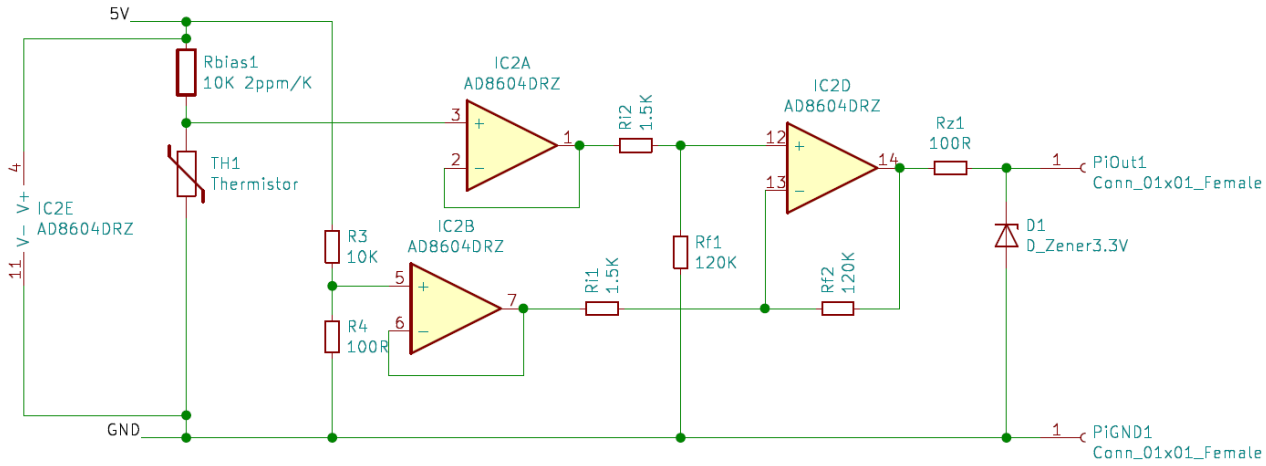

Figure 31: Figure showing the circuit diagram for the temperature sensing circuit. Output to the Pi shown on the right, 'PiOut1'; output is an analogue signal that the Pi reads. Zener diode, D1, provides over-voltage protection to the Pi Pico.

The temperature sensing circuit (see above, figure 31) is simple: it is a differential reading of the PT100 (Wheatstone bridge[37]) then a gain block (ICD2D).

[51]Later connected to coax to cross the required distance

The circuit was designed to reach a precision of 0.1°C up to 220°C[52], hence the gain block was set to approximately cap the Pi Pico voltage (3.3) at 220°C. To achieve this the resistor used with the PT100 (TH1), Rbias1, needed a very good resistance thermal coefficient[53]. The Rbias resistance should not change the analogue-digital converter (ADC) reading over the possible range of lab room temperatures (taken as a range of 10K). To calculate this limit the effect of the thermal coeficient was tracked through to Pi ADC as follows:

$$R_{\text{bias}}(\text{ppm}) = 10^3 \cdot \left(1 + \frac{\text{ppm}}{10^6} \cdot 10\right), \tag{34}$$

$$R_{\text{PT}}(\text{T}) = (1 + 3.9083 \cdot 10^{-3}\,\text{T} - 5.775 \cdot 10^{-7}\,\text{T}^2)[38], \tag{35}$$

where T is temperature and taking PT100 resistance to second order only,

$$\implies V_{in}(\text{T}, \text{ppm}) = \frac{V_s \cdot R_{\text{PT}}(\text{T})}{R_{\text{PT}}(\text{T}) + R_{\text{bias}}(\text{ppm})}, \tag{36}$$

where $V_s = 5\text{V}$ is the supply voltage,

$$\implies V_{out}(\text{T}, \text{ppm}) = G \cdot \left[V_{in}(\text{T}, \text{ppm}) - \frac{5 \cdot 100}{10^3 + 100}\right], \tag{37}$$

where, G is gain = 80, and the fraction on the right hand side (RHS) is V_0: the voltage at 0°C (generated by R3, R4, see figure 31).

Then finally:

$$\implies \text{ADC}(\text{T}, \text{ppm}) = V_{out}(\text{T}, \text{ppm}) \cdot \text{ADC}_{\text{res}}//3.3, \tag{38}$$

where '//' is the floor division operator, the Pi Pico analogue to digital converter resolution: ADC_{res} is 2^{12}, and 3.3(V) is the Pi Pin voltage.

Using this the ppm can be varied and the point at which the ADC reduces by 1 step can be recorded. At 200°C (target temperature), the ADC value for a ppm = 0 is 3664. The ADC drops to 3663 at 4.9 ppm, hence a limit of 4 ppm/K was taken. For the PCB a value of 2 ppm/K was used.

The PT100 resistance coefficients (eq^n 35) vary from device to device and as such where calibrated manually against a commercial thermocouple[54]. After doing this it was found on subsequent days that the PT100 would no longer be calibrated with the thermocouple. After an investigation it was found the issue was when the circuit was power cycled. Although the

[52]Target temperature is 200°C, this allows pad.
[53]A high value was also chosen to limit current flow to prevent the PT100 self-heating.
[54]This thermocouple only ran to 100°C, limiting the test range

voltage regulator had very good precision and drift over time qualities its repeatability was poor: it maintained its initial voltage to a precision good enough for the application, hence the PT100 could be calibrated; however, the initial voltage varied within a range too high for the application, hence after power cycling the calibration was lost as the supply voltage changed. Since the entire circuit drew very little current the voltage regulator could be replaced with a voltage reference instead[55].

B. Coil Control

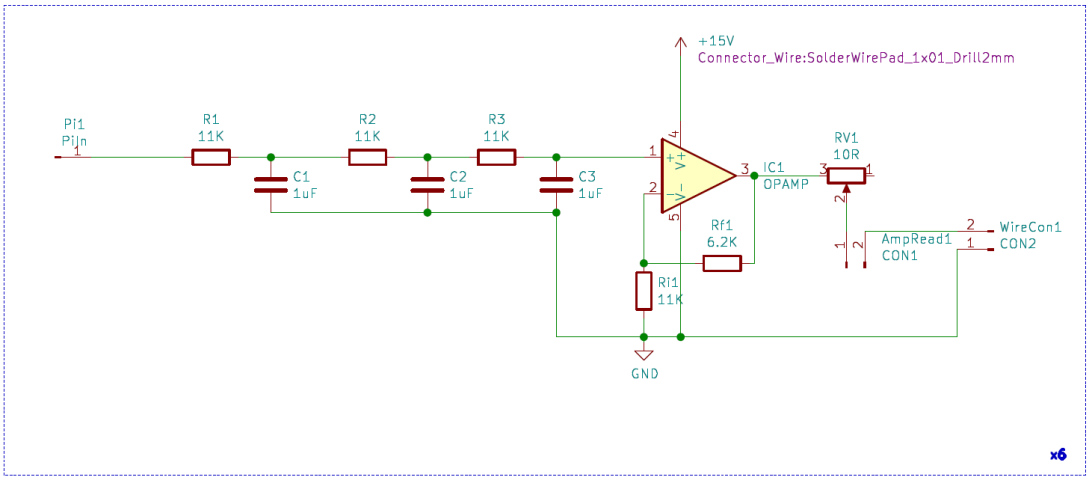

Figure 32: Figure showing the circuit diagram for the coil controller. For simplicity this is the circuit controlling an individual coil, the board is 6 of these combined. The Pi connects via the 'Pi1' connector on the left.

The coil control circuit was used to convert the control PWM signal from the Pi Pico to a current to drive the coils. The simple circuit (see figure 32) converts the PWM to analogue and then uses a high power op-amp, LM675[39], to drive the coils. A variable resistor, RV1, is used to adjust the load resistance to be the same for each coil (which have varied wire lengths) such that the same control signal from the Pi Pico results in the same current output. The Pi Pico signal to current to field strength (centre) had to be calculated manually as detailed in section X A.

The choice to use an op-amp over MOSFET was taken to avoid having to complicate the circuit to bias the MOSFET to overcome its threshold voltage. Since power op-amps existed that could drive the needed 0.5 A max[56] already they were selected.

[55] These have very high voltage accuracy and very low drift: orders of magnitude above voltage regulators, however, they are not designed to supply current.

[56] Aiming for a 100 μT field in any axes as specified as a requirement for both experiments.

In practice the full 0.5A range was not needed, the field strengths needed were far lower than specified, as such a series resistor was added into the load to reduce the current to increase the use-able range of the Pi Pico signal. This was easier than adjusting the op-amp gains and allowed for several series resistors to be tested easily.

X. SECONDARY EXPERIMENT CHARACTERISATION

A. Magnetic Field

Characterising the magnetic field and coil performance was difficult for multiple reasons. Accessing the interior with a magnetic probe was difficult due to the limit access holes in the shield; the magnetic probe was single axis and its readout was not always reliable[57]; repeating the same field strength was difficult and the coil circuit power supply varied; there was a strong field gradient due to the large opening in the shield; there were further field gradients (other axes) as the shield was old and needed to be degaussed.

1. Aside: Degaussing

Plans to degauss the shield were underway. A degaussing unit was bought, however, using it on the secondary experiment would involve disassembling most of the setup to access the inner mu-metal layer to wrap the coils around it. This time commitment was compounded by the first experiment requiring it first resulting in the decision being made there was not time to re-assemble the setup to degauss. Instead other components continued to be tested such that the state of every component of the experimental setup could be known.

2. Setup

To attempt to characterise the coil behaviour a single axis magnetic probe was inserted into the shield via the setup described in figure 33. This setup had some limitations. The long extend of the metal rod caused it to droop by the coil assembly centre, this limits the span of vertical and makes the readings have a fixed offset[58]. Coil geometry (despite efforts to fix in place) could be moved accidentally. The translation setup is likely not perfectly perpendicular to the shield, which over the large extent of the metal rod likely creates a slight combining of translation axes (with respect to the shield). This made efforts to measure the required extent to reach centre difficult and reaching a central position required a combination of measurements and alignment by eye against marked points on the coil geometry. Despite this repeatability of the setup without changing the probe orientation (accessing shield) was good as nothing could be moved, and the relative position was accurate within a run due to the translation stages[59] (even

[57] Probe frequently would have its reading drop (assume loose connection) without changing anything.
[58] Assuming the probe is not pushed into the coil floor.
[59] Precision of 500 μm per revolution (PT1/M[40].

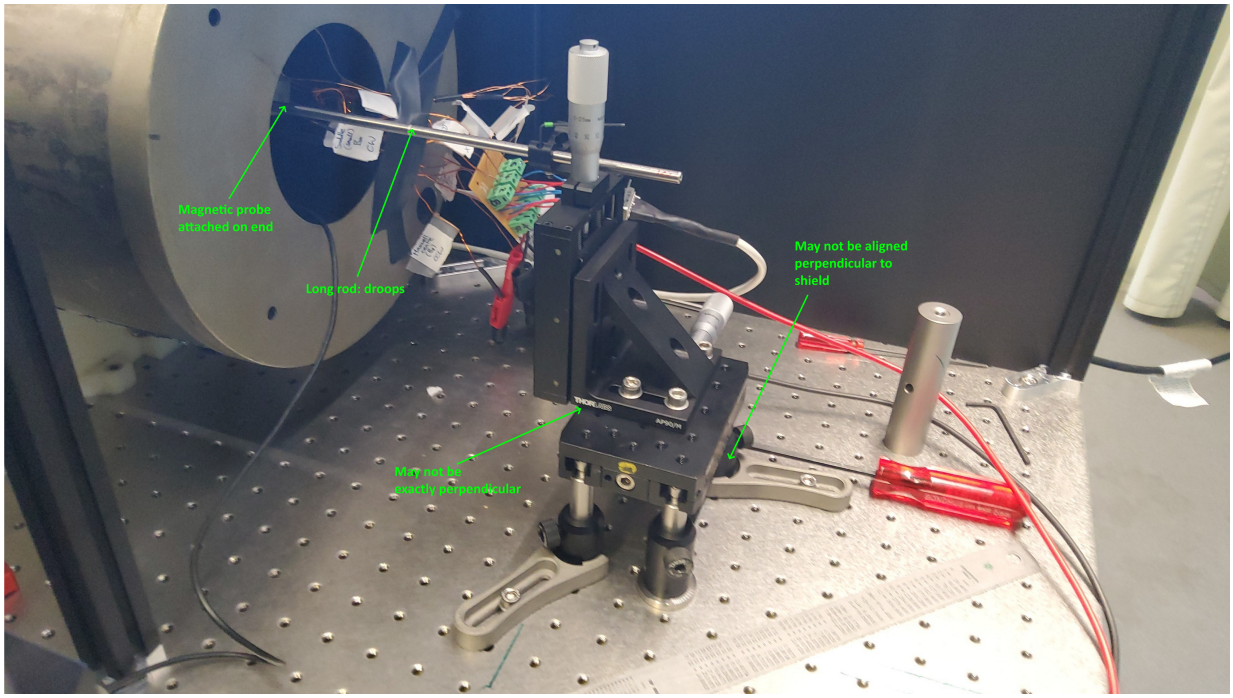

Figure 33: Image showing setup used to probe magnetic field inside shield. Setup consists of a 2-axis translation stage mounted in front of shield (optics have been removed). A metal rod is inserted into the shield with the magnetic probe attached to the end (black probe wire seen exiting on bottom), the probe is attached via a 3D printed cradle that screws onto the rod. As the probe is single axis 3 different holders exist: one for each direction of probe. In green are annotations including known issues with the setup.

if absolute accuracy is poor).

