UNIVERSITY OF OXFORD

Development of a SQUID Magnetometry System for a Cryogenic Neutron Electric Dipole Moment Experiment

by

Amy Cottle



A thesis submitted in partial fulfillment for the degree of Doctor of Philosophy in the

Department of Physics

Trinity Term 2015

Abstract

A measurement of the neutron electric dipole moment (nEDM) could hold the key to understanding why the visible universe is the way it is: why matter should predominate over antimatter. As a charge-parity violating (CPV) quantity, an nEDM could provide an insight into new mechanisms that address this baryon asymmetry. The motivation for an improved sensitivity to an nEDM is to find it to be non-zero at a level consistent with certain beyond the Standard Model theories that predict new sources of CPV, or to establish a new limit that constrains them.

CryoEDM is an experiment that sought to better the current limit of $|d_n| < 2.9 \times 10^{-26}\, e\,\mathrm{cm}$ by an order of magnitude. It is designed to measure the nEDM via the Ramsey Method of Separated Oscillatory Fields, in which it is critical that the magnetic field remains stable throughout. A way of accurately tracking the magnetic fields, moreover at a temperature $\sim 0.5\,\mathrm{K}$, is crucial for CryoEDM, and for future cryogenic projects.

This thesis presents work focussing on the development of a 12-SQUID magnetometry system for CryoEDM, that enables the magnetic field to be monitored to a precision of 0.1 pT. A major component of its infrastructure is the superconducting capillary shields, which screen the input lines of the SQUIDs from the pick up of spurious magnetic fields that will perturb a SQUID's measurement. These are shown to have a transverse shielding factor of $> 1 \times 10^7$, which is a few orders of magnitude greater than the calculated requirement.

Efforts to characterise the shielding of the SQUID chips themselves are also discussed. The use of Cryoperm for shields reveals a tension between improved SQUID noise and worse neutron statistics. Investigations show that without it, SQUIDs have an elevated noise when cooled in a substantial magnetic field; with it, magnetostatic simulations suggest that it is detrimental to the polarisation of neutrons in transport. The findings suggest that with proper consideration, it is possible to reach a compromise between the two behaviours.

Computational work to develop a simulation of SQUID data is detailed, which is based on the Laplace equation for the magnetic scalar potential. These data are ultimately used in the development of a linear regression technique to determine the volume-averaged magnetic field in the neutron cells. This proves highly effective in determining the fields within the 0.1 pT requirement under certain conditions.

Acknowledgements

This thesis represents the culmination of three and a half years of work, which has seen many surprises and challenges along the way. Though words cannot do you justice, this page serves as a small token of appreciation for all those who have touched me during this chapter of my life, and helped me see it through to its end.

First and foremost, I would like thank my supervisors, Sam Henry and Maurits Van der Grinten, for their guidance throughout the DPhil process, but also for the opportunities to work in exciting areas of research. I look back fondly on our trips to ILL and LSBB – few students can lay claim to have been truly involved in experiments at such a fundamental level. I would like to thank the teams at both institutions, not only for their assistance but also for their interesting, if often non-scientific, conversations – a warm mention especially for those members of the collaboration who have looked out for me over the years. May you long continue to work hard and produce good science.

I owe many thanks to Hans Kraus, for his kind words, advice and support, even when it was not sought. He deserves special acknowledgement for all he does in the lab, the experience he offers and the care and attention taken both over his work and over members of his group. It is small wonder that he inspires admiration and gratitude amongst his students, and I am sure he will do so for many years to come.

I am grateful for the camaraderie of my fellow students in the group: even though we worked on different projects, we bonded over common problems and revelations in the lab. I extend this sentiment in particular to Alice and Xiaohe, with whom I shared much of the trials and tribulations and the ups and downs of the PhD experience. I wish all of you the success you deserve wherever the future takes you.

Special thanks go to Mark Pipe, both as a colleague and as a friend. His unfettered optimism and sometimes blind enthusiasm in the face of scientific adversity were what carried me through those early years. His advice of "don't panic" was often noted but rarely heeded. I take with me the memories of sumptuous dinners, cryogenic mishaps and unexpected successes. Here's hoping that some of your positive attitude has rubbed off on me.

Lastly, I am indebted to my parents and to David, whose patience and unwavering support throughout this saga were invaluable, and who had the toughest job of all. They kept me in check when I was straying, brought me back to ground when I was caught dreaming, and offered a shoulder to cry on at the worst of times. I may take you for granted, but I can assure you that I am grateful, more than words can express, and I hope that you are proud of me for having finished what I set out to accomplish four years ago.

Thanks to all of you and to many others, close friends and peers, whom I may have failed to mention but whom I hold in the utmost regard. Thanks for being there, and making my stay at Oxford an enjoyable one.

Contents

Abstract i

Acknowledgements ii

Introduction 1

1 The Neutron Electric Dipole Moment (nEDM), d_n 3
- 1.1 Defining the nEDM . . . 4
- 1.2 A Brief History of CP Violation . . . 7
- 1.3 The Baryon Asymmetry of the Universe . . . 9
- 1.4 An nEDM in the Standard Model and Beyond . . . 11
- 1.5 EDM experiments . . . 15

2 Maximising the Sensitivity to d_n 19
- 2.1 Determining an EDM . . . 19
- 2.2 Neutron Density, N . . . 26
- 2.3 Storage Time, T . . . 30
- 2.4 Visibility, α . . . 33
- 2.5 Electric Field, E . . . 38
- 2.6 Magnetic Fields . . . 44
- 2.7 Magnetometry . . . 49
- 2.8 Summary . . . 58

3 Concept and Design of the SQUID Magnetometry System (SMS) for CryoEDM 60
- 3.1 A Brief Introduction to CryoEDM . . . 60
- 3.2 The DC SQUID as a Magnetometer . . . 66
- 3.3 CryoEDM SQUIDs . . . 77
- 3.4 Overview of the SMS for CryoEDM . . . 79
- 3.5 Readout cabling from 4.2 to 290 K . . . 82
- 3.6 The Input Circuitry . . . 87
  - 3.6.1 PTFE Connector Block . . . 88

3.6.2 Pickup Loops . . . 89
3.6.3 Calibration Coils . . . 92

4 The Characterisation of Solder Capillary Shields for the SMS 94
4.1 The Shielding Requirement for CryoEDM . . . 94
4.1.1 Choice of Shielding . . . 98
4.2 Measuring Pickup Factors . . . 103
4.2.1 Experimental Setup . . . 103
4.2.2 Converting Signals Into Pickup Factors . . . 105
4.2.3 Unshielded Pickup Factors . . . 108
4.3 Determining Shielding Factors . . . 112
4.3.1 Shielded Pickup Values and Shielding Factors . . . 112
4.3.2 Repairing Solder Shields . . . 115
4.4 Heaters . . . 118
4.4.1 Heater Power . . . 119
4.4.2 Magnetic Pickup . . . 120

5 The Effects of Cryoperm: SQUID Sensitivity vs Spin Depolarisation 125
5.1 The Reduction of Low Frequency, $^1/_f$ Noise . . . 126
5.1.1 The Origins of $^1/_f$ Noise in SQUIDs . . . 126
5.1.2 Preparatory Investigations in Different Cooling Fields . . . 129
5.1.3 Analysing the Noise Spectra . . . 132
5.1.4 Results . . . 135
5.1.5 Consequences for CryoEDM . . . 141
5.2 The Depolarisation of Neutrons in Transport . . . 145
5.2.1 Quantifying the Spin Depolarisation . . . 145
5.2.2 Modelling the Magnetic Environment . . . 148
5.2.3 The Effect of the Cryoperm . . . 150
5.2.4 The SQUID Compensation Coil (SQCC) . . . 153
5.2.4.1 Characterising the SQCC . . . 153
5.2.4.2 Improved Designs for the SQCC . . . 155
5.3 Summary . . . 160

6 The Simulation of the SQUID and Neutron Magnetic Signals 162
6.1 The Components of the Model . . . 163
6.1.1 The Magnetic Background . . . 163
6.1.2 The Pickup Loops . . . 167
6.1.3 The Ramsey Cells . . . 171
6.1.4 The Mu Metal & Superconducting Shields . . . 173
6.2 A Finite Difference Solver for Magnetic Fields . . . 175
6.2.1 The Algorithm . . . 176
6.2.2 The Relaxation Scheme . . . 181
6.2.3 Conversion to Fields . . . 185
6.3 Conclusions . . . 189

7 Determining the Fields in the Neutron Cells **190**
7.1 Introduction to Proposed Interpolation . . . 191
7.2 Linear Regression Method . . . 196
7.3 Neural Network Approach . . . 200
7.3.1 Refining the Neural Network . . . 205
7.4 Comparison of Neural Network and Linear Regression Methods . . . 212
7.5 Characterising the Fit . . . 217
7.6 Conclusions . . . 220

Conclusions and Further Outlook **222**

A Formulae for Coil-Generated Magnetic Fields **227**

B The Finite Multigrid Method (FMG) **229**
B.1 The Algorithm . . . 229
B.2 The Transfer of Information . . . 231
B.3 Computational Cost and Performance . . . 234
B.4 Improved Relaxation Method . . . 236

Bibliography **239**

Introduction

The aim of this thesis is to demonstrate that Superconducting Quantum Interference Devices, or SQUIDs, are a viable choice for a magnetometer in an nEDM experiment. In doing so, it documents the work done towards a 12-SQUID magnetometry system for CryoEDM, both experimentally in the design and characterisation of its components, and theoretically in the simulation of its operation. These are elements and considerations that could equally apply to a future cryogenic nEDM project.

The content can be divided into three distinct parts:

Chapters 1 and 2 set the scene by introducing the concept of an electric dipole moment and its importance in a cosmological context, before discussing how it could be measured. The first will give an overview of fundamental symmetries, and how the nEDM could play a role in determining the causes of symmetry breaking that may ultimately underpin the existence of matter in the universe. The second will describe in depth the wide range of technologies that are employed together in an nEDM experiment, and how the SQUIDs could fit into this picture.

Chapters 3 to 5 focus on the experimental issues behind the construction of a SQUID system. Chapter 3 starts by explaining what a SQUID is and how it operates, and then outlines the salient points of the CryoEDM experiment and how the SQUIDs are incorporated into its design. It describes the changes that were made to components of the system during the author's time on the experiment, with some of the elements completely redesigned.

Chapter 4 takes this further and looks specifically at superconducting magnetic shielding developed for input lines of the SQUID. It details experiments undertaken to establish its shielding factor, and the efficiency of heaters intended to heat this circuitry out of the superconducting state. Chapter 5 also looks at shielding, but this time for the SQUID chips themselves. It is observed that this shielding can ameliorate the noise level of a SQUID cooled in a magnetic environment, but can also disrupt the fields that maintain the polarisation of the neutrons during

transport. The chapter notes work done to investigate the extent of both effects and to consider the possible consequences they have in CryoEDM.

Chapters 6 and 7 present computational work to devise a simulation for SQUID and neutron magnetic data, and use the results to establish a prescription for determining the magnetic fields within the neutron cells using the SQUID signals. The former describes a model for the magnetic environment, the sensing areas of the SQUIDs, the neutron cells and the magnetic shielding in place for CryoEDM. It elaborates on a Laplace equation solver developed to compute the attenuation of the fields due to this shielding. Part of the chapter, as well as the appendix, is dedicated to processes used to accelerate the convergence of this solver.

Chapter 7 compares two possible methods for finding the magnetic fields present in the neutron cells, the main function of the magnetometer. It discusses refinements made to both techniques to optimise their properties for this purpose, and illustrates their predicted performance for CryoEDM.

Overall this thesis, from start to finish, presents a comprehensive account of the many aspects involved in the realisation of a SQUID magnetometry system for an nEDM experiment. It is intended that this work will provide a useful reference for future cryogenic nEDM projects, as an aid in the design of a similar system.

Chapter 1

The Neutron Electric Dipole Moment (nEDM), d_n

Low energy, precision experiments are increasing in importance as the search for physics beyond the Standard Model (BSM) intensifies. Fundamental neutron physics is one subset of this growing field, which exploits the neutron's unique properties to test for new physics. These include its electric neutrality and its long, but finite lifetime, which confer a definite advantage to using neutrons over other, less manipulable particles. Most of these experimental efforts concentrate on one of the many neutron-related observables that are sensitive to BSM physics. The list of such parameters is extensive, but include the coupling constants and decay correlation coefficients involved in neutron β decay, which are linked to the $|V_{ud}|$ element in the CKM matrix, and the neutron mass, which can provide limits on gravity-like forces in extra dimensions. Comprehensive reviews that highlight both the importance and vastness of this field can be found in [1] and [2].

A great many of these searches complement research conducted at particle accelerators like the LHC. In some cases, precision work at the low energy frontier offers a more effective path to look for tiny signatures of phenomena that in an accelerator would only manifest at energy scales far beyond the capabilities of current

technologies. The search for a neutron electric dipole moment (nEDM) is a good example: a quantity which, by its discovery at a given threshold, or lack thereof, can constrain or even eliminate numerous BSM theories. The current limit set by the Sussex-RAL-ILL collaboration of $2.9 \times 10^{-26}\, e\,\mathrm{cm}$ [3] encroaches upon the parameter space of Supersymmetric (SUSY) theories, a popular extension to the Standard Model. The question of why we should continue the search is addressed in this chapter, explaining why the nEDM is such an important quantity, its relation to fundamental symmetries, and how it could help shed more light on new physics.

1.1 Defining the nEDM

An electric dipole is most easily visualised as a pair of discrete, equal and opposite charges ($\pm q$), whose moment can be expressed as $\vec{d} = q\vec{x}$, where x is the distance between them. Most particle systems consist of a cloud of charged components, wherein a non-uniform distribution can similarly give rise to a polarised system. If there is a separation of charge into a positive and a negative centre, and if this exists without the presence of an external field, this is equivalent to a permanent electric dipole moment (EDM). This can be expressed as $\vec{d} = \int \rho(\vec{r})\vec{r}\,\mathrm{d}^3 r$, where $\rho(\vec{r})$ is the electrostatic charge distribution. In the case of distinct charges about a well-defined centre, the formula reduces to $\vec{d_n} = \sum_i q_i \vec{r_i}$, using $\rho(\vec{r}) = \sum_i q_i \delta(\vec{r} - \vec{r_i})$.

The Wigner-Eckart theorem imposes two important consequences on the neutron, being a spin-$\frac{1}{2}$ particle: one that it is forbidden from possessing higher moments than a dipole; two, that this EDM must be proportional to the angular momentum. As the spin vector is the only intrinsic direction associated with the particle, the nEDM must be aligned parallel or anti-parallel to it such that we may write $\vec{d_n} = -d_n \frac{\vec{s}}{|\vec{s}|}$. This is also true of any fundamental particle, and it can be shown that any non-zero value for the dipole moment leads to a charge-parity (CP) violating system.

There are two easy ways of seeing this. One is to consider the effect of charge (C), parity (P) and time (T) operations on the interaction Hamiltonian for the neutron in an electromagnetic environment,

$$H = \frac{-d_n \vec{s} \cdot \vec{E} - \mu_n \vec{s} \cdot \vec{B}}{|\vec{s}|}, \tag{1.1}$$

with $\vec{E}$ and $\vec{B}$ the external electric and magnetic fields respectively. As $\vec{E}$ is a polar vector whereas $\vec{s}$ and $\vec{B}$ are both axial vectors, the magnetic and electric interaction terms transform differently under P and T (Table 1.1). The Hamiltonian is therefore not invariant with respect to these symmetries, and thus to reinstate invariance, given the neutron has a documented magnetic moment [4], d_n must be set to zero.

Quantity	C	P	T
Electric Field, $\vec{E}$	$-\vec{E}$	$-\vec{E}$	$\vec{E}$
Magnetic Field, $\vec{B}$	$-\vec{B}$	$\vec{B}$	$-\vec{B}$
Spin, $\vec{s}$	$\vec{s}$	$\vec{s}$	$-\vec{s}$

TABLE 1.1: The transformation of relevant quantities under the actions of C, P, T operations.

This argument assumes that our definition of $\vec{d_n}$ in terms of $\vec{s}$ remains true irrespective of the fundamental symmetry operations. A second means of establishing that $\vec{d_n}$ is CP violating is to notice that this relationship is in fact not static under C, P and T (Figure 1.1). Let us pick one of the two possible configurations and assume that the dipole moment is aligned parallel with the spin direction. If we use the simple representation of the dipole moment, $\vec{d_n} = \sum_i q_i \vec{r_i}$, it is clear that the direction of $\vec{d_n}$ flips under parity, as each position vector becomes the negative of itself. The spin vector, however, is unaffected by P. Therefore $\vec{d_n}$ is now anti-aligned with $\vec{s}$ and thus not the same as our initial definition. Under T, the system is similarly not invariant, as $\vec{d_n}$ remains unchanged whereas, since T can essentially be thought of as a motion symmetry, $\vec{s}$ becomes $-\vec{s}$. One might also assume that $\vec{d_n}$ also violates C, as each q_i converts to $-q_i$, and $\vec{s}$ does not

change. However, the caveat is that, in swapping the quarks for antiquarks, we transform the particle into an antiparticle. Charge conjugation therefore leads to a physical entity, as an antineutron should have an electric dipole moment equal and opposite to the neutron, relative to its spin direction. Via the CPT theorem, any observable which violates T must also violate CP, which can easily be verified in the above two cases by considering the combined action of CP.

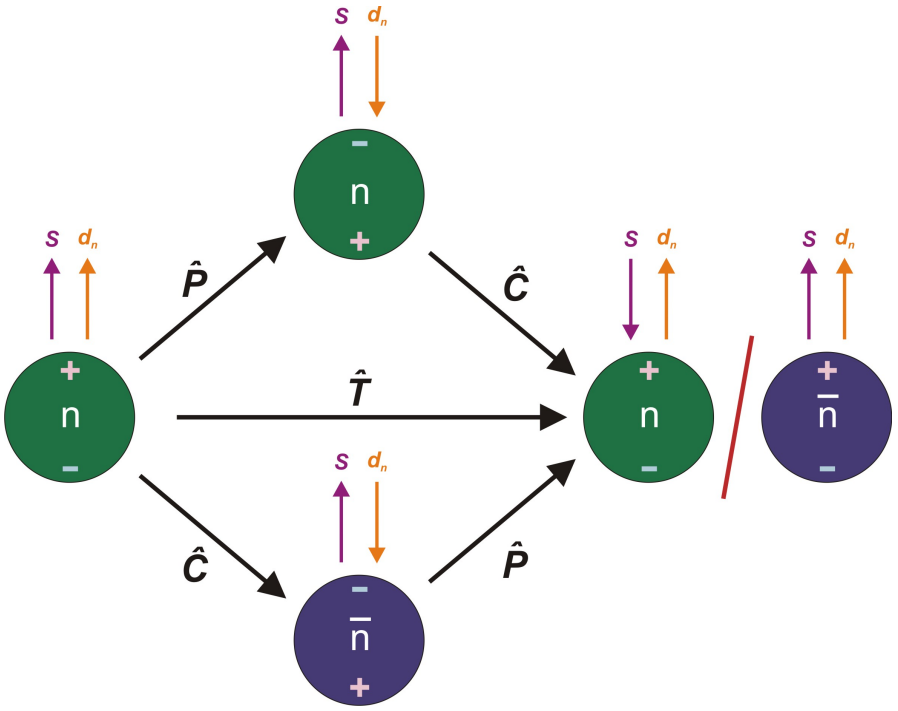

FIGURE 1.1: The neutron under the action of fundamental symmetries. Charge conjugation gives a physical antiparticle, whereas parity and time operators both produce an unphysical particle. The combined actions of $\hat{C}\hat{P}/\hat{P}\hat{C}$ yield the equivalent, symmetry-violating antineutron.

The symmetry properties of an EDM are in line with one of the five possible combinations allowed by the CPT theorem (Table 1.2) [5]. The conservation or violation of one of the symmetries C, P or T is tantamount to the conservation or violation of the product of the other two: that is, at least one of the other invariances must fail if one invariance is invalid. An EDM would be evidence of a set of preservations and violations under these symmetries that has not yet been observed with other CP-violating systems.

[1] As far as is known. There is good reason to believe that, for example, neutron decay is not CP conserving, see [6]

C	P	T	Example Occurences
✓	✓	✓	Strong & Electromagnetic Interactions
×	✓	×	None
×	×	✓	β Decay[1]
✓	×	×	***EDMs***
×	×	×	$K_L^0 \to 2\pi$

TABLE 1.2: Possible combinations of outcomes under C, P and T symmetries for CPT invariant interactions, ✓ representing conservation and × violation. The discovery of an EDM would be the first evidence of a process with the properties of row four.

The idea of CPT invariance was put forward independently by Lüders (1957) [7] and Pauli (1955) [8], such that it is sometimes referred to as the Lüders-Pauli theorem. It is yet not disproven because of the absence of any compelling evidence signalling Lorentz violation, a consequence of the breaking of CPT [9]. To fully understand the implications of CPT invariance and of CP violation, and the role the neutron EDM plays, it serves to first appreciate the rich history behind the subject, and the necessity of CP violation in modern cosmology.

1.2 A Brief History of CP Violation

It was long believed that C, P, T should be real symmetries that hold in their own right. The scientific community was rocked in the mid-20th century by a series of experiments that showed otherwise. In 1956, C.S. Wu demonstrated the non-conservation of parity in beta decay [10]. Spurred on by her colleagues Lee and Yang, whose review of the subject found no conclusive evidence that parity was unviolated in weak interactions as it was for strong and electromagnetic processes [11], Wu embarked upon an investigation into the nature of beta decay of cobalt-60. The cobalt nuclei were prepared at a low temperature in a large magnetic field to achieve a high degree of polarisation. Electrons from the decay should be emitted in roughly equal amounts parallel and antiparallel to the field if the weak

process is parity conserving, as the field and spin are constant under the action of parity. Two photons are subsequently released by the resultant excited nickel nucleus, through an electromagnetic interaction that should respect parity. By comparing the directions of the electrons with these photons, Wu deduced that the electrons were preferentially emitted in the opposite direction to the spin of the nuclei, no matter which way the field and thus the spins were oriented.

The weak force was later found to be maximally parity and charge violating, which explains, for example, why although a left-handed neutrino exists, a right-handed neutrino or left-handed antineutrino have never been observed. However, under the combined action of charge and parity (CP), the left-handed neutrino can be transformed into a right-handed antineutrino, a particle which has been seen in nature. In this new parity-violating world, Landau proposed that perhaps CP was the more fundamental symmetry that was obeyed [12].

The notion held until 1964, when this too was found to be violated in the decays of neutral kaons [13]. Two CP eigenstates had been proposed that were themselves combinations of the flavour eigenstates, $|K_0\rangle$ and its antiparticle, $|\bar{K}_0\rangle$:

$$|K_1\rangle = \frac{|K_0\rangle - |\bar{K}_0\rangle}{\sqrt{2}} \xrightarrow{\hat{C}\hat{P}} +|K_1\rangle \quad \& \quad |K_2\rangle = \frac{|K_0\rangle + |\bar{K}_0\rangle}{\sqrt{2}} \xrightarrow{\hat{C}\hat{P}} -|K_2\rangle\,. \quad (1.2)$$

These were thought to be synonymous with the physical, weak decay eigenstates, $|K_1\rangle = |K_S\rangle$ and $|K_2\rangle = |K_L\rangle$: the short-lived state, $|K_S\rangle$ should be CP-even, decaying into two pions, and the long-lived state, $|K_L\rangle$, CP-odd, decaying into three pions, at a slower rate because of the greater total mass of the product.

Cronin and Fitch observed kaon decay events at the end of a 57 ft beamline: this far from the source, the decays should be entirely those of the K_L eigenstate. They found that about 1 in 500 resulted in the production of two pions, contrary to its expected CP preserving properties. CP violation therefore leads to physical states

which are admixtures of the original CP eigenstates:

$$|K_S\rangle = |K_1\rangle + \epsilon\,|K_2\rangle \quad \& \quad |K_L\rangle = |K_2\rangle + \epsilon\,|K_1\rangle\,, \quad \epsilon \approx 2 \times 10^{-3}. \tag{1.3}$$

This indirect CP violation results from the mixing of the neutral K_0 with its antiparticle, also known as kaon oscillations. Direct CPV, however, is so called because it manifests as the direct decay of a CP eigenstate into products forbidden by CP symmetry e.g. $K_2 \to \pi\pi$. It took until 1999 to establish direct CP violation in the kaon system at the NA48 experiment based at CERN [14]. It was able to determine that the process occurs with a probability some 1000 times less than the indirect form, through the comparison of the branching ratios of the decays of K_L and K_S to two charged or two neutral pions.

Since then, CP violation has also been seen in other meson systems: the B factories, BaBar and Belle established CP is violated in the B system in 2001 [15] [16], and more recently in 2011, LHCb saw the first signs, later unsubstantiated, of direct CPV with D mesons [17]. All of these are the product of flavour-changing interactions of quarks that can be described by the Cabibbo-Kobayashi-Masakawa (CKM) matrix under the Standard Model. The fact that no deviations from the Standard Model have yet been confirmed is significant as it emerged that CPV is central to an outstanding cosmological question behind our existence, termed the baryogenesis problem.

1.3 The Baryon Asymmetry of the Universe

It is widely accepted that for the universe to exist as it does today, CP violation must have been present in its formation. CPV is one of three conditions stipulated by Sakharov in 1967 to account for the observable universe, alongside baryon number violation and thermodynamic non-equilibrium [18]. All three of these criteria are satisfied to some extent by the Standard Model. However, the search for new

origins of CPV persists as the Standard Model processes that have so far been documented lack any major sources, at least enough to explain the matter-antimatter symmetry that we see today. Planck results suggest that this asymmetry, expressed as the difference between the baryon and antibaryon densities, n_B and $n_{\bar{B}}$ respectively, is

$$\eta_B = \frac{n_B - n_{\bar{B}}}{n_\gamma} = (6.05 \pm 0.07) \times 10^{-10}, \tag{1.4}$$

where this has been normalised, as is tradition, to the photon number density, n_γ [19]. The same quantity, commonly referred to as the Baryon Asymmetry of the Universe (BAU), when calculated in the framework of the CKM matrix is some order 10^{16} smaller for the Standard Model prediction [20].

The Standard Model can also accommodate CP violation in the lepton sector through mixing analogous to that of quarks described by the CKM matrix. Within the traditional parameterisation of the Pontecorvo-Maki-Nakagawa-Sakata (PMNS) matrix, all mixing angles have now been established as non-zero, following the recent measurement of θ_{13} by Daya Bay [21]. This is a prerequisite for CPV, and allows for the possibility of leptogenesis that would, similar to baryogenesis, lead to a dominance of leptons over antileptons. Through sphaleron processes, this asymmetry can be converted into an equivalent in the baryon sector [22]. However, estimations of the CP violating phase, δ_{CP} in the PMNS matrix, which describes lepton oscillations, have so far been consistent with zero.

Can the observed BAU be reconciled without the need for new sources of CPV? One possibility is that this imbalance was inherent in the universe, but akin to the anthropic principle, this leads to a sterile hypothesis which moreover is not compatible with most favoured inflation theories. It could also be that the universe formed with pockets of antimatter and therefore that it is baryon symmetric at a 'global' scale. This has its own problems: causality forbids antimatter and matter from separating faster than they would annihilate, the 'annihilation catastrophe.' A signature of such a universe would be annihilation radiation produced at the

interfaces between matter and antimatter domains, which has not been observed. No instrument has yet found any trace of a significant source of antimatter: the Alpha Magnetic Spectrometer, for instance, set an upper limit on the ratio of antihelium to helium of $< 10^{-6}$ [23]. The visible universe, at least, remains matter-dominated.

The solution to the baryogenesis problem therefore appears to lie in as yet unestablished CP violating processes beyond the Standard Model (BSM). Efforts continue to study meson systems, placing tighter constraints on the CKM unitarity triangles and thus restricting the parameter space of viable BSM theories. Future neutrino projects, such as the Long Baseline Neutrino Experiment, have been given high priority and will have vastly improved sensitivities to δ_{CP}. Lastly ATLAS and CMS can study CPV in the production and decay of top quarks and, with its recent discovery, the Higgs boson.

However, perhaps the most sought after proof of CPV comes in the form of the existence of electric dipole moments of fundamental and composite particles. These measurements, because of their high precision, should be sensitive to signatures of BSM physics at the TeV scale. How the nEDM fits into this picture, how it can be determined and the consequences if it were to exist are discussed in the subsequent sections.

1.4 An nEDM in the Standard Model and Beyond

There are two potential sources of CP violation in the Standard Model relevant to the neutron: the first results from the constituent quarks via the CKM matrix. A quark EDM is only generated at the three-loop level, since there are no flavour changing processes and any two-loop contributions accidentally cancel when summed over all quark flavours. The EDMs of both the down and up quarks are found in this framework to be of the order of $\approx 10^{-34}\, e\,\mathrm{cm}$ [24]. The neutron

EDM is estimated under the naïve quark model (NQM) as a combination of these quark EDMs[1]: using the SU(6) coefficients,

$$d_n \sim \frac{4}{3}d_d - \frac{1}{3}d_u, \tag{1.5}$$

which leads to a d_n of similar magnitude as that of one of its quarks. This approach overlooks the multi-parton structure of the neutron, which can be addressed by applying perturbation theory to processes involving two or three quarks. This parton quark model uses a different weighting for the quark EDMs and includes contributions from strange sea quarks. The method yields an estimated neutron EDM of $10^{-32}\,e\,\mathrm{cm}$ [25]. Whichever treatment is used, the Standard Model EDM arising from the electroweak sector is well below the current limit and remains out of reach for the foreseeable future.

Another contribution to a hadronic EDM within the Standard Model could come from CP violation present in the strong sector. It was only with the emergence of QCD that this became theoretically possible. The generalised QCD Lagrangian includes a CP violating term introduced by t'Hooft as a resolution for the $U(1)_A$ problem:

$$L_\theta = \theta\frac{\alpha_s}{8\pi}G^{\mu\nu}\tilde{G}_{\mu\nu}, \tag{1.6}$$

where $\alpha_s = {}^{g^2}/_{4\pi}$ is the strong coupling constant, and where $G^{\mu\nu}$ and $\tilde{G}_{\mu\nu}$ are the gluonic field tensor and its dual. The term represents the interactions of gluons with the vacuum, and θ may be interpreted as the vacuum expectation value of the gluon field. When QCD is fully integrated into the Standard Model, the physical value, $\bar{\theta}$, acquires a contribution from the quark mass matrix, diagonalised after symmetry breaking, $\bar{\theta} = \theta + arg(\mathbf{M_q})$.

[1]The spin-up wave-function of the neutron is $|n_\uparrow\rangle = \sqrt{\frac{2}{3}}\,|d_\uparrow d_\uparrow u_\downarrow\rangle - \sqrt{\frac{1}{6}}(|d_\uparrow d_\downarrow u_\uparrow\rangle + |d_\downarrow d_\uparrow u_\uparrow\rangle)$ such that the expectation value of the electric dipole moment operator, $\langle n|\,\hat{d}\,|n\rangle$, then yields the relation in Equation 1.5. This NQM prescription thus accounts for the valence quarks, but not the quark sea or any polarisation or movement of the constituents. Its continued success lies in the fact that the neutron can be effectively modelled as a bound system of non-relativistic, heavy consituents such that higher order, non-perturbative effects do not come into play. It can be further shown, under the chiral quark model, that the orbital and spin angular momenta of the quark sea provide contributions which roughly cancel each other in the evaluation of a magnetic or electric dipole moment [26].

The theta-mechanism has the ability to generate sizeable hadronic EDMs, and under the framework of chiral perturbation theory, the neutron EDM can be estimated as

$$d_n \sim e\frac{m_*\bar{\theta}}{m_n^2} \sim \bar{\theta}\cdot(6\times10^{-17})\,e\,\mathrm{cm}, \tag{1.7}$$

where $m_* = \frac{m_u m_d}{m_u+m_d}$ is the reduced quark mass [27]. Other schemes, such as calculations via the constituent quark model and QCD spectral sum rules, lead to similar estimates of $d_n \sim f\bar{\theta}\times10^{-17}e$cm, where $4 \leq f \leq 40$ [28]. The current limit on the neutron EDM thus translates to $\bar{\theta} \lesssim 10^{-9}$, and not order unity as would be naturally expected. This fine-tuning conundrum is commonly referred to as the Strong CP problem. Several mechanisms have been proposed in the past few decades to resolve the Strong CP problem, and the relative insensitivity of the neutron EDM to the Standard Model electroweak CPV phase δ_{CKM} as compared to $\bar{\theta}$ makes it essentially a background-free diagnostic for new flavour-diagonal CP-odd sources introduced by these models.

One simple solution is to have a quark with no mass, $m_u = 0$. Then $\bar{\theta}$ becomes unphysical and is assimilated by a phase transformation into the massless quark. As light quark masses are determined indirectly, there remains the possibility that one could be massless. However, this idea is heavily discounted by chiral perturbation theory and lattice calculations, both of which conclude that all quarks are massive [27].

Perhaps the strongest candidate promotes $\bar{\theta}$ to a dynamical field, known as the axion, thereby introducing a further symmetry, the Peccei-Quinn symmetry [29]. It is the vacuum expectation value of this pseudo-scalar Goldstone boson which renormalises $\bar{\theta}$ and provides a cancellation mechanism so that its value spontaneously relaxes to zero. There has been to date no evidence for the existence of an axion despite extensive searches [30].

A final way to engineer $\bar{\theta}$ to zero is to spontaneously break CP. Left-right symmetric theories provide a natural suppression of strong CP violation, as non-zero

contributions to $\bar{\theta}$ manifest only at the two loop level [31]. Nelson-Barr models form another set of solutions which assume CP invariance up to high energies, and thus have an associated heavy scale that requires heavy quarks [32]. Such models generate small values of $\bar{\theta}$ via radiative corrections. Both of these popular theories have variants compatible with the notion of Supersymmetry (SUSY).

SUSY theories are well-motivated as they address issues that are noticeably lacking in the Standard Model: that is, a natural solution to the hierarchy problem, gauge coupling unification and promising candidates for dark matter. The essential feature which characterises all Supersymmetric extensions to the Standard Model is the introduction of fermionic superpartners for each boson described by the latter, and vice versa. Since these spartners have not been observed, SUSY must be a symmetry which holds for high energies but is broken below the electroweak scale. The extra particles are associated with additional phases that provide further potential for CP violation. In fact, there is typically an overgeneration of CP breaking mechanisms, with estimates of d_n from SUSY models generally around and above the current experimental limit: large contributions to EDMs can arise at the one-loop level, as opposed to the three-loop level in the Standard Model. The need to suppress these in line with EDM bounds is the basis of the SUSY CP problem [33]. The constraint can be expressed, in a general form, as [27]

$$\delta_{\mathrm{CP}} \left(\frac{1\mathrm{TeV}}{M_{\mathrm{SUSY}}} \right) < 1 \ : \tag{1.8}$$

in other words, any CP violating phases δ_{CP} must be small, or the SUSY mass scale must be much greater than $O(\mathrm{TeV})$. Insisting that the phases be small leads to a new fine tuning problem, whereas maintaining that the spartners be heavy is disfavoured as then they cannot provide effective corrections to the Higgs mass. Fortunately, there are numerous modifications that can help reconcile these stringent conditions on SUSY, such as the proposition that CP is an approximate symmetry, or ensuring accidental cancellations between contributions of CP-odd sources.

Despite these setbacks, given the many variants of SUSY models, it remains an attractive theory, compelled by its resolution of Standard Model problems.

A wealth of other theories exist that introduce further CP violating sources, such as multi-Higgs models, all of which predict values for a neutron EDM. It is fair to say that the neutron EDM has ruled out more theories than any other parameter in physics, let alone particle physics. The results from the past decades of experiment have constrained the parameter space of SUSY, and have begun to touch upon the limits of other BSM theories. Clearly, further determination of the neutron EDM is beneficial to discriminate between feasible and infeasible new physics models.

1.5 EDM experiments

One might ask the question of why the neutron EDM, and not the EDM of another particle, has turned out to be so instrumental. A primary reason is the relative ease of interpretation of an nucleon EDM as opposed to a more complex system. Let us first consider atomic EDMs. Under the Schiff theorem, the EDMs of atoms are heavily suppressed by the rearrangement of their constituents to counteract the effect of an applied electric field, leading to no net EDM [34]. With heavy diamagnetic atoms, approximations regarding the point-like nature of the particle and its charge distribution break down, such that the corrections due to its finite size and non-electrostatic interactions between the constituents become non-negligible. In heavy, paramagnetic atoms, Schiff shielding is ineffective due to relativistic motion of the atomic electrons. Such systems can therefore have a predicted EDM that is within experimental reach, but whose origins are complicated by the interplay between atomic, nuclear and hadronic physics.

Likewise, molecules are rather complex systems capable of having an EDM. However, molecular EDMs can sometimes exist without being T and hence CP violating as a result of any degeneracy between its states. Ammonia, for example, exists

as two chiral isomers that are transformed into one another under the action of T. Polar molecules (and atoms) are suitable for searching for an electron EDM that would otherwise be elusive. The internal electric fields are as high as tens of GV/cm, which enhances the EDM of an unpaired electron orbiting the molecule. There are several viable molecules, including YbF and ThO, investigated by the Imperial College [35] and Advanced Cold Molecule Electron EDM (ACME) [36] experiments respectively. The latter holds the current limit, taking advantage of the 84 GV/cm field inside the ThO molecule. The story for the electron EDM is similar to that of the neutron, with a non-zero EDM arising only at high-order loop level, a small Standard Model prediction, and a much higher estimate from many BSM theories.

Particle	Limit (e cm)	SM Prediction (e cm)
n	2.9×10^{-26} (90% C.L.) [3]	$< 10^{-32}$ [27]
e	8.7×10^{-29} (90% C.L.) [37]	$< 10^{-38}$ [38]
μ	1.8×10^{-19} (95% C.L.) [39]	$< 10^{-36}$ [40]
^{199}Hg	3.1×10^{-29} (95% C.L.) [41]	$\sim 10^{-33}$ [27]
^{129}Xe	4.1×10^{-27} [42]	$\sim 10^{-34}$ [33]

TABLE 1.3: Current experimental limits for a selection of particle EDMs, compared with their expected values in the Standard Model.

In fact, there are many particle EDMs whose current bound is well above the Standard Model estimate, and which represent good searching grounds for new physics: the status of some of these is highlighted in Table 1.3. Should a permanent EDM be found, it alone cannot elucidate which mechanism is responsible for its presence. Multiple EDM determinations are required to differentiate between the possible sources of CP violation or, in the case of composite systems, the prevalent EDMs of their constituents. What makes the neutron EDM stand out from that of elementary particles is that it can be probed directly, but most significantly, it is most sensitive to any CP violation from the strong sector and thus the root of the Strong CP problem. Given the theoretical implications, it is small wonder that there are a number of experiments worldwide dedicated to finding a nEDM.

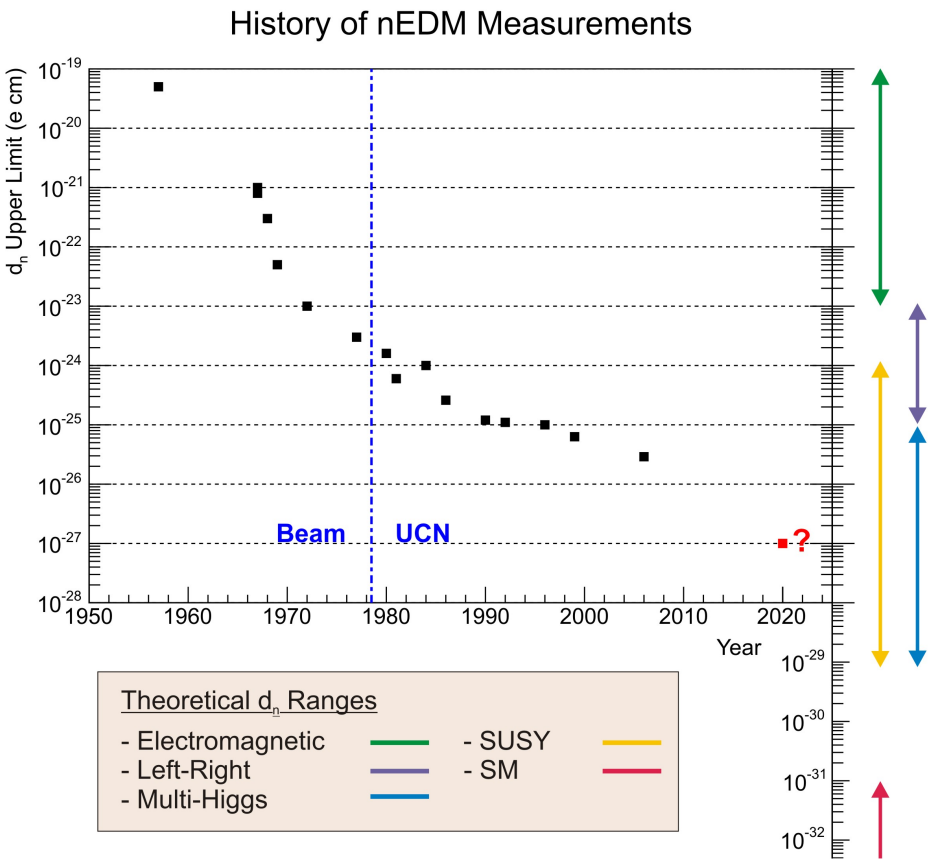

FIGURE 1.2: The status of nEDM experiments, with predictions of d_n from various classes of theoretical model. Current and near-future projects should produce another magnitude of improvement by the end of the present decade.