This setup was used to characterise coil current - B field relations as well as coil field shape.

3. Current - magnetic field relations

The current-magnetic field (I-B) relation for every coil was calculated, this was done by measuring the current with an ultra-precise ammeter and the field using the probe setup describe above (with the probe in a fixed position in the centre). For brevity only the Maxwell coil is reported here, though all coils gave I-B relations are similarly high precision. For gradient coils the relation was defined versus the field gradient in μT/cm. This was measured using 3 points to calculate the gradient. As seen in figure 34 the current-field relation is very well linearly fitted, as expected.

The relation between Pi Pico signal duty cycle (sig) and current (and hence B-field) was also found. However, at low currents a strange hysteresis was found in the sig-I relation as seen below in figure 35b. This was due to the ammeter switching circuits internally to deal with the

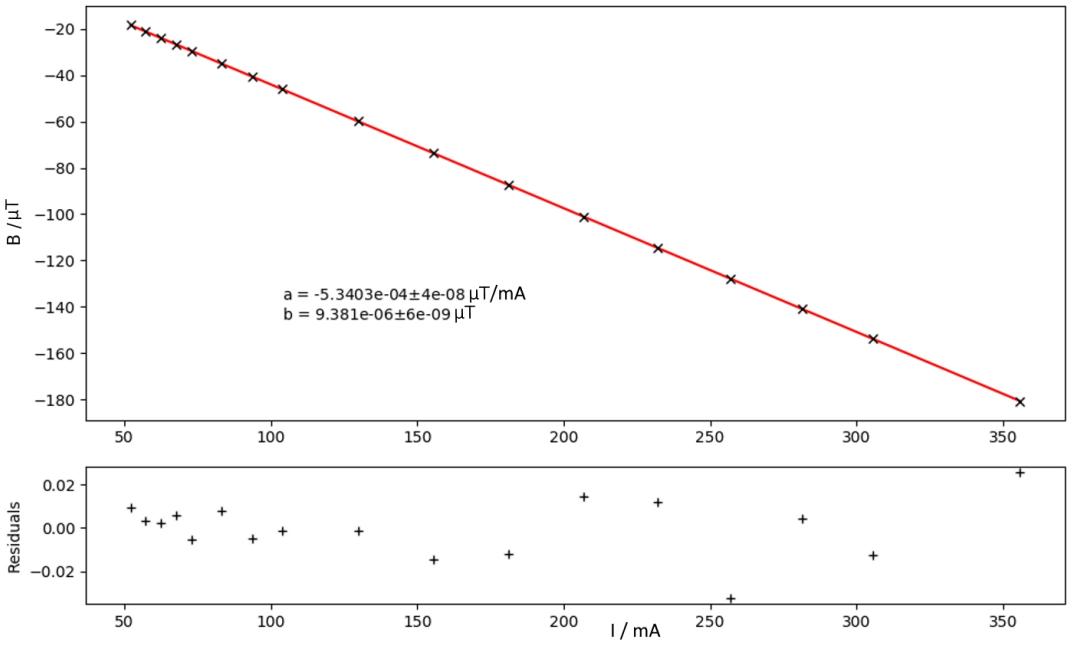

Figure 34: Figure showing the I-B relation for the Maxwell coil in the secondary experiment coil assembly.

range change. The sudden jump increase makes sense as the circuit switches to accommodate the higher power. The return response varied more, in figure 35b an example of an unclean return is shown, though some hysteresis curves were shown with the clean drop down were also seen (see figure 35a). A clean curve supports the switching ammeter hypothesis, however the unclean curve leaves questions. Without knowing the internal circuits of the ammeter it cannot be known if this unclean curve is some failure mode of the ammeter, though it seems unlikely.

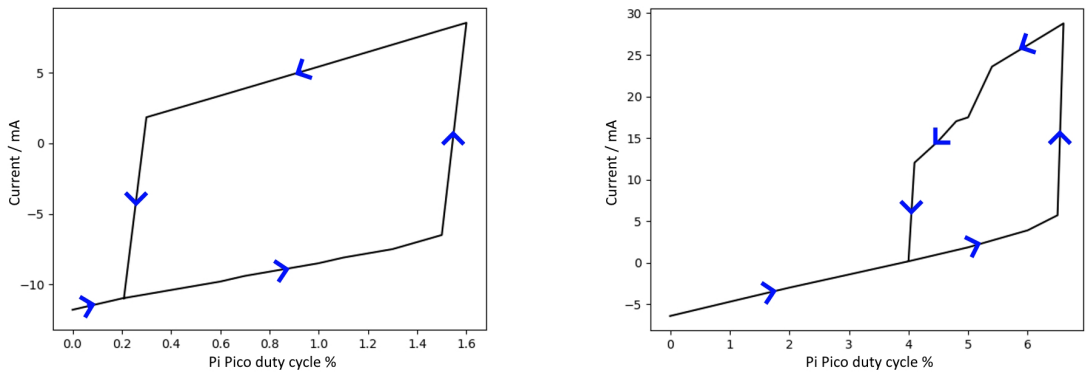

(a) **(b)**

Figure 35: Figures showing the hysteresis relation between Pi Pico PWM duty cycle and the coil current at low currents. (a) showing a clean curve whilst (b) showing an irregular curve.

No hysteresis was seen in the I-B relations, this indicates a possible issue lying at the coil control circuit. This could possibly be a precision issue for the Pi generating low duty cycle PWM signals. To avoid this issue, and increase the useful range of the Pi control signal, series resistors were added in the coil load (see section IX B) to reduce the current and hence shift the Pi control signal to higher duty cycles away from this problem range.

4. Field Shape

Getting an accurate field shape from a coil was difficult. This was due to the probe setup limits discussed previously as well as the gradient fields present. However, by utilising the assembly gradient coils these fields could be counteracted to allow shape profiling of the coils.

Utilising the Maxwell gradient coil the Maxwell coil was shape profiled along the main axis (cylindrical main axis), see figure 36.

The function for fitting was derived from the function for a single coil[41]:

$$B_{\text{coil}}(z, n, R, I) = \frac{\mu_0}{4\pi} \cdot \frac{2\pi R^2 \cdot nI}{\left(z^2 + R^2\right)^{\frac{3}{2}}}, \tag{39}$$

where R is the coil radius, z is the displacement from coil (motion constrained to axis through coil face), I is the current through a wire and n is the number of turns.

Then a Maxwell coil is simply a specific sum of 3 coils:

$$\begin{aligned} B_{\text{Maxwell}}(z, n_c, R_c, I, dz) = \quad & B_{\text{coil}}(z, n_c, R_c, I) \\ & + B_{\text{coil}}(z + dz, n_o, R_o, I) \\ & + B_{\text{coil}}(z - dz, n_o, R_o, I), \end{aligned} \tag{40}$$

where n_c is strictly the number of turns in the centre coil, and n_o for the outer 2 coils, similarly for R_c and R_o, dz is the separation between the centre coil and the outer coils (along axis through coil face, z). The relations between centre and outer coils are:

$$n_o = \frac{49}{64} n_c, \tag{41}$$

$$R_o = \sqrt{\frac{4}{7}} R_c. \tag{42}$$

To create an equation to fit a horizontal and vertical offset, dz_c, and dB are added respectively; to reduce the number of parameters for easier fitting n_c is fixed to the known value (33):

$$B_{\text{Maxwell}}(z, R_c, I, dz, dz_c, dB) = B_{\text{Maxwell}}(z - dz_c, n_c, R_c, I, dz) + dB. \tag{43}$$

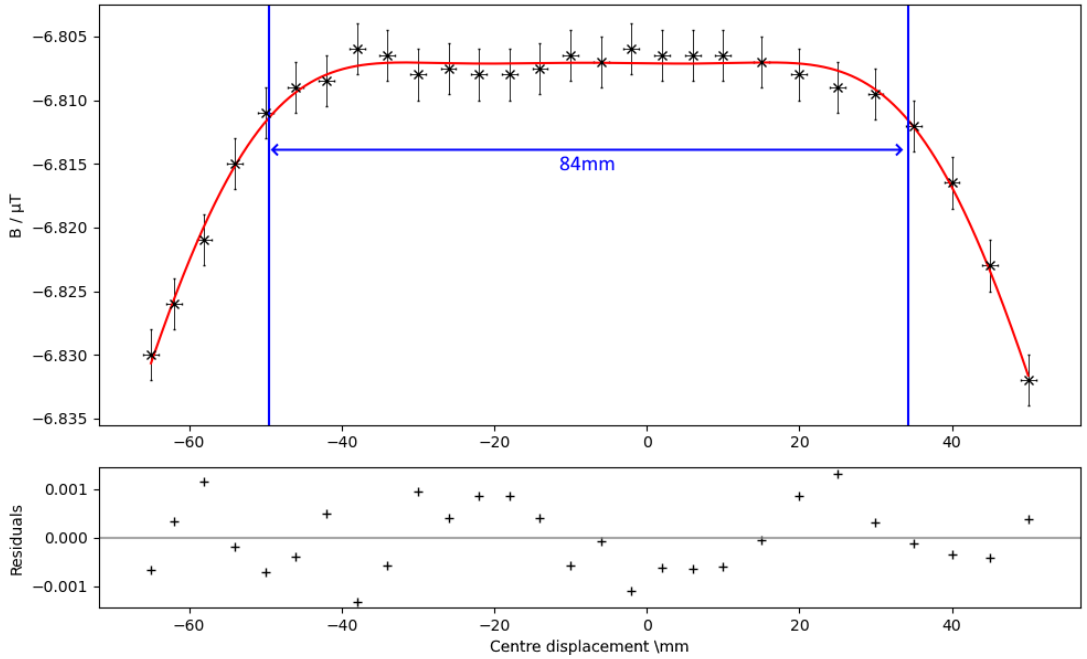

Figure 36: Figure showing the gradient compensated shape profile of the Maxwell coil (black) with fitted values (red). The z axis is through the coil assembly face; z increases as one heads further into the shield. Blue lines indicate the uniform region, please see section X A 4 a for details.

Parameter	Fitted Value	Real value
R_c	6(1)cm	8cm
I	0.2(1)mA	1.46mA
dz	4.7(4)cm	4.5cm
dz_c	-7.8(2)mm	0mm
dB	-6.94(5)μT	-7.04(3)μT[a]

[a] Measured with all coils off.

Table IX: Fit parameters for Maxwell shape profiling shown in figure 36. Real life values shown too.

As seen above in figure 36, with the gradient field corrected (dB_z coil at $\sim 30\mu A$), this fits extremely well. The fit parameters are detailed in table IX below:

R_c, and dz are the main physical parameters being fitted and match up with reality remarkably well, especially dz, with R_c being within 2σ.

The centre offset is quite large, some of this may be a result of a poorly defined centre coordinate in the data due to issues with the characterisation setup detailed previously (see section 33), however it seems unlikely to account for it all. The centre of the coil may not line up with the centre of the assembly since during winding the coils (for Maxwell) the winds did not finish an entire layer (part of the layer was unfilled) and the filled side was towards the forward

facing (negative z) side of assembly. This will shift the the magnetic field in this direction and would likely account for a lot of the fitted dz_c.