One of these, the Crystal Diffraction nEDM Experiment, uses a cold neutron beam. The neutrons are reflected within the quartz crystal under the Bragg condition, causing spin rotation. The interaction may be short, but the strong electric fields present in the lattice due to the interference of the nuclear and the electromagnetic parts of the scattering amplitude can be a few orders of magnitude greater than that in vacuum, and therefore magnify the effect. They predict an achievable sensitivity of $\sim 2.5 \times 10^{-26}\, e\,\mathrm{cm}$ [43].

A further experiment has been proposed that uses a pulsed neutron beam that could in theory compete with the sensitivity of ultra-cold neutron (UCN) experiments [44]. However, UCN are generally favoured by the neutron EDM community

because of the greater neutron densities and interaction times possible. It is therefore hardly surprising that the five other contemporary projects, the nEDM experiments planned at the Paul-Scherrer Institute (PSI), by the Technische Universität München (TUM) at the Forschungsreaktor München II reactor (FRM-II), at the Spallation Neutron Source (SNS) at the Oak Ridge National Laboratory (ORNL), by the Petersburg Nuclear Physics Institute (PNPI) at the Institut Laue-Langevin (ILL), and at the Canadian facility, TRIUMF, will all use bottled UCN.

Since the first nEDM measurement by Smith, Purcell and Ramsey in 1957 [2] [45], each decade has seen their limit of $d_n = 5 \times 10^{-20}\, e\,\mathrm{cm}$ surpassed by roughly a further order of magnitude (Figure 1.2). That is, until the present decade: pushing the limit further has proved difficult, and the current generation of experiments plan to incorporate further technologies, not least cryogenics, in an attempt to realise EDM sensitivities of $10^{-27}\, e\,\mathrm{cm}$ and better. Further details about these efforts are given in the next chapter.

[2] The experiment was conducted in 1951, but the results were only published in 1957 in a climate where parity violation was recognised.

Chapter 2

Maximising the Sensitivity to d_n

This chapter aims to give an overview of the recent technological developments that will be employed to improve the sensitivity to d_n. As mentioned in Chapter 1, almost all current experiments will involve UCN. Moreover, they will nearly all revolve around the same central principle to make their nEDM measurements, formulated by Ramsey some 60 years ago. An outline of this method is given, not only to explain how an EDM is derived, but also to qualify the important quantities that need to be bettered to achieve a greater sensitivity. This is followed by a review of the relevant state of the art, with a particular emphasis on cryogenic experiments and magnetometry, and thus how SQUIDs can form an influential part of any EDM experiment.

2.1 Determining an EDM

The first challenge in detecting a neutron EDM is to find a way of quantifying its value. Following on from Equation 1.1, it can be seen that a neutron in the presence of an electric and magnetic field will precess with a frequency of

$$\nu_{\uparrow\uparrow \text{ or } \uparrow\downarrow} = \frac{2(\vec{\mu_n} \cdot \vec{B} \pm \vec{d_n} \cdot \vec{E})}{h}, \tag{2.1}$$

with an electric field parallel to the magnetic field (↑↑) enhancing the precession frequency, and anti-parallel (↑↓) subtracting from it. The difference between the two configurations is thus directly sensitive to the nEDM:

$$d_n = \frac{h(\nu_{\uparrow\uparrow} - \nu_{\uparrow\downarrow})}{4|\vec{E}|}. \tag{2.2}$$

This differential method accounts for the inevitable existence of some magnetic field, and also boils down the dependence of d_n on just two parameters: assuming the electric field is known, there remains the question of how to extract the precession frequency. This can be done using a magnetic resonance technique known as the Ramsey Method of Separated Oscillatory Fields. Unlike its original incarnation in Ramsey's experiment of 1951, in a bottled UCN experiment, the pulses are separated temporally rather than spatially.

The sequence of steps involved is visualised in Figure 2.1. The neutrons are prepared so that initially their spins are along the z-direction (defined as spin-up), aligned with a static magnetic field, $\vec{B_0}$. The application of a burst of AC magnetic field in the x-y plane causes the spin vector to rotate. In the spin's frame of reference, even when it begins to move, this additional field appears static and perpendicular to $\vec{B_0}$, so that it begins to precess around it. The neutron spin describes a helical spiral until it reaches the x-y plane, whereupon the oscillating field is switched off: this period, t, of exposure to the AC field is therefore often termed a $\pi/2$ pulse. A much longer period, T, of free precession ensues, after which the AC source is reengaged to apply a second $\pi/2$ pulse coherent with the first, and again for time t. The frequency of the AC field is engineered to match the Larmor frequency of the neutrons when exposed to just the initial, static B-field i.e. $\omega_{AC} = \gamma|\vec{B_0}|$, where γ is the neutron gyromagnetic ratio, in which case, the product of the cycle is a neutron in the spin down state. However, the presence of any additional fields will cause a deviation of the neutron precession frequency, leading to an accumulation of a phase difference during the period of free precession (and

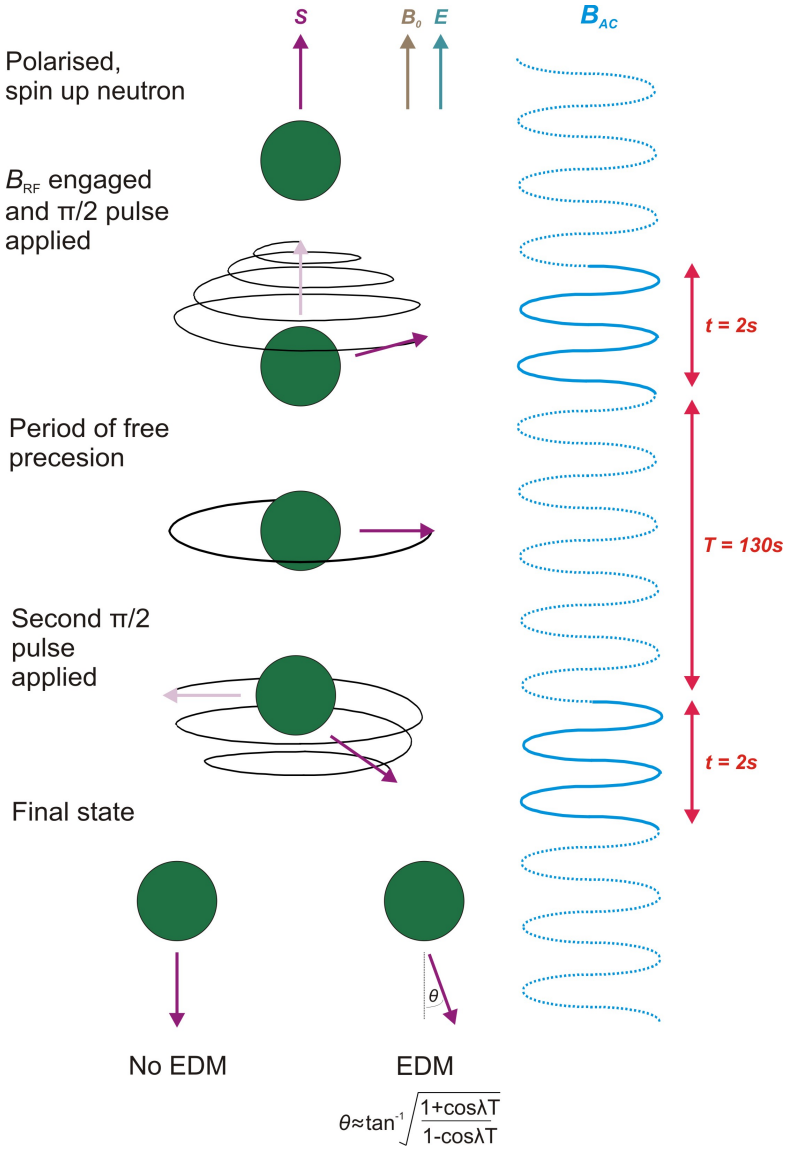

FIGURE 2.1: The Ramsey resonance sequence, with typical $\pi/2$ pulse and free precession times. The phase accumulation during free precession due to an EDM will result in a small spin up component in the final state. An approximation for the angle between the final spin direction and the spin down state is given, if working in the central fringe, where $\lambda = \omega - \omega_0$ is the difference between frequency, ω, of B_{AC} and the Larmor precession frequency, ω_0, of the neutrons in the static field, B_0.

essentially only then, provided that $T \gg t$). In this scenario, the second $\pi/2$ pulse results in a neutron in a superposition of spin up and down states, governed by the magnitude of the phase shift.

Upon the application of a constant electric field parallel or antiparallel to $\vec{B_0}$ throughout a resonance cycle, a difference should be seen that is attributable

to an electric dipole moment contribution. However, this necessitates that the introduction of an electric field does not induce a change in the setup or a further magnetic field. In addition, the magnetic field needs to be accurately known before and during the measurement, so that any instabilities or inhomogeneities do not introduce a false EDM reading. For instance, an electric field of 10 kV/cm with an EDM of $1 \times 10^{-26}\, e\,\text{cm}$ would result in a frequency shift of $\approx 5 \times 10^{-8}\,\text{Hz}$. Referring again to Equation 2.1, we see that a misestimation in the magnetic field of $\Delta B = \frac{d_n E}{\mu_n} \approx 2\,\text{fT}$ would produce the same frequency shift. It is thus imperative that the magnetic field be homogeneous within the vicinity of the neutron cell and as temporally invariant as possible, and that should any small changes in the field occur, they should be tracked to a high level of precision (see Section 2.7).

Detection of the spin states of the neutrons after the resonance cycle enables the phase shift to be reconstructed, and therefore the perturbed frequency. Maintaining a constant static field $\vec{B}_0$ and sweeping the frequency of the AC source through values either side of the Larmor frequency, the neutron counts of a particular spin state form a pattern known as the Ramsey resonance curve. From now on, we shall refer to the resonance curve as that of the spin up state. This curve, illustrated in Figure 2.2, has notable characteristics: a modulating envelope representing Rabi resonance when there is no period of free precession between the AC pulses; a minimum when the AC frequency equals the neutron precession frequency, and fringes either side corresponding to Ramsey resonance; a height between minimum and maximum dependent on the visibility, α, which is connected to the degree of polarisation and the accuracy of the detectors. An EDM contribution would therefore be tantamount to shifting the curve slightly to the left or right. As the expected displacement due to an EDM is small, the best hope of spotting such a shift would be to operate at the steepest places on the curve, at the working points illustrated. This is far preferable to generating a new Ramsey curve when the electric field is operational, since this would rely upon the temporal stability of the environment and suppression of systematics throughout data taking. The question is how to

use the neutron counts at these points to find an estimate for d_n. First, we need an expression for the neutron counts themselves.

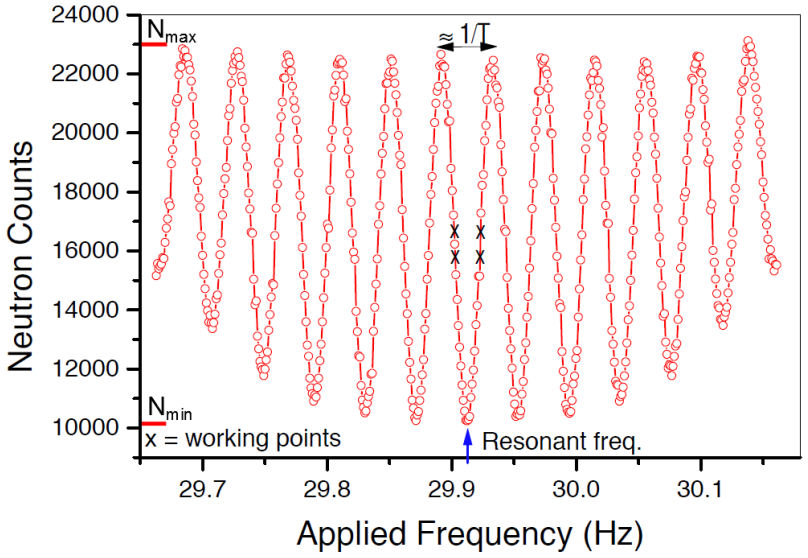

FIGURE 2.2: A Ramsey resonance curve obtained during the limit-holding experiment, with a B_0 of 1μT [46]. The working points used in data taking for an EDM measurement are shown.

The equation that underlies the final spin-state distribution is that of a probability of a spin flip as a consequence of the Ramsey resonance sequence [47],

$$P = 4\sin^2\Theta\sin^2\frac{a\tau}{2}\left[\cos\left(\frac{\lambda T}{2}\right)\cos\left(\frac{a\tau}{2}\right) - \cos\Theta\sin\left(\frac{\lambda T}{2}\right)\sin\left(\frac{a\tau}{2}\right)\right]^2, \quad (2.3)$$

with $\lambda = \omega - \omega_0$ being the difference between the frequency of the oscillating field, ω, and the Larmor frequency associated with the static field, $\omega_0 = \gamma_n B_0$, i.e. the resonance frequency; τ is the total duration of the $\frac{\pi}{2}$ pulses and T the period of free precession; $a = \sqrt{\lambda^2 + \left(\frac{\omega_l}{2}\right)^2}$, $\cos\Theta = \frac{\lambda}{a}$ and $\sin\Theta = \frac{\omega_l}{2a}$ related to the Larmor frequency in the oscillating field, $\omega_l = \gamma_n B_{\mathrm{AC}}$. All other factors being equal, the probability is maximised when $\lambda = 0$: that is, when the AC frequency matches that of the precessing neutrons in B_0. This corresponds to the central minimum of the resonance curve, where the spin up counts are fewest.

In general, after one Ramsey cycle, the number of spin up neutrons is given by $N_u = N(1-P)$, with N the total number of neutrons and P the probability that a neutron spin flips during the cycle, if all the neutrons begin in that state. Of

course, the situation is not quite so simple as we have neglected any possibility that the neutrons could have changed spin state before or after the cycle or that their spin states could have been misidentified as a result of any deficiencies of the detectors (see Section 2.4). As these are independent events, their likelihoods can be amalgamated into a second probability (much as in Equation 2.4,P_S – the non-Ramsey spin flip probability. Denoting the original, Ramsey resonance probability as P_R, the expectation for a spin-up count is given by

$$P_u = (1 - P_S)(1 - P_R) + P_S P_R = 1 - P_S + P_R(2P_S - 1). \tag{2.4}$$

Thus the expected value for the number of spin up neutrons becomes

$$N_u = NP_u = N(1 - P_S + P_R(2P_S - 1)). \tag{2.5}$$

If P_S is independent of the parameters in the Ramsey resonance technique, such that it is the same for all frequencies of the applied AC field, then the extrema of the spin up counts are still governed by those of P_R.

Ideally we would like a simple way of expressing the spin up counts in the vicinity of the central fringe, encompassing the working points in which we are interested. The central minimum where $\lambda = 0$ can be explicitly verified by substitution, with $P_R = 1$, using the identity $\omega_l = \pi/\tau$. Intuitively, the first maxima should occur when the applied frequency waveform is in antiphase with that of the neutron precession frequency, i.e. where $|\lambda| = \pi/T$. We can use some approximations to simplify P_R within the range of the two maxima, where $|\lambda| < \pi/T$. First, $\lambda \ll \omega_l$ because $T \gg \tau$ and thus a is dominated by the ω_l term, leading to $a \approx \frac{\omega_l}{2}$ and $\sin\Theta \approx 1$. It follows that $\cos\Theta \approx 0$, and thus we can ignore the second term in brackets in Equation 2.3. P_R therefore reduces to

$$P_R = \cos^2\left(\frac{\lambda T}{2}\right) = \frac{1}{2}(1 + \cos(\lambda T)). \tag{2.6}$$

The number of spin up neutrons can therefore be written as

$$N_u = \frac{1}{2}N(1 + (2P_S - 1)\cos(\lambda T)), \tag{2.7}$$

using Equation 2.5. Further, we can define the visibility of the central fringe as

$$\alpha = \frac{N_{\max} - N_{\min}}{N_{\max} + N_{\min}} = 1 - 2P_S, \tag{2.8}$$

again using Equation 2.5 for $N_{\max} = N(1 - P_S)$ when $P_R = 0$, $N_{\min} = NP_S$ for $P_R = 1$. This gives the final expression of

$$N_u = \bar{N}(1 - \alpha\cos(\lambda T)), \tag{2.9}$$

where $\bar{N} = \frac{N}{2}$, the average spin up counts across the fringe or equivalently at half-height. A similar equation can be derived for the spin down counts, with the minus replaced by a plus, assuming the same α, although in practice this could be different if there were any bias towards one spin state in the measurement system.

To determine an EDM, we also require the neutron counts in the presence of an electric field. The expression becomes

$$N_u = \bar{N}(1 - \alpha\cos(\lambda \pm \omega_E)T), \tag{2.10}$$

where ω_E represents the shift to the precession frequency introduced by $\vec{E}$, adding or subtracting from the Larmor frequency depending on whether it is parallel ($\uparrow\uparrow$) or antiparallel ($\uparrow\downarrow$) to $\vec{B_0}$. Expanding the cosine, the difference between the two becomes

$$N_{\uparrow\downarrow} - N_{\uparrow\uparrow} = 2\alpha\bar{N}\sin\lambda T\sin\omega_E T. \tag{2.11}$$

Assuming that $\lambda T = \frac{\pi}{2}$ if we operate at the steepest point on the resonance curve, and that the perturbation caused by the EDM is small such that $\sin\omega_E T \approx \omega_E T$,

we have

$$\Delta N = N_{\uparrow\downarrow} - N_{\uparrow\uparrow} = 2\alpha\bar{N}\omega_E T = \frac{2N\alpha d_n E T}{\hbar}, \tag{2.12}$$

reinstating N in place of $2\bar{N}$, and using $\omega_E = \frac{2d_n E}{\hbar}$: again assuming that we are at half-height, N could be estimated by adding the two counts. Using these two counts alone, we arrive at

$$d_n = \frac{\hbar \Delta N}{2\alpha E T N}. \tag{2.13}$$

To estimate the statistical uncertainty in d_n, that is the error which arises solely from that in counting the neutrons, we must derive an expression for the uncertainty in the neutron counting. This is

$$\sigma = \frac{\Delta N}{N}\sqrt{\left(\frac{\sigma(\Delta N)}{\Delta N}\right)^2 + \left(\frac{\sigma(N)}{N}\right)^2} \approx \left(\frac{\sigma(\Delta N)}{N}\right), \tag{2.14}$$

since we expect $\Delta N \ll N$. The uncertainty in ΔN is the same as that for N, governed by Poisson statistics i.e. $\sigma(\Delta N) = \sigma(N) = \sqrt{N}$. This gives the well-known expression of

$$\sigma(d_n) = \frac{\hbar}{2\alpha E T\sqrt{N}}. \tag{2.15}$$

From this it is apparent that to increase the sensitivity to d_n, all of α, E, T and N must be maximised. We must also bear in mind that, whilst this is the main limiting factor of current EDM experiments, this is only the statistical uncertainty: there will be sources of systematic uncertainty that will be addressed in the rest of this chapter.