The current value is the largest discrepancy otherwise. A possible reason for this could be because the field is within a mu-metal shield, the formulas work for coils in vacuum (/air). Within the Mu-Metal shield the return path of the magnetic field is shortened through the Mu-Metal due to its permeability. This results in an effective trapping of the field lines within the shield and hence amplify the fields strength.

Interestingly the dBz term is very close to reality too. The 'real value' was not measured but comes from the I-B relations defined before (see section X A 3) for 0 current. This does not take into account any offset created by the dBz coil. For an order of magnitude calculation if you take the applied gradient field to be approximately the magnitude of the DC field it is applying at the centre of the system then the 'real value' dB would be -6.86 μT. The fact that the fitted dB value is within this limit (in [-6.86, -7.04] μT) and very close to begin with is a good sign it is accurate.

Overall the fit performs very well and lines up very well with physical values (excluding current) which was an expectantly good result, especially given all the troubles fitting the function to the data and gradient compensating the coil.

a. Discussion

To determine whether the physical Maxwell coil is performant it will be compared to simulated Maxwell coils. No simulations were performed specifically for the secondary experiment coil assembly as it was effectively a shrunk down version of the first experiment's coil assembly (with gradient coils added) and as such expected to perform similarly. Hence, it shall be compared to simulations of the first experiment's coil assembly.

The 1% region actually encompasses all of the data seen in figure 36. This is an unfair metric since the simulations were ran at absurd magnitudes (see figure 47). As such a comparison based on the shape will be performed. By taking the derivative of the fitted function the bounds will be where the magnitude of the derivative increases beyond the standard deviation of the derivative. These are marked in blue on figure 36.

This gives a value of an 84 mm uniform region along the z-axis. This value slightly exceeds the predicted value in table V. Given that the Maxwell coil is smaller than that used in the simulations it would be expected to have a smaller uniform region. Although the comparison metric is not exact due to the field magnitudes in simulation and the experiment it seems highly unlikely that

could explain the increase versus the expected decrease in region size. It is possible that enclosing the field in mu-metal is acting to focus the field and hence improving field uniformity (similar to suggested above for current discrepancy in fit, section X A 4). In would be interesting to see future work test this hypothesis and if true to calculate to what degree field uniformity and/or strength is improved within a field insulating environment. Unfortunately, this was outside the scope of this work.

B. Magnetometer signal

The magnetometer signal from the balanced photodetector is processed through a lock-in amplifier to isolate the signal pump amplitude modulation (driving) frequency. The pump amplitude modulation frequency is swept across a frequency ramp with a range such that a spectrum response for the magnetometer can be plotted. This response is expected to be a Lorentzian peak centered at the Larmor frequency (see equation 6). The lock-in amplifier used only had X/Y outputting and not the automatic phase removed R mode of modern amplifiers. This created a problem when fitting the signal to characterise it as detailed below.

1. Aside: Lock-In Amplifier Theory

Lock-in amplifiers are effectively a mixer and a low pass. The mixer combines the signal with a reference to create the sum and difference signals[60]. When the signal is at the reference signal the difference frequency will be 0 and is isolated from the sum frequency via a strict low pass filter.

There is a trade off with lock-ins resulting from the low-pass: a lower cut-off creates better frequency selection however it increases the time constant of the filter and therefore the responsiveness of the lock-in (i.e. you have to sweep the reference frequency slower).

By using the trigonometric multiplication identity it can be seen the output of a lock-in amplifier is[42],

$$V = \frac{1}{2} V_{sig} V_{ref} \cos\big(\phi_{sig} - \phi_{ref}\big), \tag{44}$$

where V_{sig}, V_{ref} are the signal and reference amplitudes respectively and ϕ are their phases.

Commonly a second quadrature is taken (adding a $\frac{\pi}{2}$ phase) such that the too can be combined

[60] Trigonometric multiplication identity.

into the 'R' mode,

$$\begin{aligned} V_R = \sqrt{V_X^2 + V_Y^2} &= \sqrt{\frac{V_{sig}^2 V_{ref}^2}{4}\cos^2\big(\phi_{sig} - \phi_{ref}\big) + \frac{V_{sig}^2 V_{ref}^2}{4}\cos^2\left(\phi_{sig} - \phi_{ref} - \frac{\pi}{2}\right)} \\ &= \frac{V_{sig}V_{ref}}{\sqrt{2}}\sqrt{\cos^2\big(\phi_{sig} - \phi_{ref}\big) + \sin^2\big(\phi_{sig} - \phi_{ref}\big)} \\ &= \frac{V_{sig}V_{ref}}{\sqrt{2}}, \end{aligned} \tag{45}$$

where the phase term is removed by adding the quadratures.

The lock-in amplifier used was a single quadrature output meaning the magnitude of the signal ('R') cannot be calculated in post processing as the signals for each quadrature would need to be taken at separate times and the difference in data between takes is too large to allow this.

As such all signal data was fitted with a term accounting for the phase:

$$\mathcal{F}\left(x, x_0, \Gamma, a, \phi, dy, \frac{dy}{dx}\right) = \mathcal{L}(x, x_0, \Gamma, a) \cdot \cos\big(2\pi T(x - x_0) + \phi\big) + y_0 + \frac{dy}{dx} \cdot x, \tag{46}$$

where x_0 is the signal centre, Γ is the half-width full-maxima, a is the Lorentzian height, ϕ is the phase difference between signal and reference, y_0 is an offset for fitting, $\frac{dy}{dx}$ is a gradient term for fitting[61], T is the and the Lorentzian, $\mathcal{L}(x)$, defined as[43],

$$\mathcal{L}(x, x_0, \Gamma, a) = \frac{a}{1 + \left(\frac{x - x_0}{\Gamma}\right)^2}, \tag{47}$$

which is a formulation of a Lorentzian with a peak height parameter, a, noting $a = \frac{A}{\pi\Gamma}$ where A is the integral over the function[44].

The phase fit performs well, as seen in figure 37, although it is not perfect. The main regions of issue are around the 90^o and 270^o (lock-in) phases, though, this correlates with the out of phase quadrature where the profile is least similar to a Lorentzian so this is not unexpected. There are also jumps in the data around the 90^o, 180^o, and 270^o lock-in phase points with the data curving towards these points forming 's' shapes in between them. The lock-in only allows continuous phase adjustment for 90^o before a button is used to add a flat 90^o to access the next quarter-cycle (with the continuous dial needing resetting to 0). A point was included for the signal from the full dial rotation to $+90^o$ as well as the button increase (with dial at 0); as seen

[61] This term was added as fitting was poor without it. A gradient is seen from the lock-in output without a magnetometer signal input however it did not fully account for it. Further causes of the gradient were not investigated.

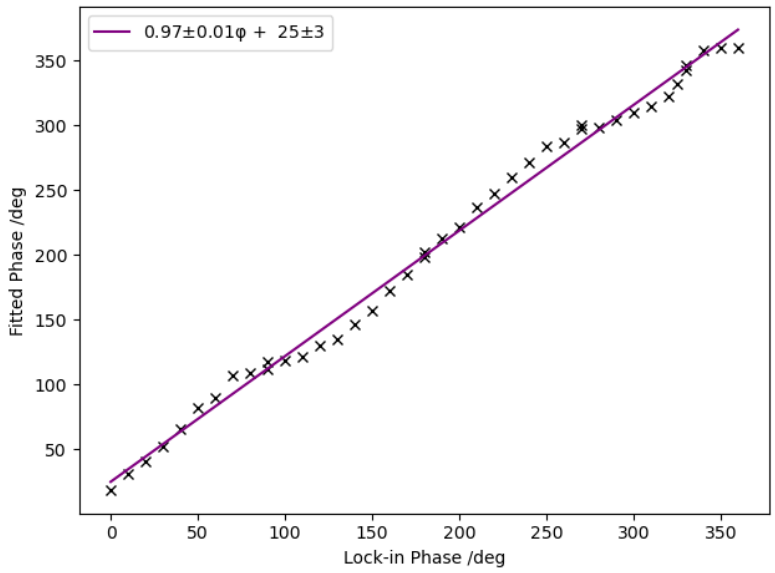

Figure 37: Figure showing plot of lock-in set phase versus fitted phase (fit equation 46. Linear fit to data shown in purple with fit parameters and error in legend.

in figure 37[62]. Given this discrepancy and the fact these regions seem to skew the surrounding points (appearing as an oscillatory behaviour) it possibly indicates the lock-in amplifier phase shift is no longer accurate/linear, especially around the quarter-cycle shift points. Continued work using this lock-in amplifier should verify if this is the case, if the oscillatory nature stems from the fitting methodology or some novel non-linear physics from the vapour cell.

The fit from figure 37 shows a near one to one relation (gradient of 1) between the lock-in set phase and the fitted phase which is a sign the phase relation and fit process is performing.

Crucially this model allows good fitting of data with a non-zero phase, as seen in figures 38 and 39. This allows the phase of the lock-in reference signal to be set such that the phase difference is 0 (determined by fit), which allows for the above out of phase region issues to be avoided for later data taking.

It importantly also allows for the correct full-width half maxima (FWHM) to be determined. The actual FWHM is of the Lorentzian term, in equation 46. By including the phase term, the signal actually widens and hence if the FWHM is taken straight from the signal (assuming the phase difference is low such that a Lorentzian fit is viable) then it is larger than the true magnetic resonance line.

As seen in figure 39, the model performs very well even in out of phase regions. Further evident in this graph is how poorly a Lorentzian only fit predicts the centre frequency[63] (and FWHM) of the signal.

[62]Double points, vertical of each other at 90, 180, 270.

[63]Note the true centre frequency is again from the Lorentzian component, hence the centre of the green line in figure 39.

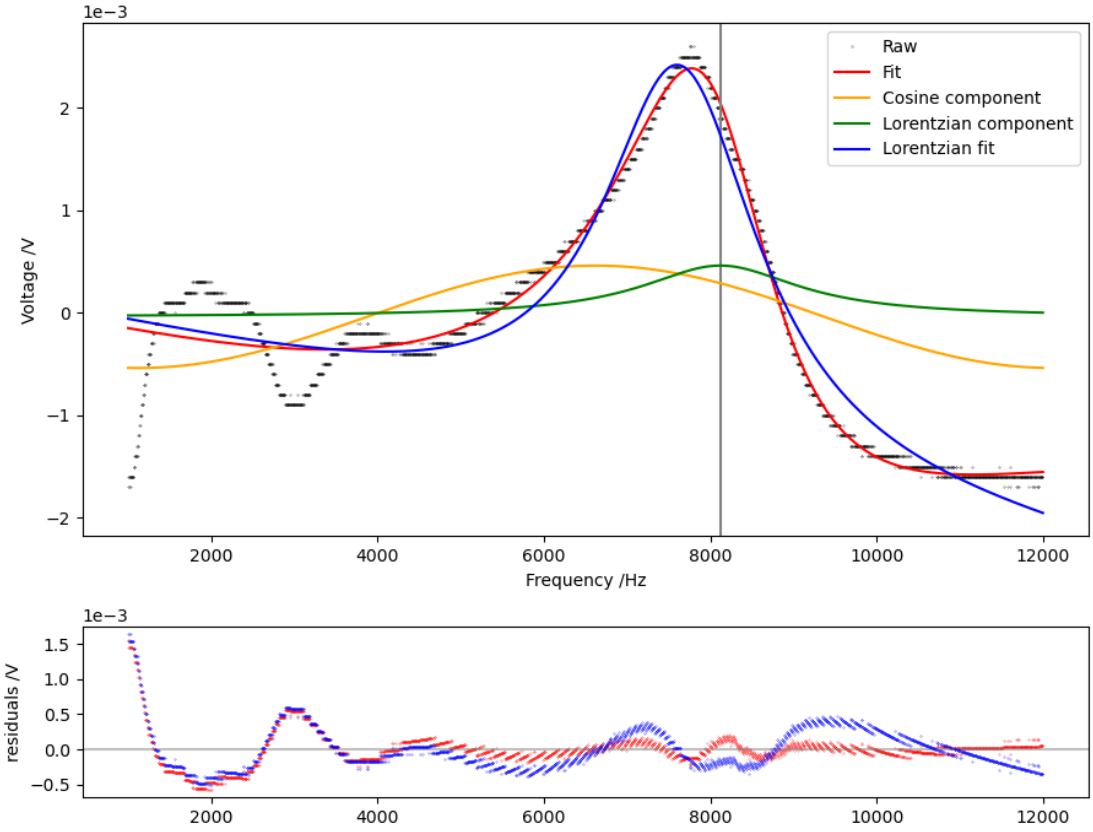

Figure 38: Plot showing the affect of the phase fitting model from equation 46. This plot is taken at a **lock-in phase of 0**^{o} which was a **fitted phase of 50**^{o}. Graph legend details line data but for further clarity, the Lorentzian fit (equation 47 with gradient term and offset) is shown for comparison with the modelled fit; the modelled fit is red, its Lorentzian component is green (normalised unity integral), and its phase component is orange.