2.2 Neutron Density, N

Beam experiments as they were in the 1980s became obsolete with the advent of significant improvements in the production of UCNs. Despite the larger count rates and higher electric fields attainable when using neutron beams, UCN experiments have the clear advantage of a greater interaction time between the neutrons and

the electric field, that would therefore magnify the effect of an EDM. This stems from the ease of storage of UCNs and their manipulability.

The first methods of UCN production relied on the extraction of the slowest particles from the tail of the Maxwell-Boltzmann distribution, exploiting the energy gains and losses of neutrons subject to a gravitational potential. At the limit-holding, Sussex-RAL-ILL nEDM experiment, neutrons encountered a liquid deuterium moderator near the reactor core, and those that emerged were guided up a vertical tube. These neutrons were directed towards a corotating nickel-bladed turbine, undergoing multiple reflections and losing momentum in the process. At the time, this was the most intense UCN source in the world, producing around $25\,\mathrm{UCN\,cm^{-3}}$ [48].

Theoretically more effective UCN generation has been proposed, that should yield UCN densities at least an order of magnitude greater than those possible via neutrons in thermal equilibrium with a moderator. Golub and Pendlebury described one such 'superthermal' technique exploiting the properties of superfluid ^{4}He (He-II) [49]. Excitations are produced in the superfluid as the neutrons enter and lose energy. The Feynman-Landau dispersion curve describes the energy-momentum relationship of these quasiparticles, with the linear, low energy component corresponding to phonons, and the higher energy part to rotons. Figure 2.3 shows that this dispersion curve coincides with that of the neutron at a single point, representing a neutron with wavelength 8.9 Å or energy $\sim 1\,\mathrm{meV}$. Such a neutron can deposit most of its energy in the superfluid by single phonon emission, downscattering to become a UCN; multi-phonon processes occur around 20-25% of the time, but still contribute positively to the UCN population. The reverse process of upscattering begins to be a problem at a bath temperature of around 0.7 K, so the superfluid is maintained at around 0.5 K, slightly below the peak UCN production rate. Then upscattering is highly suppressed by the appropriate Boltzmann factor, since the difference in energies between the UCN and cold neutron states is much greater than the surrounding helium temperature. In other words, the likelihood

of a UCN encountering a 12 K phonon in the superfluid becomes very small. At the H53 cold neutron beamline at ILL, a UCN production rate of $(0.91 \pm 0.13)\,\mathrm{cm}^{-3}\,\mathrm{s}^{-1}$ was measured [50]. This was the first confirmation of the theoretical production rate due to single-phonon emission, estimated to be $(1.19 \pm 0.18)\,\mathrm{cm}^{-3}\,\mathrm{s}^{-1}$ in these experimental conditions.

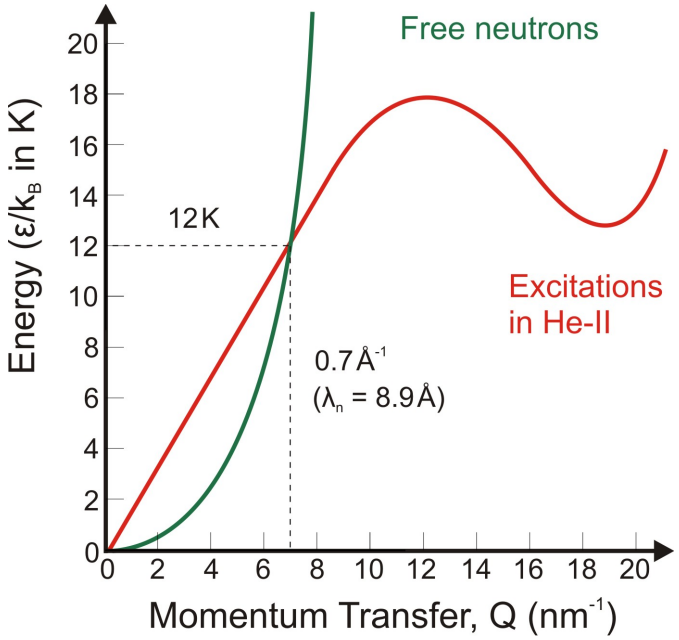

FIGURE 2.3: The dispersion curves for neutrons and phonon-roton excitations in superfluid helium are seen to cross at a 'critical momentum' corresponding to a neutron wavelength of 8.9 Å.

Using superfluid helium confers other advantages, having zero neutron absorption cross-section and zero magnetic moment for that of the neutron to couple to help preserve their polarisation. He-II also has a superior dielectric constant compared to vacuum so that in theory higher electric fields can be supported. However, commercial helium also contains a significant proportion of ^{3}He, which is not quite so desirable. ^{3}He has a magnetic moment and a large neutron capture cross-section when its spin is antiparallel to that of the neutron. The natural abundance of ^{3}He is around 1 part in 10^6, which if present in the experiment would reduce the UCN storage time to mere seconds. The helium is isotopically purified via passage through a superleak, whose main feature is its low porosity. Superfluid can flow easily through the superleak, whilst 'normal' fluid remains

mostly trapped at the other side. As ^{4}He and ^{3}He have substantially different lambda points (2.17 K for ^{4}He versus 2.5 mK for ^{3}He), the technique can be used to separate out the two isotopes. The level of ^{3}He contamination can be reduced by three orders of magnitude [51]. However, other impurities may be present in the superfluid helium that could impact on the maximum achievable electric field before breakdown occurs.

A key aspect of the above UCN production method is that, unless there were an effective way of retrieving the neutrons from the superfluid, the experiment would need to run in cryogenic conditions: this can be useful for a variety of reasons, as shall be mentioned throughout the text, but can also provide a technical challenge in its construction and operation. A second superthermal option that could avoid this would be to use a solid deuterium (sD_2) moderator. The concept is similar to that of superfluid helium, with the neutrons being downscattered by single phonon processes. A complication similar to that with ^{3}He contamination arises as deuterium exists as two spin isomers, the ortho and para forms, decided by the relative orientations of spins of the two nuclei in the molecule. The ortho state is the ground state and corresponds to a total spin of 0 or 2, and the para state a spin of 1, with the energy difference between the states being 7 meV. The exothermic conversion from para to ortho forms is more than enough to upscatter a UCN. Moreover, even in its lowest rotational state, a para molecule is excited with an energy equivalent to about 90 K. The lifetime of a UCN within the sD_2 is therefore governed by the fraction of para molecules. The ratio of ortho to para molecules at room temperature is 2:1, as all spin states are evenly populated and there are two times as many ortho as para states. At lower temperatures, it becomes more favourable for the deuterium to exist in ortho-form, and at a temperature of 20.4 K, the ortho content should be 98% [52]. However, the relaxation time to this percentage is on the order of a few percent per month [53]: thus to achieve a workable ratio, a catalyst is needed. A converter, such as hydrous iron (III) oxide is used to prepare the deuterium in the ortho state. Using this method,

[54] was able to achieve an ortho content of 95% at 5 K and report a maximum UCN density of $145 \pm 7\,\mathrm{UCN\,cm^{-3}}$ in a storage bottle, but noted that an increase in temperature impacts on the production rate via increased thermal upscattering. Nevertheless, the technique is promising, although is yet to be realised in a room temperature experiment.

2.3 Storage Time, T

UCNs have wavelengths longer than the internuclei spacing in a solid, and thus interact macroscopically with a material. The effective potential seen by the neutron is the volume average of the repulsive pseudopotential associated with each of the scattering nuclei. This can be represented as

$$V = \frac{2\pi\hbar^2}{m_n} N\alpha, \tag{2.16}$$

where m_n is the neutron mass, N is the number density of the nuclei and α the coherent scattering length. The behaviour of UCNs on encountering a surface is analogous to that of light, with total reflection occurring for an energy of incidence lower than this characteristic potential of the material. In the field of neutron optics, it is also possible to define an index of refraction for the material, as well as a critical velocity, v_c, which represents the maximum velocity a neutron can have and still be reflected from the material.

However, if we extend the picture beyond a classical potential barrier, there remains some quantum mechanical probability that the neutron can penetrate into the wall material. This leads to UCN losses via absorption, where neutron capture occurs, or via inelastic scattering, where the neutron is upscattered, gaining energy by collision with a vibrating nucleus at the greater, wall temperature. The

total effective potential, V_F, known as the Fermi potential, then becomes

$$V_F = V - iW = \frac{2\pi\hbar^2}{m_n} N(\alpha + i\alpha_l), \tag{2.17}$$

where $\alpha_l = \frac{\sigma_l k}{4\pi}$ is the equivalent scattering length due to losses, dependent on the loss cross section per nucleus, σ_l. A useful parameter to define is therefore the loss coefficient, $\eta = \frac{W}{V}$, which gives a measure of the loss probability.

The material of choice should therefore have a suitably high V_F, whilst maintaining a low η. The typical UCN energy is around 100-300 neV and thus materials with V_F around this magnitude are desirable to contain the UCNs whilst providing a window for higher energy neutrons to escape. Commonly used materials include beryllium and beryllium oxide, diamond-like-carbon (DLC) and nickel, whose properties are given in the Table 2.1.

Material	V_F/neV	v_c/ms^{-1}	η
Beryllium	252	6.89	(2.0-8.5)$\times 10^{-4}$
Beryllium Oxide	261	6.99	(1.0-3.0)$\times 10^{-5}$
Nickel[1]	252	6.85	5.1×10^{-4}
Diamond	304	7.65	1×10^{-7}
Graphite	180	5.47	9.5×10^{-5} [55]

TABLE 2.1: The Fermi potential, critical velocity and loss coefficient for various viable wall materials. Those of DLC should lie somewhere between diamond and graphite, depending on the ratio of its sp^3 to sp^2 content (sources: [56], [57], [58]).

Each of these has its own problems. Nickel has a high neutron absorption cross section and therefore a larger η, so is only useful in cases with few wall collisions. The toxicity of beryllium and beryllium oxide lead to complications in their manufacture and disposal. Lastly, the properties of DLC are highly dependent on its exact structure, which can make it difficult to characterise, and its cost of synthesis is not inconsiderable. Nevertheless, it is an attractive choice for wall coatings in fundamental UCN experiments, being established and widely used in industry. An

[1] In its natural form. Pure ^{58}Ni has the largest V_F known: 335 neV and a critical velocity of 8.14 ms^{-1}.

sp^3 (the bonding structure of natural diamond) content of 60% is thought possible, which would yield a Fermi potential of around 260 neV [59]. A typical Ramsey cell would consist of a ceramic or quartz ring, coated on the inside with a material like DLC, and closed at either end by similarly treated electrodes.

Aside from wall interactions, there are other loss mechanisms detrimental to the possible storage time of UCN in a bottle. For instance, any gaps or holes will lead to leakage of UCN. Of course, the storage time is ultimately capped by the lifetime of the free neutron, $\sim 886\,\mathrm{s}$ [60]. There will also be other loss pathways specific to the experiment. In the case of one employing a superthermal superfluid helium source, we have already discussed the possibility of UCN absorption due to ^{3}He content, and upscattering by phonons if operating at a temperature close to 0.7 K. Each of these mechanisms can be associated with a loss rate: $\tau_{\mathrm{wall}}^{-1}$ due to wall interactions; $\tau_{\mathrm{leak}}^{-1}$ for leakages; τ_{β}^{-1} for the neutron lifetime via β decay; $\tau_{^3\mathrm{He}}^{-1}$ for ^{3}He absorption; and $\tau_{\mathrm{phonon}}^{-1}$ for phonon upscattering. These contributions can be combined to give the overall storage time, τ, with

$$\frac{1}{\tau} = \frac{1}{\tau_{\mathrm{wall}}} + \frac{1}{\tau_{\mathrm{leak}}} + \frac{1}{\tau_{\beta}} + \frac{1}{\tau_{^3\mathrm{He}}} + \frac{1}{\tau_{\mathrm{phonon}}}. \tag{2.18}$$

Similar considerations apply for any neutron guide system, as these loss mechanisms will also affect neutrons during transport. Further, the length of the superthermal source needs to take into account the mean free path of cold neutrons in the superfluid before downscattering, which is some 10 m. The source in CryoEDM was 3 m in length, which would see around 70% of the incoming neutrons lost before they reached the rest of the guide system.

2.4 Visibility, α

The visibility is the product of the polarising and analysing efficiencies, and characterises the peak-to-peak height of the Ramsey resonance curve, or in other words, how 'visible' these fringes are. This encompasses not only the effectiveness of the polariser itself, but also how well polarisation is maintained until detection. Akin to storage time, loss mechanisms result in depolarisation of neutrons during transport and confinement.

In discussing the Fermi potential, we neglected an extra term that would be present if the material were magnetisable: $V_F \rightarrow V_F \pm \vec{\mu} \cdot \vec{B}$. If the material is highly magnetised, it can be made invisible to one spin state, with a very small angle of acceptance for reflection. This property can be used to great effect in a supermirror arrangement. A supermirror is formed from multiple, alternate layers of two materials, with varying thickness. In neutron transport, the idea is to increase the critical angle of the guide wall, by some factor m compared to pure ^{58}Ni, to maximally retain UCN in the guide. Instead for a polariser, each other layer is magnetised to trap neutrons of one spin state. The polariser for CryoEDM consisted of iron deposited on bands of silicon substrate and was around 3 m in length, to catch neutrons at a variety of angles. It was set between a yoke magnet designed to provide a field of around 300 G, to preserve the magnetisation of the iron. The design is efficient as there is a strong contrast between the parallel and antiparallel spin states, with the potential seen by the latter at the iron interface being almost identical to that at the silicon interface: a polarisation of 95% was obtained for 9 Å neutrons. Other designs look at using other pairs of materials (e.g. cobalt and titanium) and explore the possibility of using remnant magnetisation of the layers, such that no active magnetic fields need to be applied during operation [61].

Alternatively, the UCN can simply be polarised if exposed to a high magnetic field. As the neutron magnetic potential is $\approx$60 neV T^{-1} and UCN have energies

<300 neV, a magnetic field of 5 T should be more than sufficient to achieve 100% polarisation [62]. However there are a number of difficulties, not least how to generate a stable field of that magnitude. The proposal is to use a superconducting solenoid that should have a persistent supercurrent and therefore a constant field provided it is maintained at the required temperature. It is not only temporal invariance that is of importance, but also spatial homogeneity, since any significant gradients will only serve to accelerate the neutrons. Finally, there is the problem of how to contain any stray fields such that they do not disturb the main EDM experiment. A large separation is likely required, which then exacerbates the problem of depolarisation and UCN losses in the guide system.

As with storage time, depolarisation can result from a number of different factors. Interactions with the guide or surfaces, whose Fermi potential may have a small but non-negligible magnetic term, could lead to wall relaxation, especially if magnetic impurities are present. The depolarisation probability per bounce can be summarised in the spin-flip probability, β. Experiments have found that this coefficient is around $(2-3)\times10^{-5}$ for beryllium, and 3×10^{-6} for DLC, which again illustrates the superior qualities of the latter material [63].

As a consequence of the parallels between the Ramsey resonance method and a standard nuclear magnetic resonance (NMR) technique, depolarisation is often quantified in terms of two spin relaxation times that are commonly used in discussions of NMR: T_1 and T_2 in the longitudinal and transverse directions respectively. T_2, sometimes called the spin-spin relaxation time, is the more often referred to in conjunction with the Ramsey cycle. T_2 represents the dephasing that occurs due to interactions between the neutron spins, with each sampling a different local magnetic field. Each spin accumulates a different phase, which results in decoherence in the neutron ensemble. T_2 is thus sensitive to variations in the magnetic field that can occur across the bottle. For a monoenergetic set of neutrons, if they are assumed to sample the field randomly, the phase distribution would be broadened, but still centred on the same average phase. By considering the effect of gravity

however, the average height of the neutrons during the cycle will depend upon their energy, and thus an energy dependent, relative phase decoherence occurs.

T_1 describes the decay of the component of the spin parallel to the direction of the external field, and is thus a useful parameter to characterise depolarisation during neutron transport, where a holding field is typically used to maintain the initial spin orientation. The main requirement of this field is that it varies slowly enough that the neutron spins can adiabatically follow the changes: this is a considerable challenge given that the field must be reduced from its value at the polariser to that of the static field in the Ramsey cells, generally a 10^4 drop in magnitude from hundreds of gauss to tens of milligauss. The neutrons will move ballistically down the guide tube, and thus inhomogeneities in the magnetic field across the cross-section of the guide, which will enhance the decay of this polarisation component, need to be avoided. The issue is discussed in more detail in Chapter 5 in relation to possible field distortion in the guide system caused by the ferromagnetic materials in SQUIDs.

The last influence on the visibility is the detection efficiency. Neutron detectors are generally based around so-called neutron converters. These are nuclei of element isotopes which have a particular affinity for neutron capture, and which subsequently decay spontaneously. There are relatively few viable species, but those which have been used include ^{3}He, ^{6}Li and ^{10}B. A footprint of a neutron interaction could be any scintillation released, or ionisation caused by charged decay products.

The standard UCN detector consists of a proportional gas tube, which contains gaseous ^{3}He that acts as both the converter and the counting gas, which is ionised by the reaction products. Such a detector is not suitable for a cryogenic experiment since the detection principle would not hold if the helium were in liquid form, and because a thicker entrance window would be needed that would further attenuate and scatter any incoming neutrons. This having been said, ^{3}He **is** used as a

neutron converter by the SNS experiment, albeit as a secondary function of their comagnetometer, acting as a spin analyser [64].

Gas detectors have been superseded by those which incorporate solid converters as they can provide a greater count rate acceptance. The higher density of converter nuclei easily compensates for any reduction in neutron absorption cross section. The detectors proposed for CryoEDM were ^{6}LiF coated silicon sensors designed to measure the ionisation signal caused by the reaction products, ^{3}H and α (see Figure 2.4). A study was carried out by Lynch coupling the ^{6}LiF to a scintillator rather than a semiconductor, collecting light with a photomultiplier tube [65]. She found ZnS to be a suitable scintillator in cryogenic conditions, and that a major factor in improving the sensitivity to UCNs was the thickness of both layers: that of the ^{6}LiF to balance neutron capture and range of the conversion products, and that of the ZnS to limit the γ background.

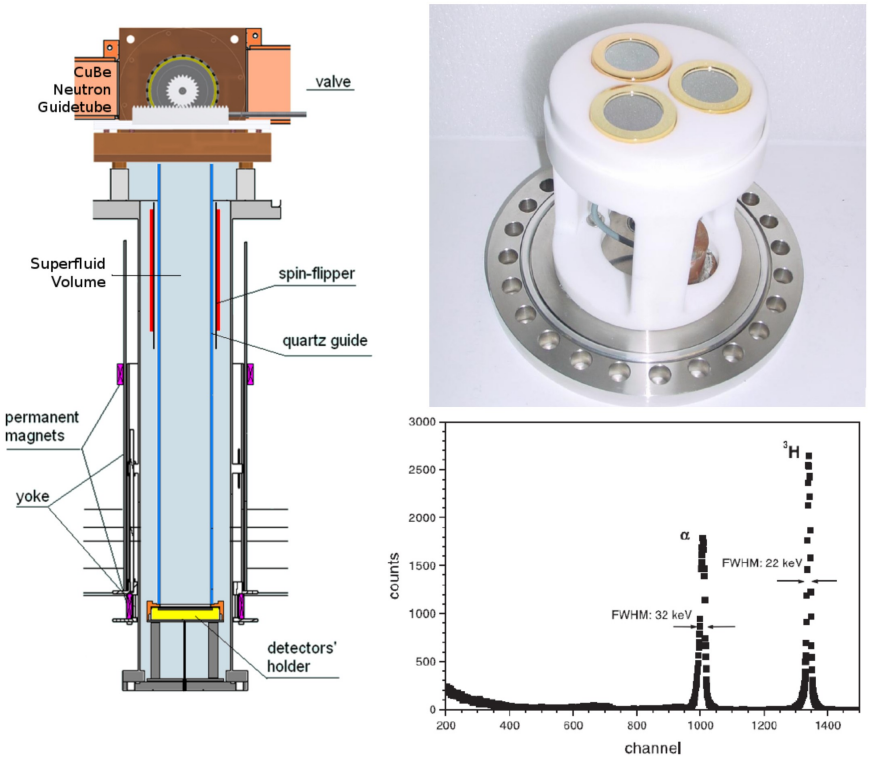

FIGURE 2.4: From left, clockwise: a schematic of the UCN detector scheme for CryoEDM; the three detectors mounted with differently magnetised foils to act as a spin analyser; a sample spectrum from a detector, obtained with a 10 μ g cm^2 deposited layer of ^{6}LiF converter [66]. The peaks for the products of the reaction $^6Li + n \rightarrow ^3 H + \alpha$ are clearly visible.

A fundamental limitation of both designs is that, since the products are released back-to-back, only half of them are detected. One way of addressing this problem is to sandwich the converter in between the detector, provided that the detector has a low Fermi potential such that it is essentially transparent to the UCN. A UCN detector formed of a ^{6}LiF converter placed in between Si wafers was developed by the group at PNPI, and showed an analysing power up to 75% in the main part of the UCN spectrum [67]. Another popular option is doped scintillators: again, a bilayer structure can be used, one half of which is lithium enriched, where the conversion will take place, and the other of which is lithium depleted. Glass scintillators (GS), which contain e.g. SiO_2, Ce_2O_3, Li_2O and other oxides, are widely used, with a thinner layer providing better neutron/γ discrimination. GS10 scintillators are the most promising, demonstrating an efficiency close to that of gas detectors, with a faster counting rate [68].

Lithium is typically used instead of boron because, despite its smaller neutron capture cross section, the daughter nuclei have higher energy and thus are more likely to penetrate into the detector and cause ionisation/scintillation. However, so long as the signal can be amplified, boron could be the better choice. CASCADE [69] is a detector that has been developed by Heidelberg and uses Gas Electron Multiplier (GEM) technology [70]. The detector consists of a cascade of GEM foils each with a deposited layer of ^{10}B. When the conversion products (^{7}Li or α) reach the holes of the GEM foils, any ionisation they cause is intensified by gas amplification. Around 10^5 electrons are transferred to the readout per conversion event. The same efficiency issue related to the thickness of the converter layer prevails, with the fine art of balancing the probability of neutron capture with that of the escape probability of the reaction products.

The last question is how to transform these into spin sensitive detectors. A spin analyser can be made by taking advantage of the spin dependence of the Fermi potential of a magnetised material. The analysing foil, typically magnetised iron, is placed upstream of the detector and lets through only one spin state. CryoEDM

used three detectors that would receive approximately the same neutron flux, two beneath a foil polarised to allow detection of one of the two spin states, and a third with no foil to act as a control, counting both spin states. The neutrons of the spin state to which the detectors are insensitive can be stored above the foil as the others are measured, then converted using an adiabatic spin flipper and subsequently detected, allowing for sequential analysis of the two states. The biggest problem with the setup is losses, both in polarisation and UCN numbers. Further, UCN which do not infiltrate the analysing foil will bounce back, and the likelihood of them reinteracting with the foil could be slim without the correct geometry. A novel idea to improve the flux of neutrons reaching the detector is centred around a U-shaped design with two arms, one with a detector for the spin up state and the other for the spin down [71]. The roof is coated with a high V_F material to reflect rejected neutrons to the opposite detector, thus maximising the UCN catchment.

2.5 Electric Field, E

A maximal electric field is desired both to minimise σ_{d_n}, and to enhance any frequency shift due to an EDM. The circuitry to generate this field is generally of simple design, commonly consisting of a Cockroft-Walton generator to act as a voltage multiplier, with suitably rated, high voltage (HV) lines connected to a feedthrough to deliver the voltage to one of the Ramsey cell electrodes. Additional considerations include a cryogenic HV feed to limit thermal conduction in low temperature experiments, such as the PTFE covered, stainless steel feed designed by Karamath for CryoEDM [72]. The electrodes are manufactured in a Rogowski profile, which ensures that the largest field is in the gap between the electrodes, and that it has a high degree of uniformity [73] (see Figure 2.5). They are usually made of stainless steel, though studies have been conducted into non-metallic electrodes with a thin metal coating that would reduce Johnson noise, which could affect

the sensitivity of the magnetometers [74]. The electrodes are electropolished to remove asperities on their surfaces that can form points of high local field, where breakdown is more likely to occur, and then coated with a high Fermi potential to reflect the UCNs. In vacuum at room temperature, electric fields upwards of $100\,\mathrm{kV\,cm^{-1}}$ should be easily achievable with such apparatus, primarily limited by field emission from the electrodes. One significant challenge is how to create a comparable field upon the introduction of the Ramsey cell and the liquids or gases contained therein.

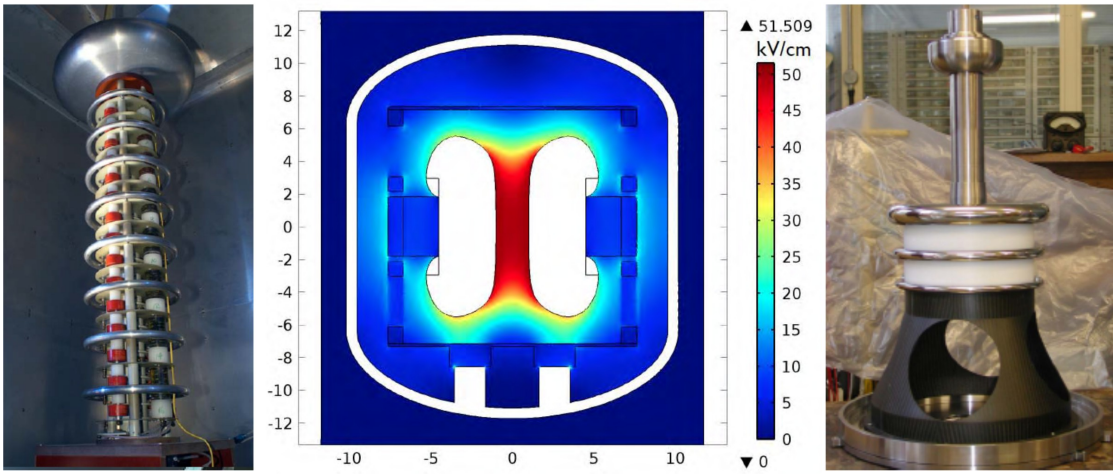

FIGURE 2.5: Designs to maximise the electric field and minimise the systematic effects derived from it. From left: the high voltage feed constructed for the limit-holiding experiment, with a Cockroft-Walton generator; finite element analysis of the electric field for the SNS test setup [75], showing the relative uniformity of the field between the Rogowski-shaped electrodes; the double Ramsey cell for CryoEDM, where the top electrode is to be charged by the high voltage feed, keeping the bottom cell a neutral cell. A neutral cell is useful as a cross-check to confirm the presence or absence of systematics related to the electric field.

One big issue is how to affix the cell between the electrodes. Grooves are etched into the electrode surfaces, which detract from the ideal electrode geometry. The location where the spacer and an electrode meet, the cathode-insulator junction, can be identified as a weak point where the field is more concentrated and thus electric breakdown is more likely to occur. Breakdown events can cause damage to the materials: craters can form on both electrodes, which leads to deterioration in their performance and a greater likelihood for further discharges and repeat

breakdowns; carbon tracks can form down the sides of the ceramic spacer that could degrade the neutron storage time [76].

When the neutron volume is not evacuated, further pathways for breakdown become available. Electric discharge in gases occurs due to ionisation. The occasional spark will happen at low voltage because of background radiation. A liberated electron can be accelerated by the field and with a sufficient gain in energy, can cause a chain of ionisation events that lead to an electron avalanche across the chamber. At the much higher voltages of interest, total ionisation of the gas occurs as it becomes conductive beyond some breakdown voltage, leading to what is known as glow discharge, since the resultant plasma luminesces. The behaviour of gases under the application of an electric field is underpinned by Paschen's law [77]. It turns out that the gas cannot be defined by a constant dielectric strength, but that this is dependent upon the pressure and the spacing between the two electrodes. The shape of the Paschen curve, which denotes the relationship between the breakdown voltage and the product of the pressure and the electrode spacing, is similar for all gases (see Figure 2.6). Much of this rests upon the interplay between the mean free path of an electron and its collision rate with gas particles: more collisions would lead to more ionisation, provided that the electron has accumulated enough energy; a larger mean free path allows the electron to build up greater energy to cause further ionisation, but conversely means collisions are less likely. Of concern to current EDM efforts are the properties of ^{199}Hg and ^{129}Xe as they are to be used for comagnetometers in future room temperature EDM experiments. A greater density of these gases i.e. a higher partial pressure is more desirable to reduce statistical errors from their readout, but could lead to more interactions with the UCN and greater likelihood for breakdown. As xenon is a noble gas, it has a much higher first ionisation energy than mercury and thus has a higher breakdown voltage in general. Measurements by the RCNP/TRUIMF experiment have shown that breakdown only occurs at pressures greater than 10^{-2} Torr [78].

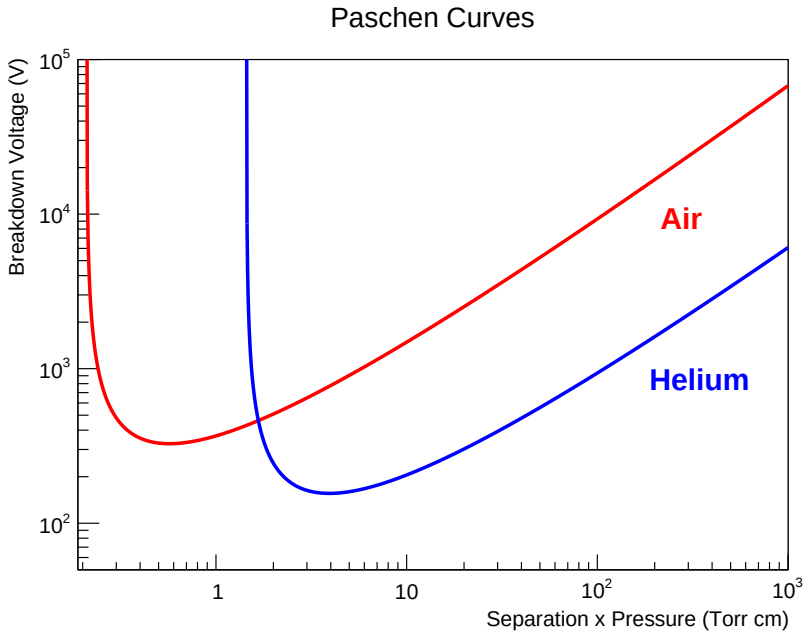

FIGURE 2.6: Paschen curves for air and gaseous helium showing the variation of the breakdown voltage with the product of the gas pressure and electrode separation. The relationship is characterised by the equation $V_B = \frac{apd}{\ln pd+b}$, where p is the pressure, d the separation, and a and b are constants specific to each gas. Tests have shown that He-II appears to follow a similar expression, despite the helium being in a different phase [72] [76].

In contrast, electric breakdown in liquid helium is poorly understood. The theoretical breakdown voltage of He-II is thought to be higher than even vacuum, if one considers the expected mean free path for electrons [75]. It is widely accepted that breakdown occurs via a gas 'streamer' that is formed in the liquid phase. Gas particles could nucleate around an asperity, since it has a higher local current density and therefore is a point of field emission. Excessive heat transfer can lead to a bubble forming, which as it increases in size has a greater hydrostatic pressure and surface tension counteracting its growth. The bubble tends towards a conical shape that can eventually form a column between the two electrodes through which breakdown occurs. Impurities in the liquid exacerbate the problem, being more likely to become charged and then accelerate to an electrode, all the while causing local heating. Despite these potential setbacks, research and development efforts by the SNS experiment have reported demonstrable progress, initial results indicating that 100 kV/cm fields can be applied across a 1 cm gap between two 12 cm diameter electrodes for a range of liquid helium pressures [79].

Aside from breakdown, characterised with an intense spark of energy, there will also inevitably be some long-term movement of charge in the form of a leakage current, as the two electrodes effectively form a capacitor. Leakage currents on the order of nA are not uncommon, and will gradually diminish the possible voltage and thus the highest electric field. This effect is negligible over the time scale of a measurement: the much greater problem is any possible spurious magnetic field the current can cause. An axial current will generate a field in the transverse direction, whereas a current which spirals or has any tangential component will produce a field along the static B_0 field direction. This has the potential of creating a false EDM, which can be estimated as

$$|\vec{d_n'}| \propto \frac{If}{r|\vec{E}|}, \tag{2.19}$$

where f is the fraction of the circumference that the current travels, I is the size of the current, r is the bottle radius and $|\vec{E}|$ is the magnitude of the electric field applied across it.

Whilst any gradients in the magnetic field have obvious ramifications on an EDM determination, as is further discussed in Section 2.6, the homogeneity of the electric field is also of importance. The major systematic uncertainties arise from the $\mathbf{v}\times\mathbf{E}$ effect. Neutrons experience an additional magnetic field in their reference frame as they travel an electric field. Explicitly the magnetic field, $\vec{B}'$, seen by a neutron is given by

$$\vec{B}' = \gamma\left(\vec{B} - \frac{\vec{v}\times\vec{E}}{c^2}\right) - \frac{\gamma-1}{v^2}(\vec{B}\cdot\vec{v})\vec{v}, \tag{2.20}$$

i.e. the Lorentz transformation from the laboratory to the neutron's frame of reference. With the UCN moving predominantly perpendicular to $\vec{B}$ with a speed $v \ll c$, the dominant correction is given by

$$\delta B' = \left|\frac{\vec{v}\times\vec{E}}{c^2|\vec{E}|}\right|. \tag{2.21}$$

If the electric field is not perfectly aligned with the magnetic field, any component which is mutually orthogonal to B_0 and the given neutron's velocity will provide an enhancement of the static field for that neutron, which will affect its precession frequency. Considering the entire ensemble of neutrons, movement of its centre of mass would register a false EDM signal of magnitude

$$|\vec{d_n'}| = \frac{\mu_n(\vec{v}_{\mathrm{CoM}} \times \vec{E}) \cdot \hat{B}_0}{c^2}, \tag{2.22}$$

with $\hat{B}_0$ the unit vector in the direction of the static magnetic field. The neutron centre of mass can acquire some small net translational velocity in the vertical direction due to Brownian motion of impurities on the walls causing 'warming' of the neutrons. It may also have some rotational motion due to circulation currents on filling. The clear advantage of using UCN over a neutron beam is the minimisation of this effect because of their slower speed and, if they are bottled, through the changing of their direction of travel upon each reflection from the walls. To further combat this systematic, homogeneity of the electric field over the bottle is imperative. Significant variations in $\vec{E}$ are only likely to occur near the edges of the wall and electrodes because of departures from the Rogowski geometry and material effects (see Figure 2.5). The estimated size of this effect is therefore small, but not negligible: for CryoEDM, $|\vec{d_n'}| \approx 0.3 \times 10^{-28}\, e\,\mathrm{cm}$.

A last aspect to consider is that of the changing refractive index of a material exposed to a large electric field, a phenomenon known as the Kerr effect [80]. This needs to be taken into account for any optical monitoring of the species inside the bottle, which is the modus operandi of some of the comagnetometers proposed by EDM experiments.

2.6 Magnetic Fields

In the example given in Section 2.1, we saw that a perturbation in the static magnetic field, B_0, of 2 fT over a measurement can result in a similar frequency change as that caused by an EDM of magnitude $d_n \sim 10^{-26}\, e\,\text{cm}$. A larger change can be tolerated if an EDM is searched for over a number of measurements, but nevertheless, this drives a constraint that $\vec{B_0}$ should be small and kept stable – a magnitude of 1-5 μT is commonly used. Moreover, as we shall see, it should be relatively homogeneous to limit systematic effects that could induce a false EDM.

There are a number of ways to generate $\vec{B_0}$ so that it is primarily oriented in one direction, as is required for the Ramsey measurement. Primarily, a spherical coil, or in cylindrical geometry, a solenoid or cosine-theta ($\cos\theta$) coil is used. A spherical coil offers the greatest fiducial region, as a fraction of its total volume, for a specified level of homogeneity compared with the cylindrical designs with similar numbers of loops [81]. The major drawback, aside from the greater complexity in its manufacture, is that any aperture to access its interior is limited to the diameter of the smallest winding. The $\cos\theta$ coil outperforms a solenoid, with the field generated perpendicular to instead of along the cylindrical axis, and thus with the Ramsey cells also lying with their primary axis transverse to this direction. Both provide a more uniform field in their interior with a larger aspect ratio (i.e. length to radius proportion). A suitable magnetic shield would be coaxial with the cylinder, which would provide a greater shielding factor in the transverse rather than axial direction. However, a considerable advantage of using a solenoid over a $\cos\theta$ coil is that, with the Ramsey cells coaxial to the shield, it is easier to configure a system to deliver the neutrons into the cell and apply a high voltage to an electrode. Further, in a cryogenic experiment, not only could the coils themselves be made out of a superconductor, set up with a persistent current that should theoretically suffer little loss, but the solenoid could itself be wound around a coaxial superconducting shield with a lower critical temperature that could be

used to 'freeze' in the axial field. Needless to say, both the field generating coils and the shielding need to have diameters much greater than the Ramsey cells such that the field in the neutron volume is relatively homogenous.

Shielding is of primary concern to help maintain the field at the chosen level: the Earth's field alone is around 40-60 μT, which would easily swamp the static field. Passive shielding is constructed out of high permeability materials, most notably mu metal, metglass and permalloy, which concentrate the magnetic field lines within them, letting a lower density of flux lines permeate to the interior. As we shall see in Chapter 4, these shields are well-suited for the magnetic environments that are expected for most EDM experiments, characterised by low-frequency, quasi-static fields. We will define a shielding factor as the ratio of the field exterior to the shield to that at its centre: this is in the context of a static, uniform field, whereas in reality, the shielding factor will depend upon the frequency of the external source and its exact field distribution.

A big issue for the performance of the shield is its uniformity, since any patches of varying permeability will provide a different shielding factor at their immediate interior. This could not only result from impurities in the material itself, but also from any temperature gradients which exist across its surface. Remnant magnetisation can also suppress the shielding factor, since the ability of the shield to provide a path for field lines will decrease the closer it is to saturation. The shield could be degaussed by heating it to the Curie point, which is not readily achieved, or by applying an AC field with gradually diminishing amplitude. The latter will always lead to some small residual magnetic induction, since magnetic hysteresis implies that an infinite number of cycles would be needed to reduce this to zero in zero magnetising field. Degaussing is generally done in-situ, since transporting the shield after the process will expose it to a new field environment in which it will become magnetised once more. Induced magnetisation, which is linked with the E field direction, will also result in a false EDM.

Even with treatment, a ferromagnetic shield may only provide a shielding factor of a few hundred, depending on its quality and its dimensions. This falls short of the upwards of 10^5 shielding factor needed by most experiments to curb ambient fluctuations down to the fT level and below. Multiple layers of shielding can vastly improve the shield's effectiveness if they are of appropriate thickness, radius and position [82]. For both spherical and cylindrical shields, the transverse shielding factor for a single layer in a uniform field can be expressed as

$$S^T = A\frac{\mu_r t}{r} \begin{cases} A = \frac{1}{2}, & \text{cylinder} \\ A = \frac{2}{3}, & \text{sphere} \end{cases} \tag{2.23}$$

in the limit of $\mu_r \gg 1$ and $r \gg t$, with μ_r the relative permeability of the shield, r its inner radius and t its thickness. An analytic approximation for the transverse shielding factor of well-separated multiple shields, again assuming a high μ, is given by

$$S^T = \prod_{i=1}^{n-1} S_n^T S_i^T \left[1 - \left(\frac{r_i}{r_{i+1}}\right)^2\right], \tag{2.24}$$

where S_i^T is the transverse shielding factor of an individual layer as in Equation 2.23 [83]. The axial shielding factor for a cylindrical shield is more difficult to estimate: for workable dimensions, it is smaller than the transverse shielding factor, being related to S^T but suppressed by a geometric factor that is dependent upon its aspect ratio. A similar expression to Equation 2.24 can be used to describe the axial shielding factor of a multi-layered shield in terms of that of each layer. These theoretical predictions are limited to the scope of ideal geometries and uniform fields. Efforts, both numerical and analytical, have looked to explore how the shielding factor is altered with the addition of other features, such as end caps or stovepipes, that are necessary for installing other apparatus. Further studies have looked to generalise external fields as multipole expansions [84].