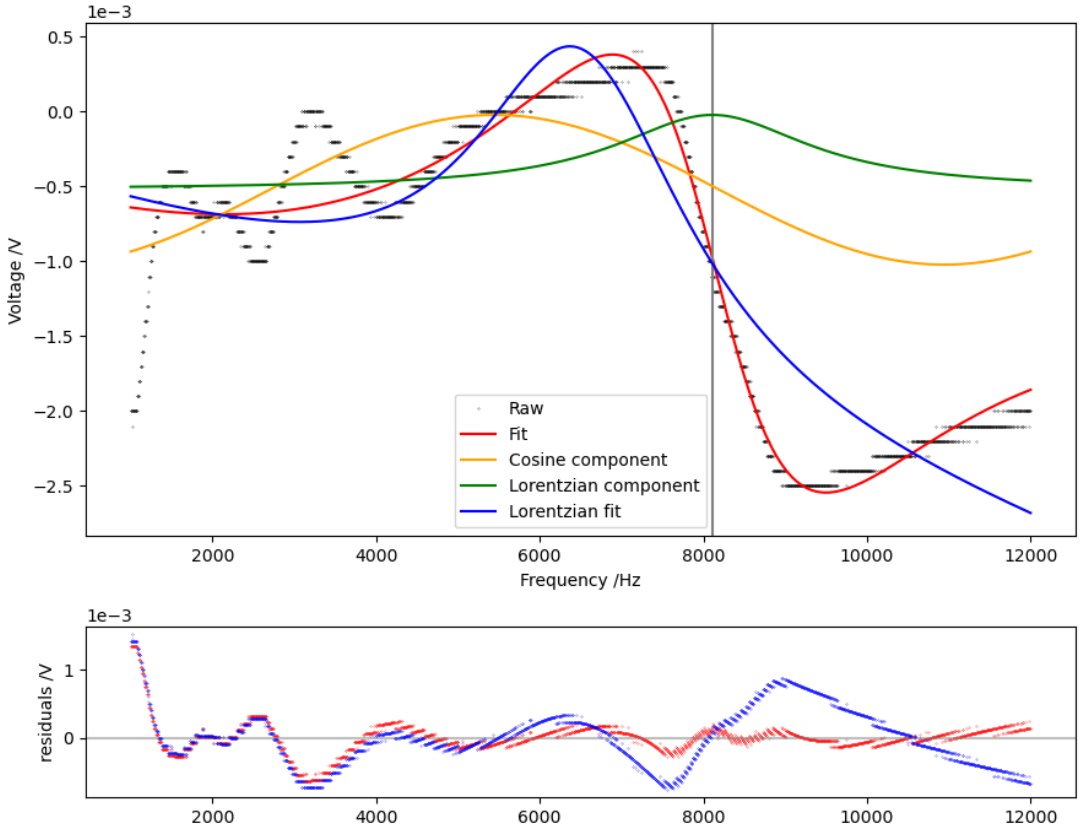

Figure 39: Plot showing the affect of the phase fitting model from equation 46. This plot is taken at a **fitted phase of 90^o**, to show the hardest fit, which came from a **lock-in phase of 0^o**. Graph legend details line data but for further clarity, the Lorentzian fit (equation 47 with gradient term and offset) is shown for comparison with the modelled fit; the modelled fit is red, its Lorentzian component is green (normalised, unity integral), and its phase component is orange.

2. *Pump Power - FWHM Relation*

With the phase difference set to ∼0 [64] the pump power was varied[65] and the FWHM of the resonance calculated from the fitted model (equation 46). The probe was detuned at 30 GHz (see section VI B 1 for reasoning).

The expected power broadening relation is given by[2],

$$\Delta\text{FWHM} = \Gamma\sqrt{1 + \frac{P}{P_{sat}}}, \tag{48}$$

where Γ is the linewidth without power broadening, and P_{sat} derives from the saturation intensity, I_{sat}[66], defined as (for a 2-level system)[2]:

$$I_{sat} = \frac{\pi}{3}\frac{hc}{\lambda^3\tau}, \tag{49}$$

where h is Planck's constant, c the speed of light in a vacuum, lambda the transition wavelength, and τ the transition lifetime.

Figure 40 shows the power relation for the experiment along with a fit following equation 48. The fitted value for P_{sat} maps to a saturation intensity of 9.5 W/m^2 which predicts a state lifetime of 46 ns, from equation 49. The lifetime of the Rb-87 D2 line is 26 ns[31]. This is far off the fitted value but it is the same order of magnitude. This indicates the fit is performing for the data.

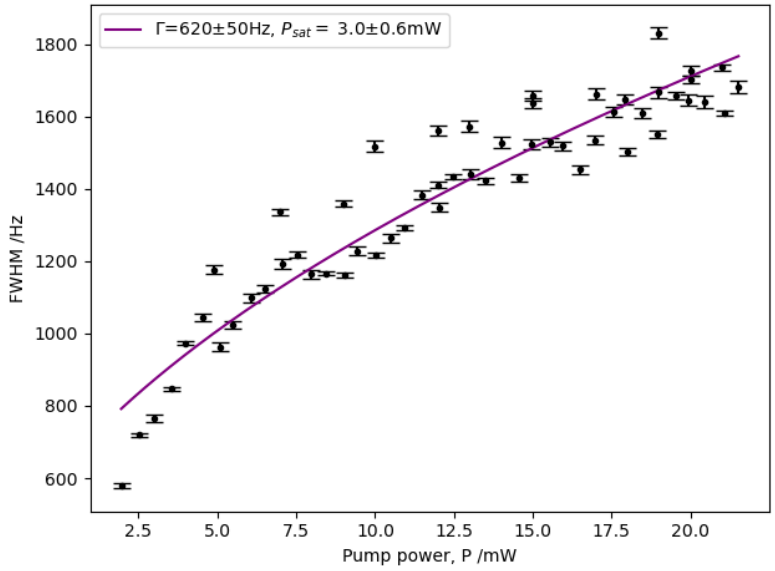

Figure 40: Plot showing the relation between the pump power and resonance FWHM. The fit uses equation 48 and its parameters are shown with error in the legend.

[64] Predicted using figure 37 and then refined and verified via test signal capture and model (equation 46) fitting.
[65] Measured via power meter before each signal measurement.
[66] I_{sat} equals P_{sat} divided by the cross-sectional area.

Parameter	Fitted Value
Γ	770(50)Hz
P_{sat}	5.2(9)mW

Table X: Fitted parameters of equation 48 applied to data shown in figure 40 with the first 5 points ommitted.

The fit performs poorly in the low power region, this is likely the result of poor laser performance (e.g. poor frequency stability) at low powers (before lasing threshold). If these points (first 5) are removed when fitting the fit improves drastically (see table X below for parameters) predicting a lifetime of 26±5 ns, which does agree with literature. This makes the minimum FWHM of 770 Hz more reliable. This value seems reasonable too since it is known the setup needs better optimising to reduce gradient fields inside the shield.

XI. CELL REGIONS AT HIGH DENSITY

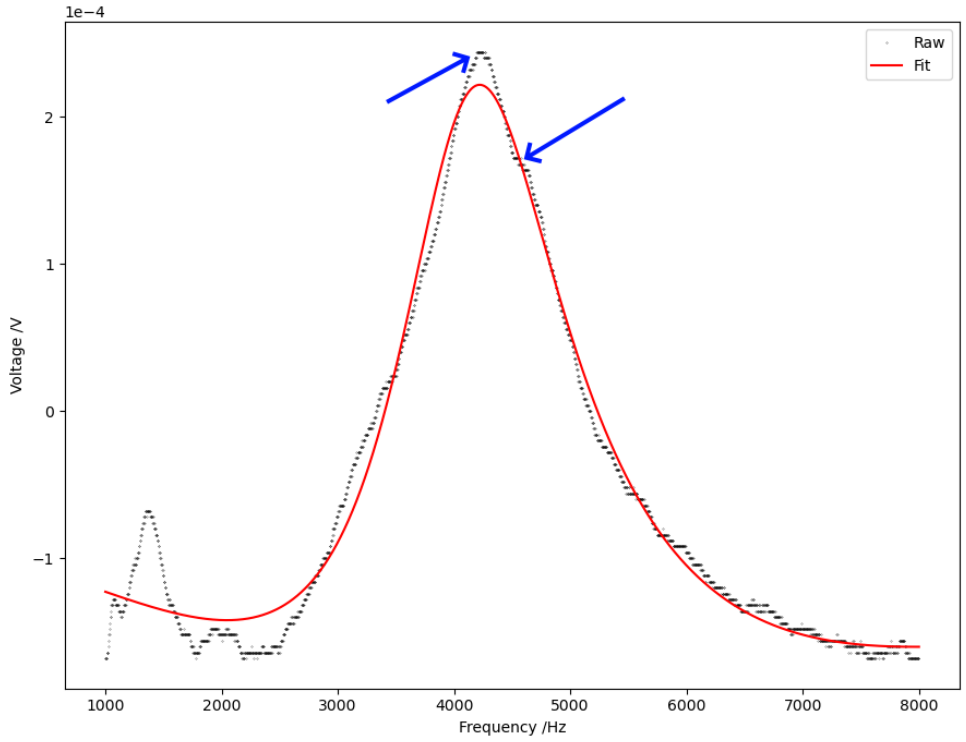

Figure 41: Plot showing a clear example of where the data presents a resonance peak with a thinner peak appearing out the top of the main resonance. Blue arrows are added to the graph to try to indicate this peak and the inflection in the resonance indicating it is superimposing on a broader peak. A fit is shown in red to show the performance of single peak fitting via equation 46.

Frequently in the signals the resonance profile appeared to have a thinner peak superimposed on it, see figure 41 for clarity. It was theorised, given this profile, that a thinner resonance was superimposing on a broader resonance.

A. Hypothesis

Due to the very high density of the vapour in the cells used, the diffusion time for the vapour is extremely high. This thesis hypothesises that the high diffusion time could lead to regions of atoms within the vapour cell that will not be able to interact with each other (via SECs) during their magnetic coherence time (diffusion time ¿ coherence time). Hence these regions could exhibit separate Larmor precession.

The difference in precession frequency could originate from imperfect field uniformity over

the vapour cell, or from the atomic vapour interaction with the vapour cell wall[67] causing a slight shift in effective magnetic field experienced by those atoms. The effective magnetic field could be altered by interaction with atomic spins within the cell wall causing an effect similar to interactions with the noble gas species within the vapour as described by the $\hat{\mathbf{I}}$ term in equation 23. No matter the cause of the resonance centre shift, interaction with the cell walls will lead to more spin destruction collisions[6] resulting in a lower coherence time. This will lead to less averaging of the magnetic states of vapour atoms in this region resulting in a broader magnetic resonance, as seen in figure 42.

B. Data analysis

Attempting to fit a double Lorentzian with a phase term leads naively to the adding an extra Lorentzian to equation 46:

$$\begin{aligned}&\mathcal{F}\left(x, x_{0a}, \Gamma_a, a_a, x_{0_b}, \Gamma_b, a_b, \phi, y_0, \frac{dy}{dx}\right)\\ =&\mathcal{L}_a(x, x_{0a}, \Gamma_a, a_a) \cdot \cos\big(2\pi T(x - x_{0a}) + \phi\big)\\ &+ \mathcal{L}_b(x, x_{0_b}, \Gamma_b, a_b) \cdot \cos\big(2\pi T(x - x_{0b}) + \phi\big)\\ &+ y_0 + \frac{dy}{dx} \cdot x,\end{aligned} \tag{50}$$

where the a, b subscripts have been added to distinguish the 2 Lorentzians and $\bar{x_0}$ is the average of the centre positions $\{x_{0a}, x_{0b}\}$ (see equation 46 for other symbols).