Active shielding is sometimes employed to lessen the fields seen by the passive shielding. A set of coils can generate a field to counteract the ambient field, with

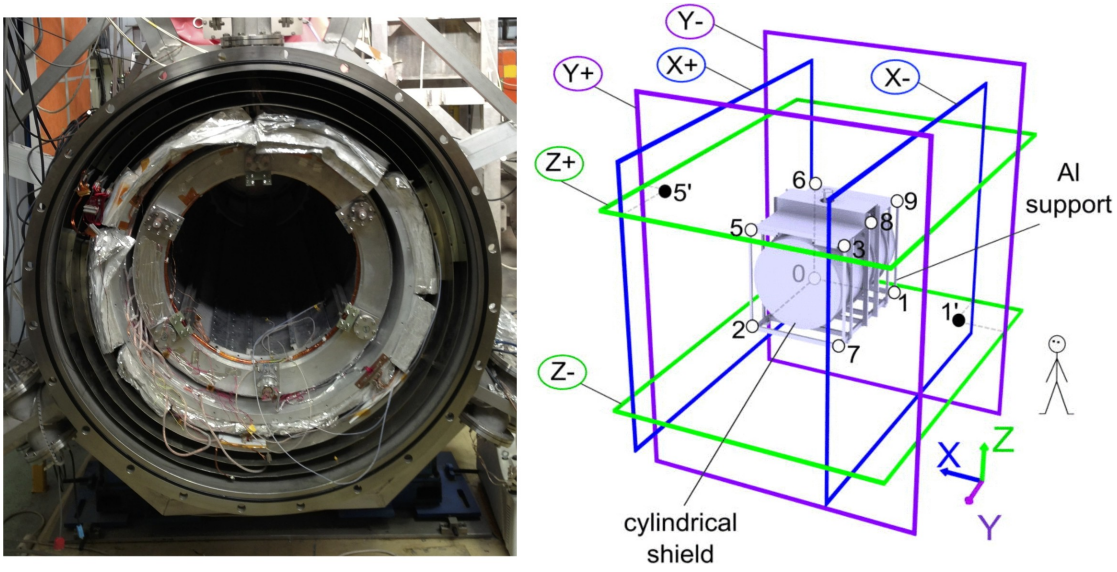

FIGURE 2.7: Passive (left) and active (right) magnetic shielding approaches. Left - a look inside the cryostat in CryoEDM shows three layers of mu metal shielding surrounding the central, superfluid containment vessel. Right - schematic of the active compensation system to be used at PSI [85].

readings of the field at various positions fed back to the coils to adjust their currents suitably. Both the experiments at PSI and FRM-II have such active compensation systems [85] [86], with a cube of three orthogonal Helmholtz coils situated around the experiment, and fluxgates at the corners (see Figure 2.7). A shielding factor of ~ 10 is thought achievable.

Whether shielding in the axial or transverse direction is more important will depend on the configuration of the Ramsey cells and other components to be housed within them. Controlling gradients in the direction of the static field is important since differences will directly impact on the precession frequency of the neutrons and therefore on the derived value of d_n. However, as the sensitivity to d_n is pushed further, inhomogeneities in the radial direction must also be suppressed since these contribute to another systematic effect known as the geometric phase effect (see Figure 2.8). A field gradient parallel to the static field is tantamount to a radial gradient, since in a sourceless region, $\nabla \cdot \vec{B} = 0$ and thus, assuming

azimuthal symmetry

$$\frac{1}{r}\frac{\partial(rB_r)}{\partial r} = -\frac{\partial B_z}{\partial z} \quad \rightarrow \quad \vec{B}_r = -\frac{1}{2}\frac{\partial B_z}{\partial z}\vec{r}. \tag{2.25}$$

The transverse field, B_{xy} seen by the neutrons is thus a combination of this field, B_r, and any field associated with the $\vec{v} \times \vec{E}$ effect, B_v, whose origins are detailed in Section 2.5. $B_{xy} = B_r + B_v$ is seen to be continuously rotating in the reference frame of the neutrons, due to their spin precession and motion: in the case of uniform field gradients and the 'garland' orbit in the x-y plane, skirting along the wall edge, the frequency of rotation of the magnetic field would be $\omega_r = \frac{v_{xy}}{r}$, with v_{xy} the neutron's speed and r the radius of the Ramsey Cell. During a 'revolution' of the magnetic field, the neutron spin direction will not only have accumulated a dynamical phase from its spin precession, but also a geometric or Berry phase associated with the solid angle subtended by the resultant magnetic field. This translates to a shift in the Larmor precession frequency as a departure from simply the ω_0 that the neutron would have had had there solely been a field in the z-direction. Often referred to as the Ramsey-Bloch-Siegert shift, this difference depends on the squared magnitude of B_{xy}. Therefore, there is a term proportional to B_r^2 that will exist in the absence of an electric field; one that scales with B_v^2 and therefore E^2, that is a second-order $\vec{v} \times \vec{E}$ effect; and finally, one that goes as $B_r B_v$ which thus causes shifts linear in E. The shift and therefore the false EDM that manifests will be related to the neutron trajectory and the relative sizes of ω_r and ω_0. UCN nearly always fall into the regime where $\omega_r < \omega_0$, and so an analytical calculation for the garland orbit, assumed as an extreme case, yields [87]

$$|\vec{d_n'}| = -\frac{\hbar}{4}\frac{\partial B_0}{\partial z}\frac{v_{xy}^2}{B_0^2 c^2(1 - \omega_r^2/\omega_0^2)}. \tag{2.26}$$

The emergence of this systematic in addition and in relation to the already well-known $\vec{v} \times \vec{E}$ only strengthens the need to address the latter. It should also be noted that the geometric phase effect will also affect comagnetometer particles, which

are in a fully non-adiabatic regime of $\omega_r > \omega_0$. Further, differences in the orbits of the comagnetometer particles and the neutrons, resulting from a translation in the centres of mass due to gravity or otherwise, will lead to complications in determining the geometric-phase effect for each species [88].

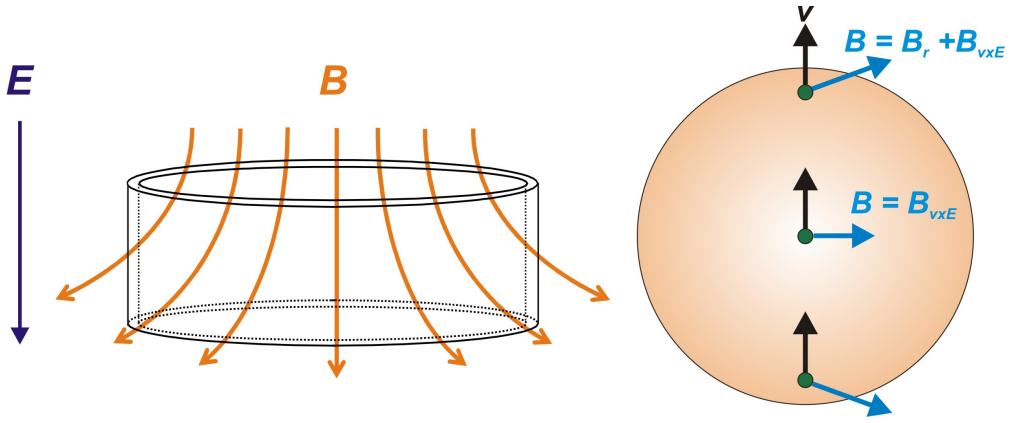

FIGURE 2.8: A gradient in B_0 in the axial direction will cause a radial component in the magnetic field in a sourceless cell. The geometric phase effect will enhance the spurious field seen in the reference frame of a moving neutron, in addition to that derived from the $\vec{v} \times \vec{E}$ effect. Moreover, this will be different dependent on where the neutron is in the cell: in the example illustrated, there is no radial component at the centre and thus no extra B component, but at the edges of the cell, the GPE contribution will perturb the net B field.

2.7 Magnetometry

Magnetometry was a relatively late development in neutron EDM projects, coincidental with the transition to bottled UCN experiments in the 1980s. Spatial variations in the magnetic field can be identified by field mapping the apparatus, but the long storage times of UCN make them also susceptible to temporal changes that therefore must be monitored. The requirements for a magnetometer in early ventures were not as stringent as they are now. The 1984 and 1990 experiments at ILL used optically pumped rubidium magnetometers placed in close proximity around the neutron volume, that were assumed to be similarly affected by any field fluctuations from outside of the shielding [89] [90]. The main concern was field changes generated internally, in particular from leakage currents that would be correlated with the electric field direction. The signals were converted into

'magnetometer EDMs' to compare with that derived for the neutron, to isolate the effects of any spurious magnetic sources in the vicinity of Ramsey cells. This treatment was consistent with obtaining an EDM limit of order $10^{-25}\,e\,\mathrm{cm}$. With current experiments aiming for a sensitivity of $\sim 10^{-27}\,e\,\mathrm{cm}$ or better, systematic effects that were previously discounted, such as the $\vec{v}\times\vec{E}$ effect, now become significant. Fluctuations in the magnetic field are now the dominant source of major systematic uncertainties in an EDM experiment. Even with substantial technological advances in shielding and field generation, inhomogeneities are inevitable at some level, inherent in the static magnetic field and arising from any external disturbances which are not perfectly shielded. Characterisation of the apparatus and field environment are insufficient to determine whether variations will be below the specified magnitude over the entire measurement. Thus, precision magnetometry has become even more crucial.

The exact role of a magnetometry system may vary from experiment to experiment. However, central to its function is an ability to both temporally and spatially resolve the field in or around the Ramsey cell. Moreover, as many of the limiting systematics are related to the interaction of the magnetic and electric fields, the magnetometer should provide an independent measurement of the magnetic field in the Ramsey cell whilst the electric field is operational. In other words, the magnetometer must

1. Be sensitive to the volume-averaged field, $\langle B\rangle$, in the neutron volume;
2. Have a sampling frequency equal to or quicker than the Ramsey measurement frequency;
3. Be accurately calibrated and have an intrinsic uncertainty much less than that due to the determination of $\langle B\rangle$ with the neutrons alone.

$\langle B \rangle$ is in principle all that is needed if the neutrons sample the field in the Ramsey cell without bias. In solely a magnetic field, the neutron counts should be dependent on the phase difference caused by any discrepancy between the AC frequency and the Larmor precession frequency in $\langle B \rangle$. This will of course be broadened due to counting statistics and sampling differences in the neutron ensemble.

The magnetometer should be calibrated so that it is able to reproduce $\langle B \rangle$ in this arrangement. If it is unaffected by the electric field, then it will yield a value for $\langle B \rangle$ true to the magnetic field that can be compared with that from the neutrons to separate out any systematic effects arising from the introduction of E. The signal could also be analysed in real-time to help decide what AC frequency should be used for the next measurement cycle. In reality, there are further complications that do not allow for such a simple interpretation. For instance, if the neutrons do not sample the field randomly, perhaps having a tendency to remain near the bottom edge of the Ramsey cell under gravity, a $\langle B \rangle$ determined by the magnetometer will not be in line with the field experienced by the neutrons. Therefore additionally, the magnetometer might also provide a spatial profile of the field, or simply information about the gradients across the cell, to get a greater handle on the potential systematics. In light of these considerations, it serves to list magnetometers that could be used in an EDM experiment, and their strengths and weaknesses. Those which can provide the required resolution include comagnetometers, atomic magnetometers and SQUIDs.

Comagnetometers are so-called as they consist of a separate particle species, which occupies the same volume as the neutrons, and are thus sensitive to the same, volume-averaged fields. They were first successfully implemented in the limit-holding experiment by the RAL-Sussex-ILL collaboration, which used a ^{199}Hg comagnetometer [91] (see Figure 2.9). To measure the precession frequency of the mercury, polarised light from a ^{204}Hg source was directed across the neutron volume and its intensity measured by a photomultiplier tube. With the ^{199}Hg spins rotated to be perpendicular to the plane of polarisation of the beam, their

absorption cross-section and hence the transmitted light intensity was seen to vary sinusoidally as the ^{199}Hg atoms precessed, modulated by an exponential decay due to spin relaxation. The ^{199}Hg comagnetometer was able to measure volume and time-averaged fields of a few mG with sub-nG precision, reducing the impact of systematic errors by a factor of ∼20 compared with previous experiments.

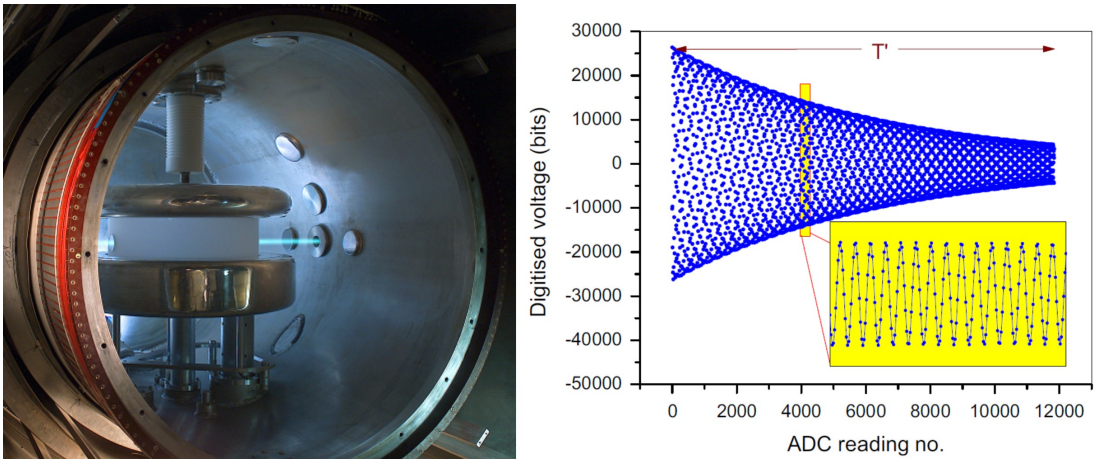

FIGURE 2.9: Photograph showing the ^{204}Hg reading light being sent across the Ramsey cell, along with the typical read out of the transmitted light intensity by a photomultiplier tube [46]. The signal shows modulation consistent with the precession of the mercury atoms, with an envelope caused by exponential relaxation of the spins.

The key features of an ideal comagnetometer are a negligible permanent EDM, such that it itself does not change its precession frequency upon exposure to an electric field; a long spin coherence time, such that no significant dephasing occurs within the particle ensemble in the duration of a measurement cycle; a small neutron absorption cross-section so that it does not deplete neutron numbers. There are only three known viable isotopes for a comagnetometer: ^{199}Hg, ^{129}Xe and ^{3}He.

^{199}Hg is to date the only comagnetometer that has been demonstrated in an nEDM experiment. It has been integrated into several nEDM experiments. However, it cannot be used in cryogenic conditions. Likewise, ^{129}Xe will not be feasible at low temperatures. Despite having vanishing neutron absorption cross-section, the present limit on its EDM is $\sim 10^{-27}\, e\,\mathrm{cm}$ [42], which means it is not useful much beyond the present nEDM limit.

^{3}He, however, is a prospect for a cryogenic experiment. ^{3}He has a large neutron absorption cross-section, and neutron capture will occur when the spins of the two particles are back-to-back. The SNS experiment plans to use ^{3}He as both a comagnetometer and a spin analyser, although interpretation of the signals may be complicated by this behaviour. The precession of the ^{3}He particles will be monitored using a SQUID gradiometer system, whilst the scintillation produced by back-to-back neutron capture is used to detect the spin states of the neutrons [92].

Comagnetometers have the distinct advantage (or disadvantage, depending on how you look at it) of being affected by the same sources of systematic uncertainties as the neutrons, in particular the geometric phase and $\vec{v} \times \vec{E}$ effects. The analysis of the signal must take into account the differing average speed of these particles and their susceptibility to these effects. However, this can offer the possibility of distinguishing between a frequency shift caused by an EDM and that caused by systematic effects under the reversal of the electric field, if the EDM of the comagnetometer species is negligible, which would be impossible to achieve with an alternative magnetometer.

Atomic magnetometers consist of optically pumped alkali vapour, typically rubidium, caesium or potassium, encapsulated in a small glass sphere (see Figure 2.10). These elements have a hydrogen-like electron configuration that implies that both the ground state and the lowest excited state are divided into a hyperfine doublet. The atoms are first polarised by using a suitable optical pumping source exciting them into 'dark' states, which can no longer absorb these photons. The vapour is then transparent to the pumping light, such that it will be maximally transmitted. In the case of caesium, in the presence of a magnetic field, the Zeeman effect will cause further splitting of the hyperfine structure into 16 energy levels. The energy differences within each hyperfine multiplet, in the limit of a small external field, are linearly dependent on the field magnitude. A small RF magnetic field can be applied, which can be used to drive transitions between

neighbouring Zeeman levels. If in a 'bright' state, the unpaired electron can once again absorb a photon, and thus the transmitted light detected will decrease. The Zeeman splitting can be determined by finding the frequency that the AC source must be tuned to to instigate the biggest change in the light observed. This process of optically detected magnetic resonance is integral to the operation of an atomic magnetometer. In the classical picture, the procedure draws parallels with the Ramsey resonance technique. Absorption only takes place when the propagation direction is in the same direction as the polarisation.

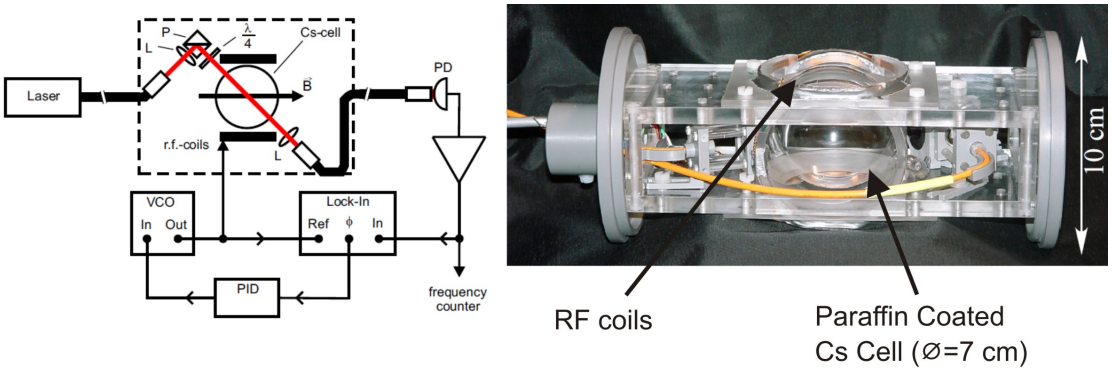

FIGURE 2.10: Schematic and photograph of the setup behind a caesium vapour magnetometer that could be installed at the nEDM experiment at PSI [93].

Conventional atomic magnetometers are scalar since they are sensitive only to the magnitude of the magnetic field, and not its individual components. They can be configured as vector magnetometers by, for instance, applying artificial fields in orthogonal directions, monitoring the magnetometer response and thus deducing the prevailing direction of the original field. Further, if these fields were time-varying with different frequencies, it could be possible to demodulate the resultant signal and establish which components are coherent with the field [94] [95] [96]. This would be difficult to incorporate in a nEDM experiment, where field contamination near the Ramsey cell needs to be minimised. An alternative is simply to apply a second probing light at an angle to the first: any vector information is encoded in the phase difference between the two signals.

A major limitation of the performance of atomic magnetometers is any spin decoherence due to spin-exchange collisions. Spin-exchange relaxation-free (SERF)

magnetometers work in a regime which diminishes such interactions between the particles by having a near zero magnetic field with a high alkali density. A SERF magnetometer may have an accompanying set of Helmholtz coils, designed to cancel some of the ambient field so as to maintain the low field needed for operation. A feedback system can be used to adjust the currents in these coils, which in turn provide a measure of the vector components of the field [97]. SERF magnetometers have shown sensitivities comparable with that of SQUIDs of $1\,\mathrm{pT}/\sqrt{\mathrm{Hz}}$ and beyond [98]. A distinct advantage of atomic magnetometers is that they do not require cryogenic conditions. The question remains as to whether, if the experiment were cryogenic, these magnetometers could be used instead of SQUIDs. They could present the better choice because they require only optical readout that will cause little magnetic disturbance. However, the vapour cells will need to be well-insulated if they are to preserve the phase of the atoms at low temperatures, and the impact of cryogenic conditions on the magnetometer performance is not well known. A further question with regard to a SERF magnetometer is whether it could be used without any feedback coils that could generate a spurious field that the neutrons will be sensitive to, because it should, in principle, be in a well-shielded position and need only be sensitive to small fields.