A single phase parameter is used to simplify the equation and since all atoms will be driven by the pump laser their phases are assumed to be similar. Even still this leads to issues quickly; as it can be seen, the parameter space is very large making fitting extremely difficult. To reduce the parameter space without loss of generality between the Lorentzians a 2 order approach was adopted for fitting.

The first order of fitting fitting a single Lorentzian as in equation 46. This is used to fix the phase parameters used in the second order fit. The second order fit is then as in equation 50 above, with the ϕ fixed and $\bar{x_0}$ set as x_0, the centre position from the first order fit.

This reduction still did not allow for reliable fitting so the offset and gradient detected from first order fitting was removed from the data before applying the second order fitting. This reduced the parameter space sufficiently to enable fitting reliably. The resulting fit equation is

[67] It is assumed that atoms near any wall will experience similar properties and therefore the peaks formed from the 2 probed regions, interacting with the cell walls (wall near laser entry, wall near exit), will superimpose and appear as the same resonance when measured. The possibility of separate cell wall interaction regions (3 peaks) is discussed later.

given below for clarity:

$$\begin{aligned}&\mathcal{F}\big(x, x_{0a}, \Gamma_a, a_a, x_{0_b}, \Gamma_b, a_b\big)\\&\quad= \mathcal{L}_a(x, x_{0a}, \Gamma_a, a_a) \cdot \cos\big(2\pi T(x - x_{0a}) + \phi\big)\\&\quad+ \mathcal{L}_b(x, x_{0_b}, \Gamma_b, a_b) \cdot \cos\big(2\pi T(x - x_{0b}) + \phi\big)\end{aligned} \tag{51}$$

where ϕ, x_0 are set from the first order fitting and the data offset and gradient were removed before (second order) fitting.

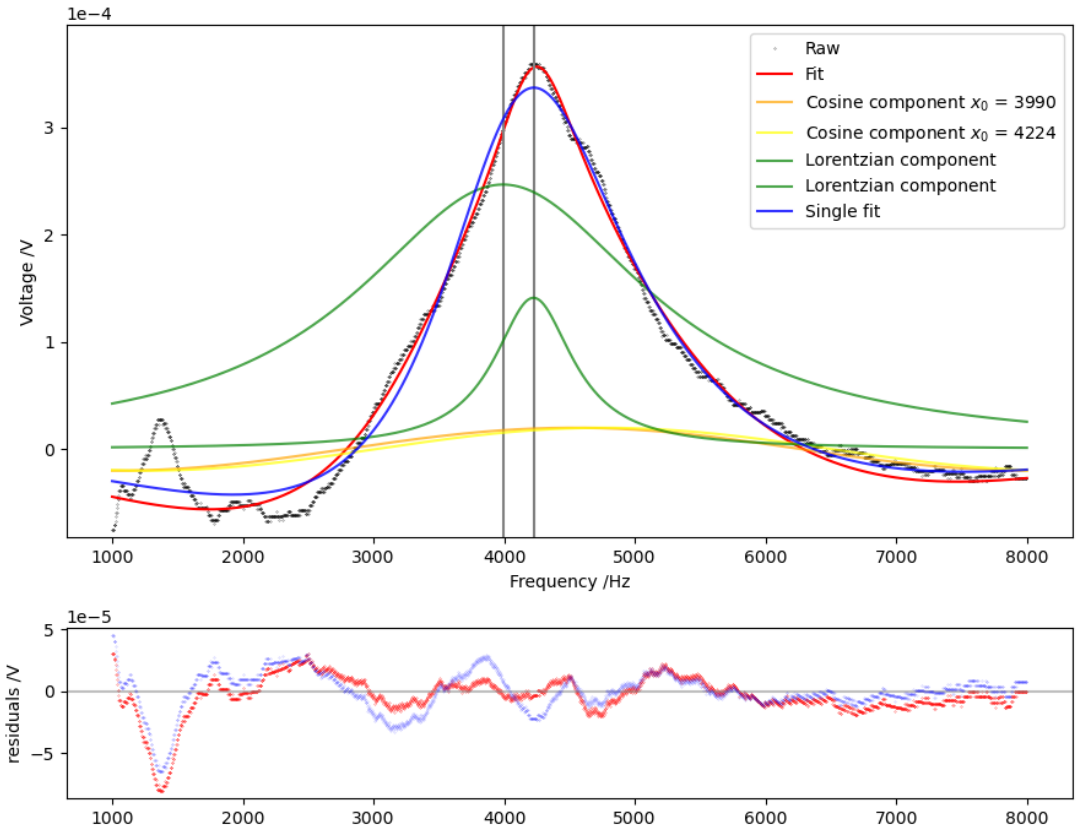

Figure 42: Plot showing the result of second order fitting of a double Lorentzian. Data is the same as in figure 41. In blue is the first order fit as in equation 46. In red is the second order fit from equation 51, see the section below for more detail on this process. The raw data is plotted with its offset and gradient removed as detailed below and required for second order fitting. In green the un-normalised Lorentzians forming the second fit (red), $\mathcal{L}_a, \mathcal{L}_b$ (see equation 51), are shown. The phase terms for the second order fit (see equation 51) are shown in orange and yellow for its respective Lorentzian as labelled in the legend (note: each phase term has the same phase but different centre).

1. Discussion

As seen in figures 41, 42 there is a clear gain in fit performance by adding the second Lorentzian. The narrower peak has been detected by the model. This has resulted in a far better match of shape profile to the data as the fit is now a lot sharper enabling it to reach the height of the data resonance. This improved fit does not trade off much accuracy at other parts of the data either, as seen on the residuals (figure 42).

An increase in fit performance is usually expected when increasing the parameter space of a model. There is a concern of over fitting the data, however at this point the data does not match up perfectly still, but it does vastly improve the shape profile fitting, hence, this level of fitting seems appropriate. This is not conclusive that sub-regions in the cell exist. A better explanation as to why the regions would form is needed first, for example there is a clear centre frequency shift too indicating that region is experiencing a different effective magnetic field.

One possible explanation for this could be that for atoms near the cell wall (as hypothesised above, section XI A) experience a weaker effective field from the noble gas species as they undergo fewer SECs with them before a SDC that destroys its state. Hence, as well as broadening the peak, it would shift the frequency as the total effective B-field the region alkali atoms experience is on average weaker. This would require the broader peak to always be found at lower frequency as well. Though this has been true through the limited data sets collected here, it is something that would need more data and to be replicated by a separate experiment before concluding anything.

One unexpected result is that the broader peak is also often larger (by area) indicating more atoms exist in this region when it would have been expected, assuming this is the near wall region, that more atoms would exist 'not near' a cell wall rather than 'near', especially given the diffusion limiting travel distance argument. Although the probe beam may likely not be perpendicular to the pump beam and/or parallel to the optical bench[68]. This could result in the beam not passing through the cell very centred leading to a higher volume of atoms 'near' the cell wall than expected if passing through the centre of the cell. Further, on some data sets these sizes are reversed, however these occur in out of phase data sets and therefore may likely be an artifact of the fit struggling in these regions. The fact this can affect these sizes though, is a sign they need to be investigated further.

[68] Due to alignment difficulties outlined in section VII C

a. Three regions

The possibility of 3 regions existing was briefly explored, this stems from the idea that, assuming the regions form based on interaction with the cell walls, that there could be a distinction between the region near the probe entry and that near the probe exit. I believe this would have to be the result of a shape imperfection in the cell to create an asymmetry between the two regions otherwise it would seem these resonances should superimpose in the signal.

Some fitting was done by adding a third Lorentzian to equation 51, utilising the same second order fitting as described there.

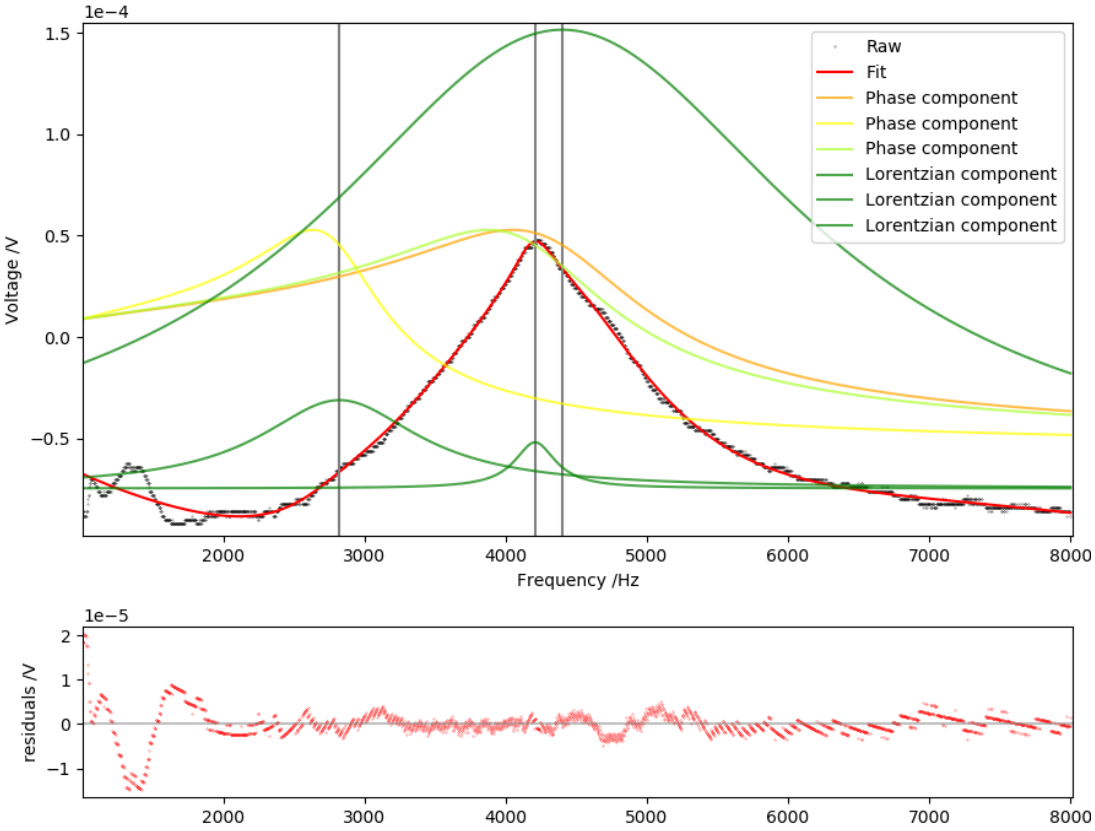

Figure 43: Plot showing the result of second order fitting of a triple Lorentzian. First order fit not shown to reduce graph clutter. In red is the second order fit from equation 51, see the section below for more detail on this process. The raw data is plotted with its offset and gradient removed as detailed below and required for second order fitting. In green the un-normalised Lorentzians forming the second fit (red), $\mathcal{L}_a, \mathcal{L}_b$, $\mathcal{L}_c$ (see equation 51), are shown. The phase terms for the second order fit (see equation 51) are shown in orange, yellow, and light yellow (centres not identified as clear to see).

Figure 43 shows an example of the triple region fitting. As seen adding a third Lorentzian causes the 3 to behave strangely. One is now massive and far is far off to the left. Further, the

fits often fail, the figure is not showing the data from figures 41, 42 as applying the process to that set creates a Lorentzian that appears as a flat line on the figure showing a fit failure.

Figure 43 shows an almost perfect fit over the resonance region (residuals on tails not considered). This combined with some fits resulting in one of the Lorentzians being set to effectively a null state by the fitter indicates this method is over fitting the data. Other colleagues continued work on the first experiment and continued fitting with 3 regions, this may perform better when the phase term does not need to be fitted as well. Their findings are not referenced in this paper as the work was ongoing at the time of writing.

XII. OVERVIEW AND OUTLOOK

Two ultra-precise co-magnetometer experiments have been setup. The practical details of design and construction have been detailed.

A compact, 6-coil, high-performance magnetic coil assembly has been designed and is scaleable to any cylindrical space. With control of magnetic fields and gradients over all 3 axes with stronger uniformity over a primary axis this design could have strong use cases within other magnetometry experiments. If more cylindrical height is available gradient Saddle coils should be replaced with Golay coils for increased performance. This design outperforms standard Helmholtz configurations commonly used due to both better usage of available volume in a cylindrical space, leading to larger coils, as well as selection of coil geometries with higher uniformity than Helmholtz coils for the same size. This design is space efficient and easy to install within a cylindrical magnetic shield, with modification to mount equipment internally being simple. Although a flexible, printed circuit board could provide a more compact coil assembly within the space it would not be possible to create a Maxwell coil with a singular board giving this design an advantage for field uniformity.

Some initial characterisation and optimisation of one of the systems has been conducted. The custom made magnetic coils were shown to work well and agree with simulations. The system linewidth, without power broadening, has been found as 770(50) Hz. Further work on the project would want to continue reducing this linewidth first before proceeding to use the experiment.

This second setup was created to investigate the impact different atomic vapour compositions (primarily high pressure) would have for applied applications (e.g. gyroscopic measurement) via the change in field precision measured. During use of high pressure vapour cells possible multiple Larmor frequencies were observed within a single vapour cell.

This thesis has presented a hypothesis for how multiple resonances can present in one vapour cell via the formation of multiple vapour regions due to diffusion times (to cross the cell) possibly exceeding the coherence time of the magnetic resonances. The location of formed regions that display differing magnetic resonances has been assumed as near the cell walls (1) and the cell centre (2). A region near the cell walls would explain why one of the resonances observed is far broader: due to more SDCs with the wall. Two possibilities for why different Larmor frequencies would form have been given. One of which is caused by poor field uniformity and future work should check this possibility first before formulating how to check for the effect of cell wall interactions.

Observation of multiple Larmor frequencies is a novel phenomena that has been observed independently within the first experimental setup by a separate, related research group. Further

work to confirm these observations is recommended. Current observations could be the result of over-fitting data or a different phenomena unrelated to needing another separate Larmor precession. This thesis encourages future work to develop a deeper theoretical understanding of this phenomena as well as testing the hypothesis presented here; this would give confidence that the observations are correct and caused by multiple Larmor precessions occurring with the vapour cell.

Testing this hypothesis would likely involve trying to alter the cell wall in a variety of ways to verify if this affects the detected Larmor resonances. This could involve using coatings on the cell walls to reduce SDCs. Coated vapour cells are difficult to produce with high vapour pressures using current techniques possibly preventing this avenue of research. Alternatively purposeful field gradients could be applied across the cell and varied to check for a change in Larmor resonances that could be predicted[69].

It is noted here that since one of the 2 resonances fitted is much broader this means the narrower Larmor resonance is far narrower than the 'single' Larmor resonance would be. It would be interesting for future work to determine if it is possible to achieve a lower linewidth by utilising the narrower resonance within a separated region within a high density cell (assuming this occurs) than using a single resonance from a lower density cell. This would lead to higher precision within applications utilising co-magnetometry, providing they can utilise high pressure vapour cells.

Determining if this multiple Larmor frequencies phenomena occurred within single species (buffer gas species not counted) high pressure vapour cells as well, without an applied gradient field[70], could give information as to the cause of the phenomena. If is not observed, then it implies the second species naturally works to shift one of the regions Larmor frequencies; if is observed, then it implies that cell wall interaction could be creating the shift.

Some discrepancy in the phase extracted from fitted resonance signals and that set by the lock-in amplifier (on the pump amplitude modulation (AM) frequency) was noted, see figure 37. This warrants further investigation in any future work to determine if this is an instrumentation issue or code issue. If it is not, then it suggests an interesting non-linear phase response from the atomic vapour.

When fitting multiple Larmor resonances it was found to perform best when fitting separate phases to each resonance (see figure 42). This implies the phase of the Larmor precession in each separated region is different. The fitted phase, showing possible non-linear response (from

[69]This would be varying the spacing in resonance centre frequency.

[70]This thesis finds no reason to think region separation would not occur in single species cells. The pertinent point being if frequency shifting of the Larmor frequency occurs without external gradients or not.

figure 37), is taken from fitting to the signal as if it were a singular resonance. The non-linearity could be the result of how the phases of the two resonances are combining within the signal. If this is the case this could provide another route to improve confidence that the two Larmor resonances observed are physical. It is noted that in order to determine separate phases for the multiple resonances functional fitting is required on the X or Y signal from the lock-in amplifier (i.e. cannot use R term or ϕ (phase) term from the lock-in amplifier). This means that the issue of it being a fitting precision/sensitivity issue about key phase points (e.g. out of phase terms, 90°, 270°) would need to be ruled out for any conclusions to be made.

[1] D. W. Gooding. Final steps to the field theory: Faraday's study of magnetic phenomena, 1845-1850. University of California Press, 1981.

[2] C. J. Foot. Atomic physics. Oxford University Press, 2012.

[3] Dr. Ramp;uuml;diger Paschotta. Faraday effect, Nov 2023. URL https://www.rp-photonics.com/faraday_effect.html.

[4] William E. Bell and Arnold L. Bloom. Optically driven spin precession. Physical Review Letters, 6 (6):280–281, 1961. doi:10.1103/physrevlett.6.280.

[5] Feb 2023. URL https://www.nist.gov/news-events/news/2014/03/magnetic-attraction-physicists-pay-homage-squid-50#:~:text=In%201987%2C%20Zimmerman%20and%20NIST,inexpensive%20and%20easily%20handled%20coolant.

[6] Dmitry Budker and D. F. J. Kimball. Optical magnetometry. Cambridge University Press, 2013.

[7] Robin F.H. Cash, Anne Weigand, Andrew Zalesky, Shan H. Siddiqi, Jonathan Downar, Paul B. Fitzgerald, and Michael D. Fox. Using brain imaging to improve spatial targeting of transcranial magnetic stimulation for depression. Biological Psychiatry, 90(10):689–700, 2021. doi:10.1016/j.biopsych.2020.05.033.

[8] Kai-Mei C. Fu, Geoffrey Z. Iwata, Arne Wickenbrock, and Dmitry Budker. Sensitive magnetometry in challenging environments. AVS Quantum Science, 2(4), 2020. doi:10.1116/5.0025186.

[9] Thomas Whitmore Kornack. A Test of CPT and Lorentz Symmetry Using a K-3He Co-magnetometer. PhD thesis, 2005.

[10] W A Terrano and M V Romalis. Comagnetometer probes of dark matter and new physics. Quantum Science and Technology, 7(1):014001, 2021. doi:10.1088/2058-9565/ac1ae0.

[11] Dmitry Budker, Peter W. Graham, Micah Ledbetter, Surjeet Rajendran, and Alexander O. Sushkov. Proposal for a cosmic axion spin precession experiment (casper). Physical Review X, 4(2), 2014. doi:10.1103/physrevx.4.021030.

[12] Interferometric measurement of resonant absorption and refractive index in rubidium. URL http://pmaweb.caltech.edu/~ph77/labs/optics/kramers.pdf.

[13] Igor Savukov. Ultra-Sensitive Optical Atomic Magnetometers and Their Applications. INTECH Open Access Publisher, 2010. URL https://www.intechopen.com/books/advances-in-optical-and-photonic-devices/ultra-sensitive-optical-atomic-magnetometers-and-their-applications#F1.

[14] jgerber (https://physics.stackexchange.com/users/128186/jgerber). π, σ - atomic transitions with respect to a quantization axis. Physics Stack Exchange. URL https://physics.stackexchange.com/q/283729. URL:https://physics.stackexchange.com/q/283729 (version: 2019-04-12).

[15] F. Bloch. Nuclear induction. Physical Review, 70(7-8):460–474, 1946. doi:10.1103/physrev.70.460.

[16] Georgios Vasilakis. Precision measurements of spin interactions with high density atomic vapors.

PhD thesis, 2011.

[17] Lev Davidovitch Landau, Lifshits Evgenii Mikhailovich, John Bradbury Sykes, and John Stewart Bell. Mechanics. Butterworth-Heinemann, 3 edition, 1976.

[18] E. A. Donley, E. Hodby, L. Hollberg, and J. Kitching. Demonstration of high-performance compact magnetic shields for chip-scale atomic devices. Review of Scientific Instruments, 78(8):083102, 2007. doi:10.1063/1.2767533.

[19] S.K. Lee and M. V. Romalis. Calculation of magnetic field noise from high-permeability magnetic shields and conducting objects with simple geometry. Journal of Applied Physics, 103(8):084904, 2008. doi:10.1063/1.2885711.

[20] N. J. Ayres, G. Ban, G. Bison, K. Bodek, V. Bondar, T. Bouillaud, B. Clement, E. Chanel, P.-J. Chiu, C. B. Crawford, and et al. The very large n2edm magnetically shielded room with an exceptional performance for fundamental physics measurements. Review of Scientific Instruments, 93(9), 2022. doi:10.1063/5.0101391.

[21] Vito Giovanni Lucivero and Morgan Mitchell. Quantum metrology with high-density atomic vapors and squeezed states of light. PhD thesis.

[22] Magnetic Shields Ltd. Mumetalproperties. URL `https://magneticshields.co.uk/technical/material-technical-data`.

[23] T. W. Kornack, S. J. Smullin, S.-K. Lee, and M. V. Romalis. A low-noise ferrite magnetic shield. Applied Physics Letters, 90(22):223501, 2007. doi:10.1063/1.2737357.

[24] 3c96 material specification. URL `https://www.ferroxcube.com/upload/media/product/file/MDS/3c96.pdf`.

[25] Helmholtz coils. URL `http://hyperphysics.phy-astr.gsu.edu/hbase/magnetic/helmholtz.html`.

[26] Maxwell coil. URL `https://homepages.abdn.ac.uk/j.s.reid/pages/Maxwell/Legacy/MaxCoil.html`.

[27] James Clerk Maxwell. articles 719-715, volume 2, page 356–359. Clarendon Press, 3 edition, 1892.

[28] Softways. Saddle coil, . URL `https://www.mr-tip.com/serv1.php?type=db1`.

[29] F. Bonetto, E. Anoardo, and M. Polello. Saddle coils for uniform static magnetic field generation in nmr experiments. Concepts in Magnetic Resonance Part B: Magnetic Resonance Engineering, 29B (1):9–19, 2006. doi:10.1002/cmr.b.20057.

[30] Maxwell coil, Aug 2020. URL `https://en.wikipedia.org/wiki/Maxwell_coil`.

[31] Daniel Steck. Rubidium 87 d line data. 01 2003.

[32] Rc low-pass filter design for pwm. URL `http://sim.okawa-denshi.jp/en/PWMtool.php`.

[33] Ning Ma, Lihong Duan, Danyue Ma, Jixi Lu, Bozheng Xing, Jin Li, and Bangcheng Han. Demonstration of a high-density alkali-metal atomic magnetometer based on the frequency-symmetrical detuning effect of two pumping lights. Optics Express, 30(25):45930, 2022. doi:10.1364/oe.474777.

[34] Softways. Mri database: Maxwell coil, . URL

https://www.mr-tip.com/serv1.php?type=db1&dbs=Maxwell+Coil.

[35] URL https://www.analog.com/media/en/technical-documentation/data-sheets/LTC6992-1-6992-2-6992-3-6992-4.pdf.

[36] URL https://liraeletronica.weebly.com/uploads/4/9/3/5/4935509/spectral_analysis_of_a_pwm_signal.pdf.

[37] Wayne Storr. Wheatstone bridge circuit and theory of operation, Aug 2022. URL https://www.electronics-tutorials.ws/blog/wheatstone-bridge.html.

[38] Konrad Scheuer Yuriy Kurtsevoy. Implementation of temperature measurements with the max22000. URL https://www.mouser.com/pdfDocs/AN7186.pdf.

[39] Lm675 power operational amplifier, Mar 2013. URL https://www.ti.com/lit/ds/symlink/lm675.pdf.

[40] Single-axis translation stage with standard micrometer. URL https://www.thorlabs.com/newgrouppage9.cfm?objectgroup_id=10211.