SQUIDs, or superconducting quantum interference devices, are the main focus of this thesis, since they are an obvious option for any cryogenic experiment. Their operation is described in greater detail in Chapter 3, but in short, relies upon the quantum mechanical nature of superconductivity. At its most basic, a SQUID consists of a superconducting ring intercepted by one or two Josephson junctions. The flux threading the ring must be an integer number of flux quanta, and the voltage which develops across the SQUID in response to a varying magnetic environment can be used to determine the field. SQUIDs are therefore sensitive only to relative changes in field rather than absolute magnitude. Flux quantisation, however, confers the SQUID a significant sensitivity to magnetic fields, with demonstrated noise levels as low as $3\,\mathrm{fT}/\sqrt{\mathrm{Hz}}$ [99].

A SQUID magnetometer is usually configured with a flux transformer, a superconducting input circuit with a much larger sensing loop to measure the field, which is coupled to the SQUID washer. The optimal size of this pickup loop will depend upon the magnetic environment, with a larger loop being preferable in a quieter setting to provide a good signal-to-noise ratio. The immediate consequence of using a flux transformer is that the SQUID signal can no longer be treated as a point-like measurement, since the result is dependent on the integrated flux over the finite area of the pickup loop. One could make the assumption that the field is uniform and thus any variations will be at the central point. However, if a small signal is detected when using a large loop, it gives confidence that the field is not varying significantly for a vast region of the experiment. Likewise, since the current is only generated in the flux transformer due to flux which traverses the loop, it provides only a measurement of the field component normal to the plane of the loop. Three SQUIDs would be needed, with orthogonal pickup loops, to configure a vector magnetometer.

Like the atomic magnetometer, a SQUID would sense the magnetic field at a single position outside the neutron cell. It would therefore not be directly sensitive to the same fields the neutrons experience, as the comagnetometer would be. An array of magnetometers could potentially be used to interpolate the field pattern across the neutron volume. Studies will need to establish how many sensors would be needed to obtain good areal coverage of the field. The magnetometers will also need to be made robust to large electric fields if they are to sit at the high voltage end. For SQUIDs, this is difficult as any sparking could destroy its sensitive electronics, which implies that the pickup loops should sit away from the electrodes, and the SQUID chips themselves even further; for an atomic magnetometer, the priority is the prevention of leakage currents across the vapour cell surface. The main issue, however, is whether this is a worthwhile vocation for these magnetometers to undertake. A complete field map, whilst certainly desirable, is most likely superfluous, particularly if the field variations are very small and if well-defined,

easily characterised gradients are thought to exist, which would be true if the sources of field are far from the cell and/or can be identified. A well-developed simulation would be required to predict the impact of the obtained field pattern on the neutron polarisation. In terms of a volume-averaged field, therefore, it would be difficult to surpass the comagnetometer, as either the field map would have to be averaged, or else perhaps a fitting method could be invoked. This would involve combining the signals of the magnetometers to recreate the volume-averaged magnetic field determined using the neutrons, training the fit on the neutron signal in the absence of an electric field. Many experiments therefore propose using atomic magnetometers in conjunction with a comagnetometer to provide the best of both worlds. For cryogenic experiments, however, where there is no clear tried and tested alternative, the viability of using SQUIDs for these functions must be given careful consideration. The outcomes of studies made for this thesis are included in Chapter 7.

There are of course options other than the three above, for example, **UCNs** and **fluxgates**. The former is used as a magnetometer in its own right during the course of the experiment. A double chamber could be used, one which would be exposed to the electric field, and the other which isn't, remaining a neutral cell. The latter would be used solely as a control, such that any difference between the two cells due to the electric field is manifest. Fluxgates do not show the required sensitivity for an nEDM measurement, but would also form a useful part of a magnetometry system that could be potentially used to monitor external fields, before they are attenuated by the shielding, and/or as part of any active compensation system. However, the comagnetometer, atomic magnetometer and SQUID are unrivalled in their sensitivity and potential integration into a neutron EDM project.

2.8 Summary

This chapter has seen a discussion of the various components needed to build a competitive neutron EDM experiment. Whilst these experiments have at their core a UCN bottle and a measurement centred around the interaction of the neutrons with electromagnetic fields, the exact methods they use to produce polarised neutrons, to generate, monitor and control both electric and magnetic fields differ. Table 2.2 summarises the main features of each contemporary UCN experiment.

Only two approaches have thus far investigated the possibility of performing a neutron EDM experiment in cryogenic conditions. Despite this, because of the advantages of higher neutron numbers and superconductors, cryogenic experiments seem the direction of EDM efforts for the future. CryoEDM has showcased the technology behind a low temperature experiment, including the possibility of using SQUIDs for magnetometry We have seen that SQUIDs are a viable choice for this niche, and their configuration and operation are the subject of the rest of this thesis.

Experiment	Sensitivity Aim (e cm)	Magnetometer	Key Features
PSI	5×10^{-28} [100]	Cs magnetometer array; ^{199}Hg comagnetometer	Uses equipment from the limit-holding experiment
TUM-FRMII	5×10^{-28} [101]	^{199}Hg comagnetometer	Active compensation
PNPI	3×10^{-28} [102]	Cs magnetometer array	Mutli-chambered detectors
TRIUMF-RCNP	1×10^{-28} [103]	^{129}Xe comagnetometer	Planned world's most powerful neutron source; spherical coil
CryoEDM	*$\sim 1 \times 10^{-28}$* [48]	*SQUIDs*	*Cryogenic; superconducting solenoid and shields*
SNS	2×10^{-29} [104] [105] [106]	^{3}He comagnetometer with SQUID readout	Cryogenic; based on Golub and Lamoureaux's ideas [107]

TABLE 2.2: The main parameters and distinguishing features of contemporary bottled UCN experiments.

Chapter 3

Concept and Design of the SQUID Magnetometry System (SMS) for CryoEDM

This chapter introduces the SQUID magnetometry system (SMS) that was constructed for CryoEDM, the topic of this thesis. After a brief introduction to CryoEDM and the expected capability of the SMS, it discusses the operating principles of a direct current (DC) SQUID, and how it could be configured as a magnetometer. The design of the SMS follows, with a focus on parts of the system that were incorporated during the author's time on the experiment.

3.1 A Brief Introduction to CryoEDM

As the inaugural nEDM experiment to operate in cryogenic conditions, a significant R&D effort was behind CryoEDM, some aspects of which have been mentioned in Chapter 2. The main reason to adopt such temperatures is the improved production of UCN via down-scattering in superfluid helium, which can yield UCN

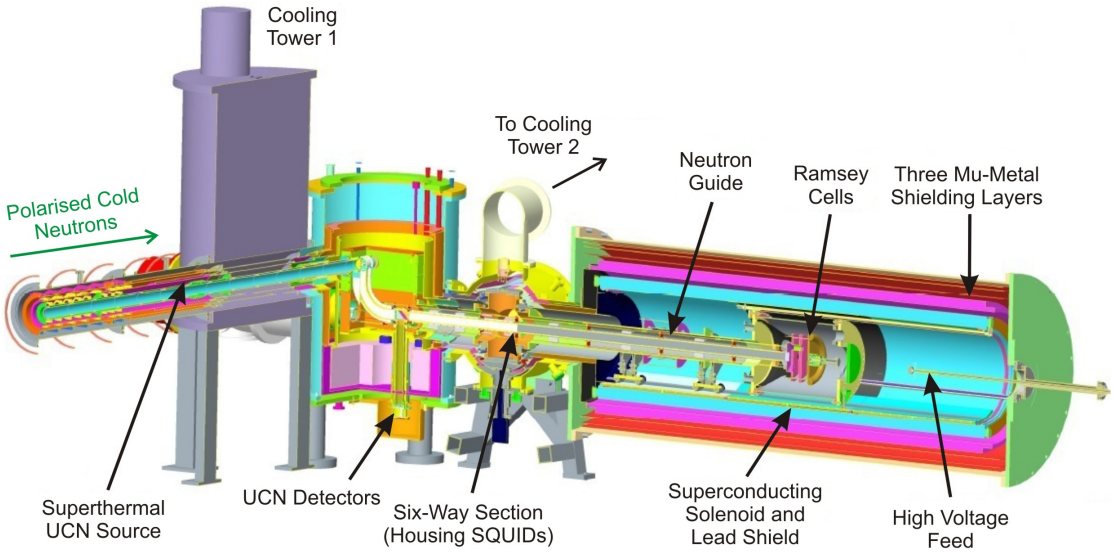

FIGURE 3.1: Diagram of the CryoEDM setup. The arrangement of the SQUID system in relation to the experiment is shown in more detail in Figure 3.6.

densities orders of magnitude greater than conventional methods. As we also have seen in Chapter 2, there are several other advantages, these being for CryoEDM

- A larger, sustainable electric field due to the later onset of electric breakdown in liquid helium compared to vacuum.
- Superconducting magnetic components: a solenoid to generate a stable B_0 field; a lead shield to freeze in the field and screen any spurious sources after cooling; SQUIDs to monitor changes in the field.
- A colder, less energetic environment, making up-scattering of UCN by collisions with each other and with impurities unlikely, leading to a longer storage time.

CryoEDM is unlikely to be the last cryogenic nEDM project, and the technological developments behind it, not least the SMS, will enter into scientific lore and could be resurrected for a future experiment.

The apparatus for CryoEDM is illustrated in Figure 3.1. It was situated on what was then the most intense neutron beam in the world, with a neutron flux of

$1.5 \times 10^{15}\,\mathrm{n\,cm^{-2}s^{-1}}$, of which $\sim 10^{7}\,\mathrm{n\,cm^{-2}s^{-1}}$ are 9 Å cold neutrons. These neutrons would pass through the polariser, detailed in Section 2.4, and on towards the superthermal source. The source, as do the other volumes the neutrons will traverse, contains isotopically-pure, superfluid ^{4}He at a temperature of $< 0.7\,\mathrm{K}$, which is required for the down-scattering process. It is in fact held at a temperature closer to 0.5 K to maximise the UCN storage time once they are produced. The 'fill' time of the source is about 160 s, which accounts for both the accumulation of neutrons in the volume, and the time for effective conversion to UCN.

After the source, they enter the transfer section that consists of two right-angled bends. Whilst these were a necessity given the spatial constraints of the experimental hall, they also help to reduce the proportion of cold neutrons that remain in the beam. With the source tube being 3 m long, and the mean free path of the cold neutrons in He-II around 10 m, some 70% of these residual particles are lost at the first bend, continuing straight on into a lead beam stop. The rest of the neutrons flow under gravity, through to the neutron guide, which is made of copper and coated in beryllium to prevent any wall interactions that can cause UCN or polarisation losses. The neutrons diffuse past a flap valve that would later deflect them down towards the spin analysers, and on to the two Ramsey cells.

The entire neutron guide system, from source to cell, is surrounded by holding field coils (not illustrated) which manage the field transition from 30 mT at the polariser to the $5\,\mu$T at the neutron cells, maintaining a gradual gradient to preserve the neutron polarisation (see Chapter 5). This latter field is generated by a superconducting solenoid, formed of 10860 turns of CuNbTi wire, that is run in persistent mode once its critical temperature is surpassed. This, and the superconducting lead shield which surrounds it, are placed in the 4.2 K helium tank, itself in a 77 K nitrogen jacket. Further magnetic shielding is provided by the three mu metal shields that lie outside of the nitrogen layer.

Parameter	Predicted Value	Gain in Sensitivity
Electric Field, E	$40\,\mathrm{kV\,cm^{-1}}$	4.0
Visibility, α	0.9	1.4
Storage Time, T	300 s	2.3
Neutron Numbers, N	$7{\times}10^5$ cycle^{-1}	7.1

TABLE 3.1: The key parameters involved in the statistical sensitivity of a nEDM experiment (see Equation 2.15), their aimed for values in CryoEDM [108] and the improvement they will provide compared to the limit-holding, room-temperature experiment.

The helium tank encloses the Superfluid Containment Vessel (SCV) in which the cells are located. The cells are formed of two beryllium oxide cylinders of height 45 mm and diameter 250 mm, separated by three titanium electrodes coated on the inside with diamond-like-carbon. The first two electrodes are grounded, whereas the one furthest downstream is charged via the high voltage (HV) feed discussed in Section 2.5. The first compartment remains a neutral cell, whereas the second has a large electric field across it and is termed the 'HV' cell. Access between the two is via a valved aperture, which lets the cells be filled/emptied separately.

Slightly further upstream of the cells lie the pickup loops that are the sensing areas of the SQUIDs. They are there to monitor the field in the vicinity of the cells as the neutrons undergo the Ramsey resonance method (Section 2.1). The neutrons precess about the net magnetic field, with their spins rotated under the influence of an AC field. This oscillating field is generated by a pair of coils, tuned to a frequency in the vicinity of 150 Hz, the resonance frequency of neutrons in a 5 μT field. The neutrons are stored for 300 s, after which the first cell is emptied and the neutrons move back upstream, and are directed toward the neutron detectors. As mentioned in Section 2.4, the detectors planned to be used for CryoEDM are solid-state silicon devices, each coated with a thin film of lithium fluoride. Neutron capture by ^{6}Li results in alpha and triton particles, which deposit their energy in the silicon layer. The detectors act as spin analysers, allowing distinction between

the two spin states. The second cell is emptied after the contents of the first cell has been analysed. The data from the two cells can be compared, assuming that the magnetic environments are similar, to see if there is any signature of an electric dipole moment. Whilst in this final phase of the measurement procedure, the superthermal source is refilled to begin the next cycle and minimise the turnaround time. In all, a single measurement takes on the order of 500 s from start to finish, with an expected 7×10^5 neutrons undergoing the process simultaneously.

The projected improvement in statistical sensitivity is then about a factor of 90, as broken down in Table 3.1. Even if some of these targets are unmet, the experiment was thought to reach its goal sensitivity of $10^{-28}\, e\,\mathrm{cm}$ after a few years of running. The overall systematic uncertainty was estimated to be $1.7 \times 10^{-28}\, e\,\mathrm{cm}$, composed of several uncorrelated sources, some of which have been mentioned in Chapter 2.

A great many of these effects impact on the magnetic field seen by the neutrons, which makes monitoring the field changes during measurement crucial. The SMS is designed to track the field not only whilst the neutrons occupy the Ramsey cells, but before and after, to advise on the frequency at which to tune the AC coils, and to characterise the field variations in general. However, with regard to the primary objective of determining the volume and time-averaged field in the neutron cells during the Ramsey cycle, with what precision should the SQUIDs be capable of measuring it?

Returning to Equation 2.9, we saw that, in the absence of an electric field, the neutron counts give a measure of ω_0 and hence the magnetic field. If there is a small change in the magnetic environment over the course of a measurement cycle, and therefore the volume-averaged field sampled by the neutrons, then the resonant spin frequency will be perturbed by a small amount, which we will call $\delta\omega_0$. Introducing the electric field, the difference between the up and down counts,

$\Delta N = N_u - N_d$, for a given measurement becomes, using Equation 2.9,

$$\begin{aligned}\Delta N &= 2\bar{N}\alpha[\sin(\lambda T)\sin(\delta\omega_0 T + \omega_E T) - \cos(\lambda T)\cos(\delta\omega_0 T + \omega_E T)] \\ &\approx 2\bar{N}\alpha(\delta\omega_0 + \omega_E)T,\end{aligned}$$

with $\lambda T \approx \frac{\pi}{2}$. At the working point, this yields, after some rearranging,

$$d_n = \frac{\hbar\Delta N}{2N\alpha TE} - \frac{\mu\delta B}{E},$$

where we have substituted in explicit forms for $\delta\omega_0$ and ω_E, and used $\bar{N} = \frac{N}{2}$. Thus the uncertainty in the EDM becomes

$$\sigma(d_n) = \sqrt{\left(\frac{\hbar}{2\alpha TE}\right)^2 \sigma\left(\frac{\Delta N}{N}\right)^2 + \left(\frac{\mu}{E}\right)^2 \sigma(\delta B)^2}.$$

As this is the result of a single measurement, when considering a number of runs, n, the uncertainty will be further suppressed by a factor of $\sqrt{n}$. Taking the duration of a cycle as 500 s, an estimated 22500 runs should be possible if there are 150 reactor days scheduled on which the experiment can operate in a calendar year. Therefore to reach a measurement sensitivity in d_n of $10^{-27}\,e\,\mathrm{cm}$ in a year, the magnetic field must be known to a precision of

$$\sigma(B) = \sqrt{n\sigma(d_n)^2 - \left(\frac{\hbar}{2\alpha TE\sqrt{N}}\right)^2} \approx 0.1\,\mathrm{pT}, \tag{3.1}$$

using values from Table 3.1. This is the ultimate requirement of the SMS, which will achieve this through a combination of its 12 SQUIDs (Chapter 7). This result also serves as a conservative estimate of the resolution required of each SQUID magnetometer. The rest of this chapter will be devoted to explaining how a SQUID can be capable of making measurements at such high precision, and how the infrastructure has been designed to enable them to be successfully operated in such a regime.

3.2 The DC SQUID as a Magnetometer

A Superconducting Quantum Interference Device, or SQUID, is one of the most sensitive devices in the world, capable of measuring, amongst other variables, current, temperature and magnetic field. It has found application in fields as diverse as biomedical imaging, e.g. magnetoencephalography; geological surveyance; noise thermometry and precision instrumentation in low temperature experiments, from astroparticle detection to quantum computing. As the name betrays, the operation of a SQUID is founded on the phenomenon of superconductivity and its quantum nature.

Superconductivity is characterised by perfect diamagnetic behaviour, a property discussed in greater detail in Chapter 4, and zero electrical resistance. The phase manifests only in certain materials, and then only when they are cooled below their critical temperature. A successful formulation to describe the underlying mechanics was devised by Bardeen, Cooper and Schrieffer (the BCS theory [109]), who postulated that the free-flowing current was not due to movement of normal electrons but that of 'superelectrons' that can occur at these low temperatures. These are effective new particles formed by the coupling of sets of two electrons in what are known as Cooper pairs. Cooper pairs can be treated as bosonic particles that therefore do not need to obey the Pauli principle: there is no reason that multiple pairs cannot occupy the same ground state. Moreover, they should not be treated independently because the high density of charge carriers and the long range of a Cooper pair interaction suggest that they have overlapping centres of gravity, and are strongly correlated spatially. The situation can be visualised as a superconducting 'condensate,' a highly coherent soup of Cooper pairs which continually break and reform. The many-body state can be described with a common centre of mass, and can be represented by a single macroscopic wave function:

$$\psi(\vec{r},t) = |\psi(\vec{r},t)|e^{i\phi(\vec{r},t)} = n^{1/2}e^{i\phi(\vec{r},t)}, \tag{3.2}$$

where n is the Cooper pair density, and ϕ the phase which underpins the centre of mass motion of the condensate.

This is of importance when considering the superconducting ring that is the backbone of the SQUID. The wave function must be single-valued across the closed loop, and therefore the phase can only change by multiples of 2π in traversing the ring. In the case of an external field, we can use this to show that this amounts to a quantised flux threading the loop. Applying the momentum operator on the wave function, we find that

$$\hat{p}\psi = (-i\hbar\nabla - q\vec{A})\psi = (\hbar\nabla\phi - q\vec{A})\psi = m\vec{v}\,\psi, \tag{3.3}$$

where q is the charge of one carrier, that is $2e$, and m and v are the mass and velocity of one Cooper pair respectively. We can take a contour integral around the loop and thus express the above as

$$\oint m\vec{v}\cdot \mathrm{d}\vec{l} = \oint (\hbar\nabla\phi - 2e\vec{A})\cdot \mathrm{d}\vec{l}. \tag{3.4}$$

The velocity can be written in terms of the supercurrent, and since this is confined to the outer edge of the loop, it can be taken to be negligible along the chosen contour and thus the LHS becomes zero. The magnetic vector field contribution can be converted via Stokes' theorem to finally give

$$\oint \nabla\phi\cdot \mathrm{d}\vec{l} = \frac{2e}{\hbar}\oint \vec{A}\cdot \mathrm{d}\vec{l} = \frac{2e}{\hbar}\int \vec{B}\cdot \mathrm{d}\vec{s}$$

$$\rightarrow 2n\pi = \frac{2e}{\hbar}\Phi, \tag{3.5}$$

and therefore the flux, Φ is quantised in units of $\Phi_0 = {}^{h}/_{2e}$, as was experimentally verified by [110] and [111].