[41] Samuel J. Ling, William Moebs, and Jeff Sanny. 12.4 magnetic field of a current loop, Sep 2016. URL https://pressbooks.online.ucf.edu/osuniversityphysics2/chapter/magnetic-field-of-a-current-loop/.

[42] URL https://www.thinksrs.com/downloads/pdfs/applicationnotes/AboutLIAs.pdf.

[43] . URL https://mathworld.wolfram.com/LorentzianFunction.html.

[44] . URL https://docs.mantidproject.org/v6.1.0/fitting/fitfunctions/Lorentzian.html.

Appendix A: Polarisation Notation

For linear polarisation: subscript shall refer to the axis of polarisation. E.g π_z refers to light with linear polarisation propagating along any axis other than z with the axis of polarisation (orientation of electric field) being along the z axis.
Whereas for circular polarisation subscript shall refer to the direction of propagation of the light. E.g. σ_z^+ refers to light propagating along the z axis with a circular polarisation. The 'σ^+' refers to the direction of rotation of the polarisation being such that it would cause a increase in atomic m_z levels when absorbed.

Appendix B: Magnetic Shield Drawing

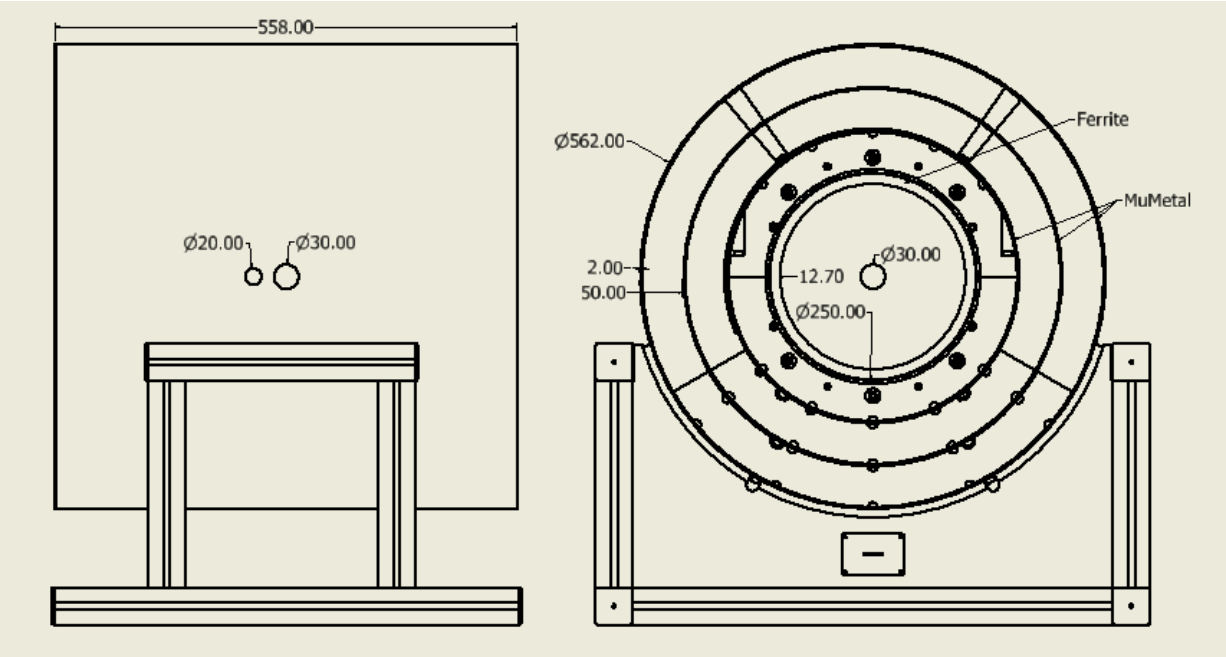

Figure 44: Drawing outlining the design of the magnetic shield as discussed in section III. End caps have been removed to see the internal layers. It should be noted that layer spacing along the central axis is also 50mm. Please note that the ferrite piece has a thick wire wound around it for degaussing purposes.

Appendix C: Coil Type Field Comparison

1. Helmholtz

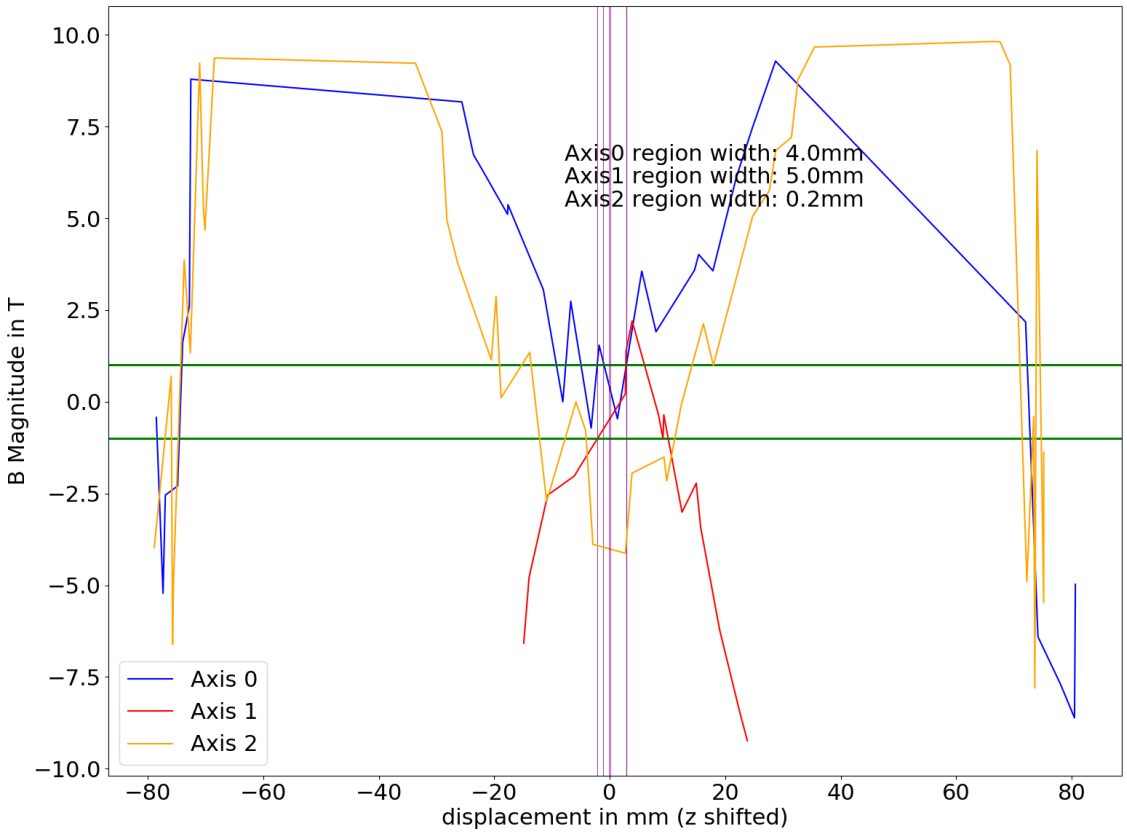

Figure 45: Graph showing the performance of a Helmholtz coil. The coil used 300A·t split between the 2 coils. It had a radius of 100mm.

2. Saddle

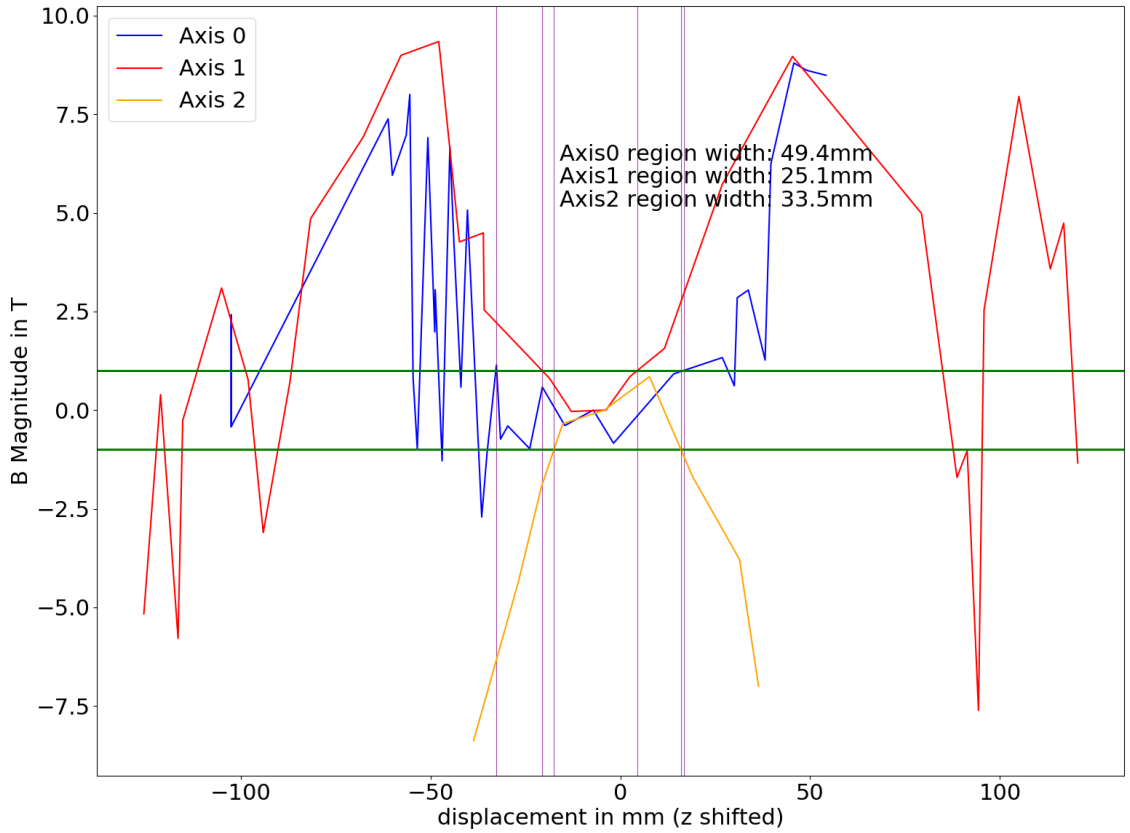

Figure 46: Graph showing the performance of a Saddle coil. The coil used 300A·t split between the 2 coils. It had a radius of 100mm.

3. Maxwell

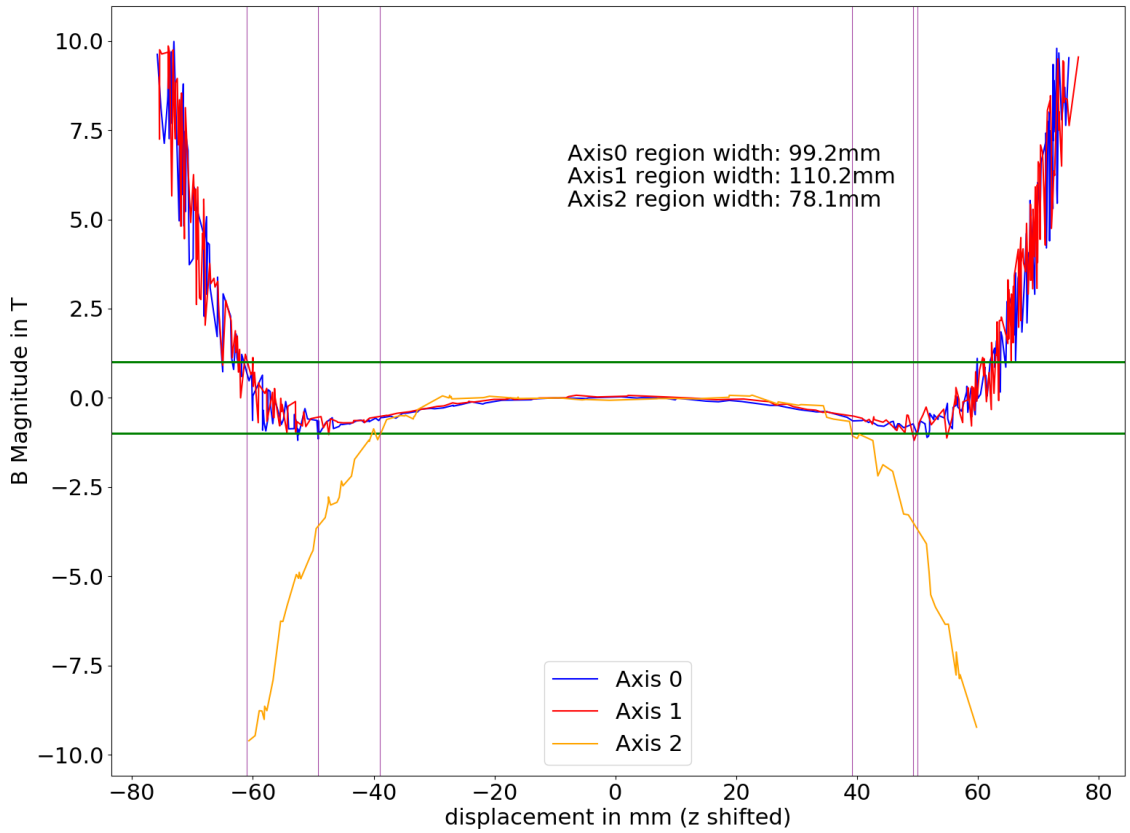

Figure 47: Graph showing the performance of a Maxwell coil. The coil used 301A·t split between the 3 coils in a ratio of 49:64 for the outer:centre coils. It had a (central) radius of 110mm. Axis 0, 1, 2 correspond to x, y, z respectively.

4. Maxwell with holes

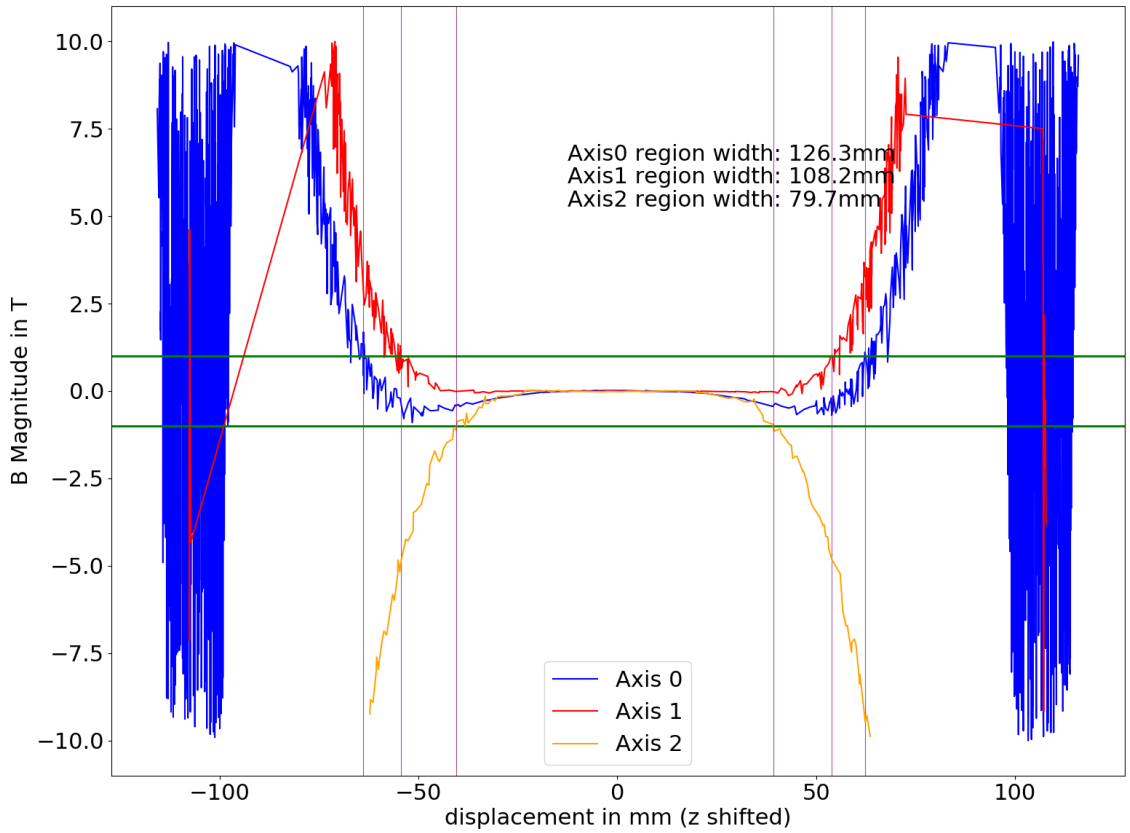

Figure 48: Graph showing the performance of a Maxwell coil with 2x30mm diameter avoidance holes on opposite sides of the central coil. The coil used 301A·t split between the 3 coils in a ratio of 49:64 for the outer:centre coils. It had a (central) radius of 110mm.

5. Final Design

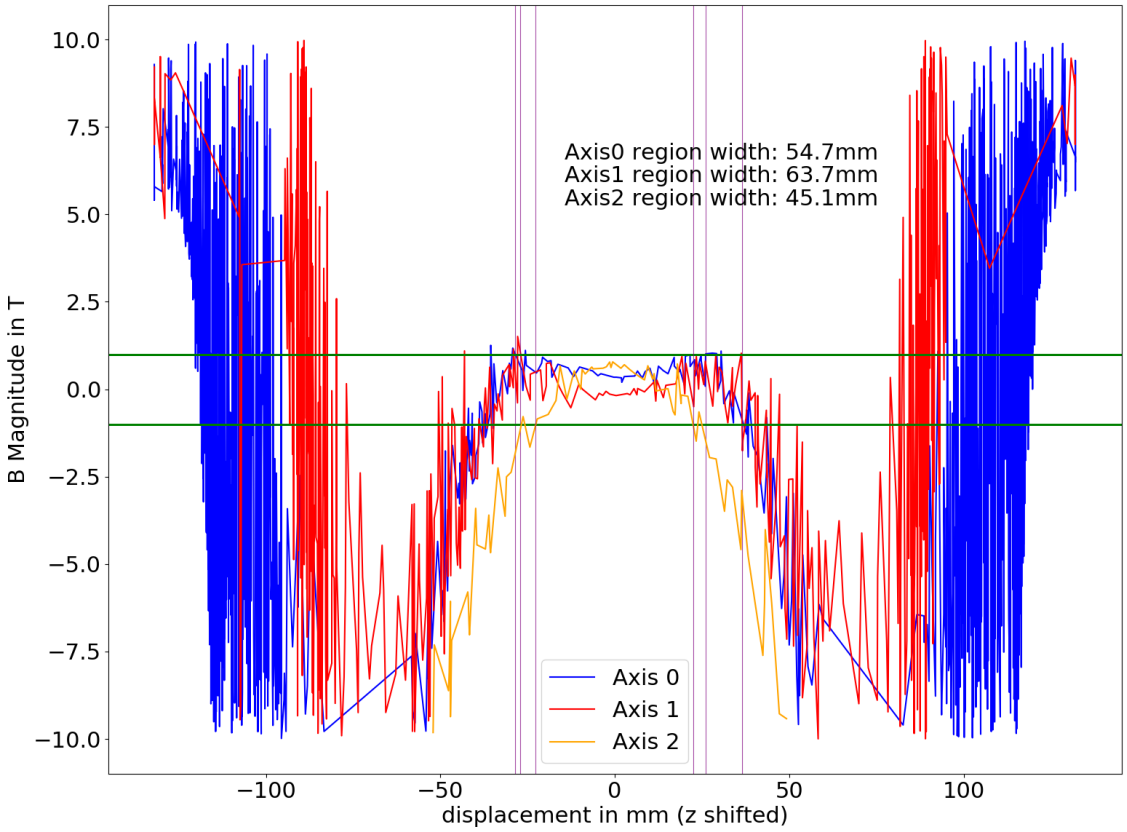

Figure 49: Graph showing the performance of a Maxwell coil with 2x30mm diameter avoidance holes on opposite sides of the central coil. The coil used 301A·t split between the 3 coils in a ratio of 49:64 for the outer:centre coils. It had a (central) radius of 110mm.

Appendix D: Probing a different state

Since the rate of spontaneous emission is higher than stimulated emission, as it is required for the optical pumping to work (see section II B 2) the atoms are far more likely to exist in the ground state than excited. Because of this it doesn't matter which excited state is probed since the optical birefringence is caused by the $\mathbf{m}_F$ population imbalance in the ground states (transferred between states via SECs).

Appendix E: Offset Lock Circuit Diagram

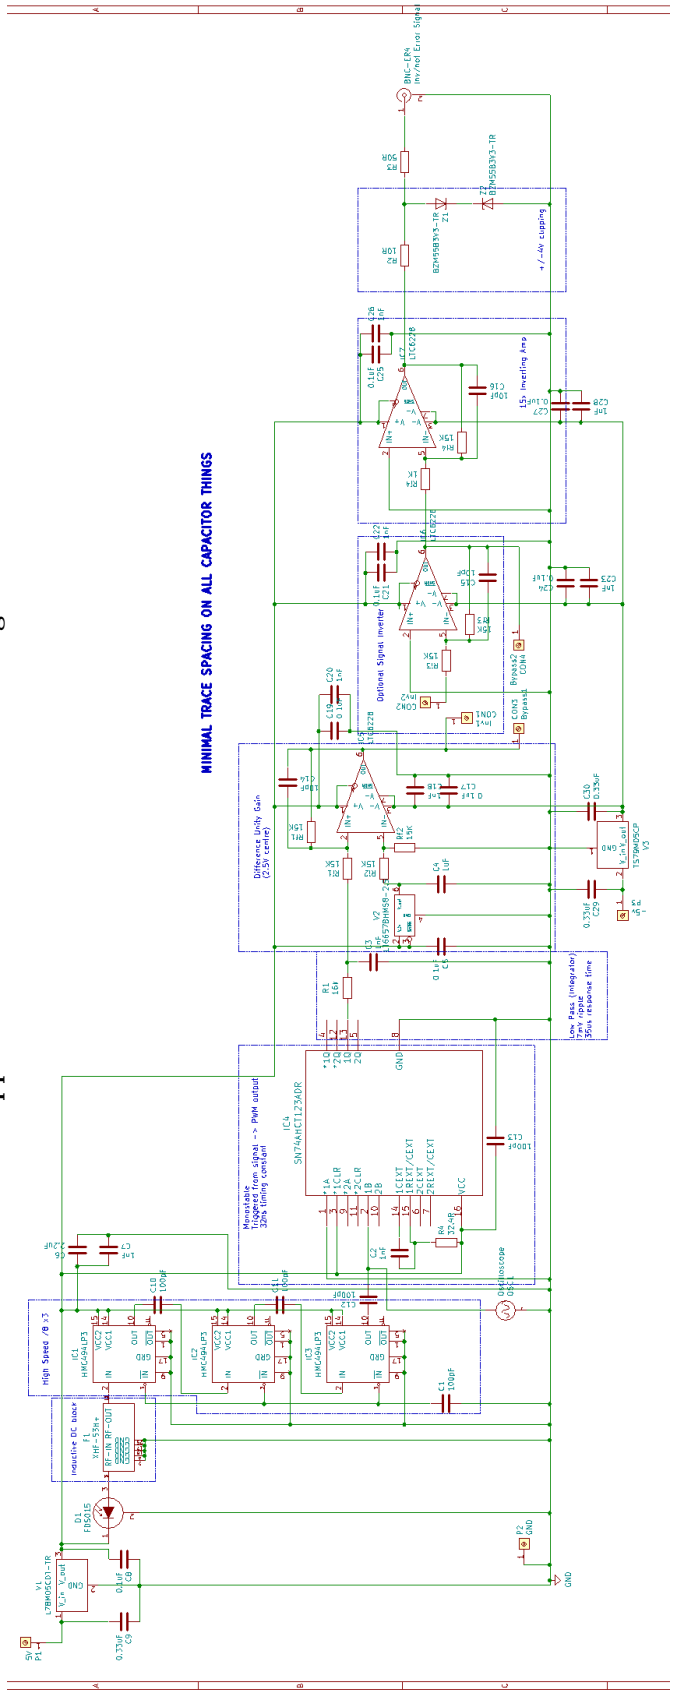

Figure 50: Drawing of the full offset lock circuit with actual components annotated.