Not only is the flux that threads the ring quantised, it will be trapped so long as the loop remains in the superconducting state. After it is cooled below its critical

temperature, any subsequent change of external field will cause a current to be induced in the loop to maintain the flux through it, in accordance with Lenz's law. The difference compared with the behaviour of a normal conductor is that this current will theoretically persist forever. The coherence of the superconducting phase implies that a decay of the current will require the entire ensemble of charge carriers to react simultaneously, which is statistically very unlikely.

As it stands, whilst we may get an idea of measuring the field from the induced supercurrent, a superconducting loop on its own is not a very useful construct since, for one, unless the changes of field are small, the induced current will exceed the supercurrent and the loop will become normal. This is where the concept of Josephson junctions comes in.

The junction represents a coupling between two regions of superconductor that are spatially close, which could result from a constriction of the material, a point contact, or the introduction of a thin layer of insulator or normal conductor between the two. The surprising observation was that the phases of two superconducting areas thus positioned are related, and that Cooper pairs can tunnel through the barrier such that a resistanceless current can flow. One can solve the Schrödinger equation for the system to find an expression for this supercurrent in terms of the phases of the superconductors,

$$I(t) = I_c \sin\delta = I_c \sin\left(\frac{2e}{\hbar}\int V(t)\mathrm{d}t + \delta_0\right), \tag{3.6}$$

where I_c is the critical current of the junction, $\delta_0 = \phi_2(0) - \phi_1(0)$ is the original phase difference between the two superconducting regions, and $V(t)$ is the voltage applied to the junction. A current will therefore flow even in the absence of an external voltage, giving rise to the DC Josephson effect; when there is a DC voltage drop, the supercurrent will oscillate with frequency $f = {}^{2eV}/_{h} = {}^{V}/_{\Phi_0}$, an observation known as the AC Josephson effect.

A DC SQUID has two Josephson junctions interrupting a superconducting loop, and it is the interference of the currents which flow through them that provides a measurement of the magnetic flux. With no external voltage, we can use the same prescription as adopted in Equation 3.5 to show that the phase difference between the two junctions must be $\phi_2 - \phi_1 = {}^{2e}/_{\hbar}\,\Phi$. Assuming the two junctions to be identical, we might then write their currents as $I_1 = I_c \sin\left(\delta_0 + \frac{e}{\hbar}\Phi\right)$ and $I_2 = I_c \sin\left(\delta_0 - \frac{e}{\hbar}\Phi\right)$, giving a total current through the SQUID of

$$I = I_c \left[\sin\left(\delta_0 + \frac{e}{\hbar}\Phi\right) + \sin\left(\delta_0 - \frac{e}{\hbar}\Phi\right)\right] = 2I_0 \cos\left(\frac{e}{\hbar}\Phi\right), \tag{3.7}$$

where $I_0 = I_c \sin\delta_0$ is the current given by the DC Josephson effect without a magnetic field. The maximum current occurs when $I_0 = I_c$, and thus the critical current of the SQUID, $I_{\mathrm{SQ}_c} = 2I_c \left|\cos\left({}^{e}/_{\hbar}\Phi\right)\right| = 2I_c \left|\cos\left({}^{\pi\Phi}/_{\Phi_0}\right)\right|$, varies periodically with the flux and therefore the applied field, with a period of one flux quantum. Further, this critical current is twice that of either of its junctions in the case of no ambient field.

The question of how we can determine the magnetic field using such a device still remains. The bias current could be adjusted until it hits the critical current value, at which a non-zero voltage appears and from which the flux can be determined by the modulation pattern. A more effective method to convert the SQUD into a magnetometer is to bias it with a constant current slightly greater than its critical current. In this case, the junctions are resistive, while the bulk material either side of them remains in the superconducting phase. A voltage develops across the SQUID which, as it happens, will also depend periodically on the flux threading the loop, in antiphase to the relationship that was seen with the critical current: it is this voltage that is read out, and therefore a SQUID serves as a flux-to-voltage transducer. This behaviour can be naïvely inferred by examining the V-I characteristic of a SQUID. Just as with a constant, or zero voltage, the current will fluctuate between the two extreme curves, the voltage will oscillate between the

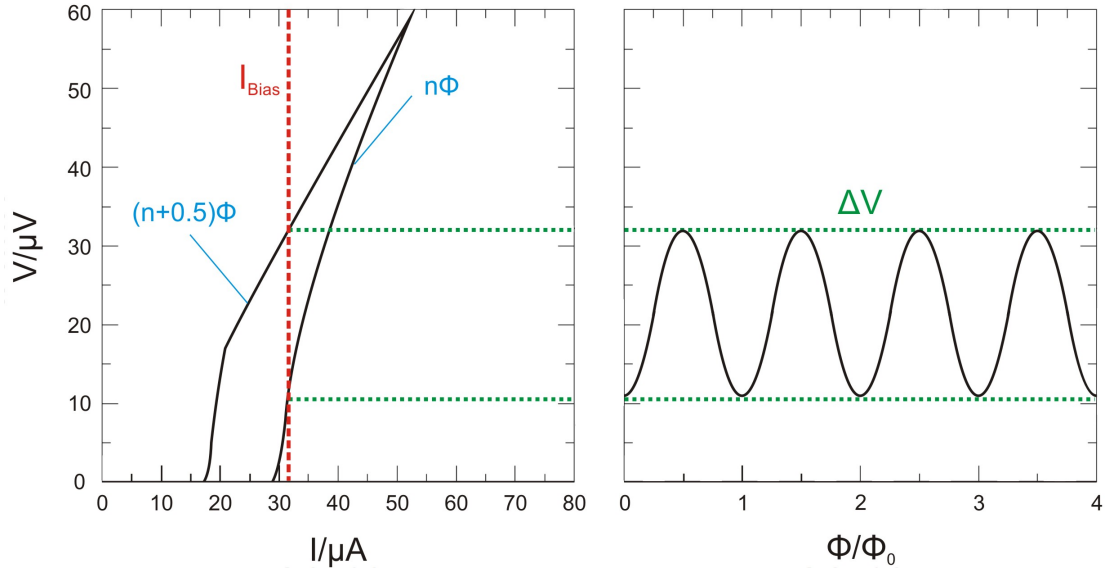

FIGURE 3.2: Typical V-I and V-Φ characteristics of the DC SQUID. Biasing the SQUID at a current just above the critical current.

same curves with a suitable static current, this time being a minimum when the threaded flux is an integer number of flux quanta, and a maximum when a half-integer. We can also see that the discrimination between the different fractions of flux quanta is greatest at a bias current just over the critical current of the integer quanta curve, as this is when the distance between the two curves is greatest (Figure 3.2).

At first glance, this seems in contradiction to the equations derived from Equation 3.6 onwards, but of course, these derivations have focussed on the supercurrent and have neglected the normal, quasiparticle current that would arise when the junction becomes resistive. We can make allowances for this by considering the effective circuit to be the junction in parallel with a 'normal' resistance, R. In fact, a junction will often be deliberately shunted with a resistor to suppress hysteresis in its V-I characteristic and enforce an Ohmic, linear relationship. The current across the junction then becomes

$$I(t) = I_j(t) + I_n(t) = I_c \sin\delta(t) + \frac{V(t)}{R} = I_c \sin\delta(t) + \frac{\Phi_0}{2\pi R}\frac{d\delta(t)}{dt}, \tag{3.8}$$

showing the explicit time dependences, with $I_j(t)$ the supercurrent and $I_n(t)$ the

normal current, and where $V(t) = \frac{\Phi_0}{2\pi}\frac{d\delta(t)}{dt}$ is the voltage across it, as derivable from Equation 3.6. The equation can be solved to yield $\delta(t)$ and therefore $V(t)$ and its time-averaged value, $V = \langle V(t) \rangle$. An ideal SQUID can be treated in a similar way, with $R \to {}^{R}/_{2}$ [112]. The expression for the time-averaged voltage across a SQUID is then

$$V = \frac{R}{2}\sqrt{I_{\mathrm{SQ}}^2 - I_{\mathrm{SQ}_c}^2} = \frac{R}{2}\sqrt{I_{\mathrm{SQ}}^2 - 4I_c^2\cos^2\left(\pi\frac{\Phi}{\Phi_0}\right)}, \tag{3.9}$$

and thus if biased with a constant current, I_{SQ}, the voltage-flux relationship will oscillate with a period of a flux quantum, being a maximum when the critical current is a minimum and vice versa.

In reality, the junctions will show some capacitive behaviour more fully described by the Resistively and Capacitively Shunted Junction (RCSJ) model [113] [114]. A more thorough treatment will include both this contribution and that of the small flux from the self-field generated by the shielding current in the SQUID loop. With a number of parameters, optimisation of the SQUID response, balancing any degradation of the critical current, modulation depth, any distortion of the V-I relation and any resultant hysteresis, usually requires numerical resolution of the associated equations. Despite this, the periodicity should remain unaffected, and thus our basic model of the SQUID is not lacking in this respect.

However, we note that, because of the sinusoidal variation, the V-Φ function is non-linear and multi-valued so that it is impossible to distinguish between flux changes that are separated by flux quanta. An external circuit may be used to couple flux into the SQUID so that it operates at the steepest part of the V$-\Phi$ curve, where the flux-transfer coefficient, V_Φ, is highest. Given $V_\Phi = |(\frac{\partial V}{\partial \Phi})_{I_{SQ}}|$ and Equation 3.9, we see this occurs when $\Phi = \frac{2n+1}{4}\Phi_0$. In this regime, small deviations in the flux due to ambient sources, both incremental and decremental, produce proportional changes in the voltage $\delta V = V_\Phi \delta\Phi$. The relationship only remains linear, and thus easy to interpret, for small flux changes. To improve

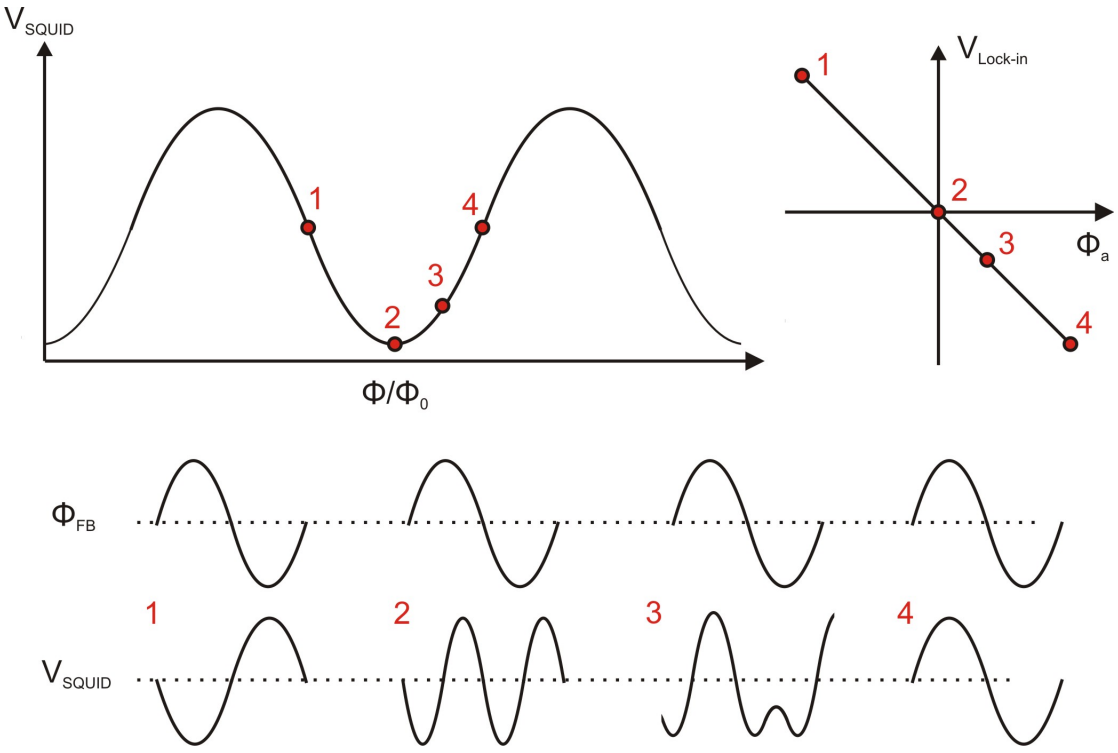

FIGURE 3.3: AC feedback modulation scheme with working point at the minimum of the V-Φ curve of the SQUID, showing the flux applied by the feedback to the SQUID (Φ_{FB}), the voltage response of the SQUID (V_{SQUID}), and the voltage output of the lock-in amplifier ($V_{\text{Lock-in}}$) when ambient flux, Φ_a, threads the SQUID to shift the curve to each of the four points shown. At 2, there is no Φ_a, so the lock-in does not need to set a compensatory voltage, and the modulation of V_{SQUID} is double the frequency of Φ_{FB}. At 1 and 4, V_{SQUID} is single-valued and oscillates with the same frequency as Φ_{FB}. At 3, there is only a small component at this frequency.

the dynamic range and conserve the linearity, the flux through the loop must be maintained at this working point. This requires a feedback system, and the SQUID must be run in so-called flux-locked loop (FLL) mode. This linearises the SQUID response over a much wider range, allowing it to track changes in flux many times a flux quantum at a sensitivity mere fractions of one. The voltage response of the SQUID is amplified and used to drive the feedback coil magnetically coupled to the SQUID. The coil produces a field to oppose a change in flux to preserve the chosen working point, and it is the voltage that is applied to this feedback system that is used as a measure of the external field. Most feedback units use AC modulation, applying a signal of typical amplitude $\leq \Phi_0/4$ at a high frequency, usually $f > 100\,\mathrm{kHz}$: a typical setup is shown in Figure 3.3. In this scenario, the

FLL keeps the SQUID at an extremum of the $V - \Phi$ curve and not at the point where $V_\Phi = |(\frac{\partial V}{\partial \Phi})_{I_{SQ}}|$ is greatest. To see why this is beneficial, we can examine the output of the SQUID when operating at these different working points. We note the $V - \Phi$ curve in the vicinity of an extremum is symmetric in the $\pm\Phi_0/2$ range. When the only field applied to the SQUID is that from the feedback system, the symmetry implies that, at a maximum or minimum, the voltage read out from the SQUID will oscillate with twice the frequency of the feedback modulation. At points of maximum V_Φ, the voltage response will have the same frequency as the feedback signal. In between these two scenarios, with a small displacement from the extremum, the output signal will contain a component at the original frequency of the feedback, whose magnitude will be proportional to the size of the shift. We can show this by returning to Equation 3.9: the total flux is a combination of the flux applied by the feedback and that from any ambient field perturbation, Φ_a

$$\Phi = \Phi_{\text{FLL}} + \Phi_a = \frac{n\Phi_0}{2} + \frac{\Phi_0}{4}\cos(\omega t + \theta) + \Phi_a, \tag{3.10}$$

where θ is an arbitrary phase in the modulation, and where we have assumed the feedback signal to have amplitude $\frac{\Phi_0}{4}$. With the SQUID bias current set higher than its critical current (i.e. $I_{\text{SQ}} > 2I_c$), Equation 3.9 can be expanded to leading order as

$$V \approx \frac{RI_{\text{SQ}}}{2}\left[1 - \frac{2I_c^2}{I_{\text{SQ}}^2}\cos^2\left(\pi\frac{\Phi}{\Phi_0}\right)\right]. \tag{3.11}$$

Substituting in Equation 3.10 gives

$$V \approx \frac{RI_{\text{SQ}}}{2}\left[1 - \frac{I_c^2}{I_{\text{SQ}}^2}\left[1 \pm \cos\left(\frac{\pi}{2}\cos(\omega t + \theta) + 2\pi\frac{\Phi_a}{\Phi_0}\right)\right]\right], \tag{3.12}$$

depending on the chosen extremum and thus whether n is even or odd. Assuming $\Phi_a \ll \Phi_0$, we can use the Taylor's expansion for a cosine function to express the voltage as

$$V \approx \frac{RI_c^2}{2I_{\text{SQ}}}[A \pm B\cos(\omega t + \theta) \pm C\cos 2(\omega t + \theta)], \tag{3.13}$$

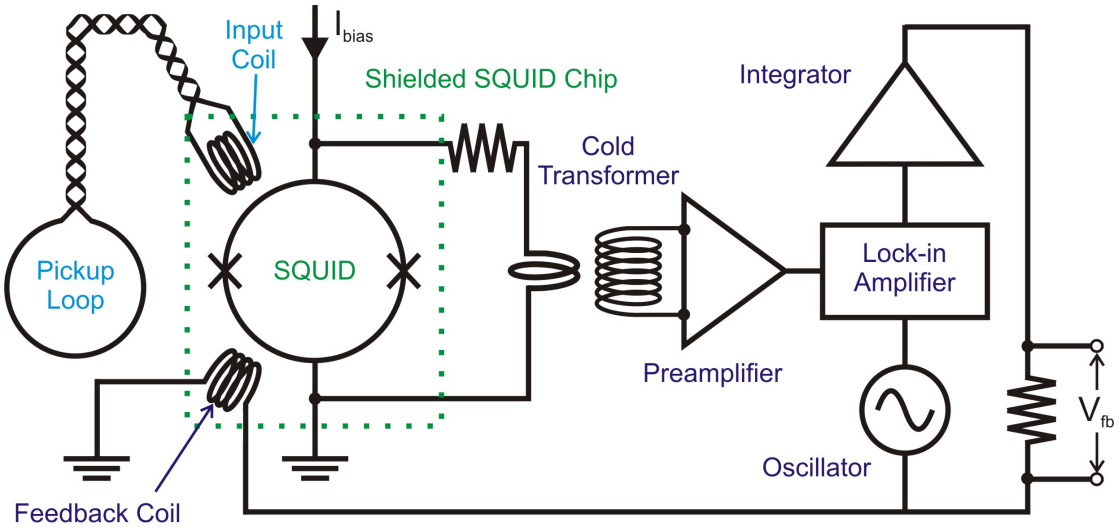

FIGURE 3.4: SQUID circuitry for a DC SQUID run in flux-locked loop mode (feedback and readout in purple) with a flux transformer (in blue).

$$\text{where} \quad A = \frac{I_{\mathrm{SQ}}^2}{I_c^2} - 1 \pm \left(\frac{\pi^2}{16} - 1\right), \quad B = \frac{\pi^2 \Phi_a}{\Phi_0}, \quad C = \frac{\pi^2}{16}.$$

The functional form of V can be shown to be the same if more terms are used in the expansion of Equation 3.9. From this, we can see that, in the case of no external field, the SQUID voltage response is periodic with angular frequency 2ω; in the case of a small perturbation, there is a component at ω that is proportional to Φ_a. It is easy to show with the same prescription and the Taylor's expansion for a sine function that the voltage when the SQUID is biased at $\frac{2n+1}{4}\Phi_0$ will not have any dependence on Φ_a and will in fact oscillate with ω whatever the applied flux. Likewise, any working point between the two will produce an output that contains a contribution with angular frequency ω, but with a smaller constant of proportionality. The sensitivity to an applied flux is greatest at an extremum, and thus the SQUID is optimally biased at this working point.

The feedback element works as follows (see Figure 3.4): the voltage signal from the SQUID is fed into a preamplifier and then into a lock-in amplifier, which works as a phase-sensitive detector, both extracting and amplifying a signal at the chosen carrier frequency. The detector is set to pick out the component at ω, and its output depends on the amplitude of this signal. If the flux threading the SQUID is unperturbed, remaining at the extremum of the $V - \Phi$ curve, there

is no component at ω and the lock-in will put out 0 V; if there is a difference, the voltage set by the lock-in amplifier, V_L is proportional to the component at ω and therefore to Φ_a. This voltage controls the current sent to the feedback coil, which will generate a field in the opposite sense to compensate the change in flux, resulting in $-\Phi_a$ to maintain a constant flux in the SQUID loop. As this current is therefore also proportional to Φ_a, we can measure a quantity related to it by putting a resistor in series with the feedback coil and measuring the voltage drop across it, V_{fb}. V_{fb} is typically sent to an analogue-to-digital converter (ADC) for further analysis. The important thing to note is that, however it is operated, the SQUID is **not sensitive to the absolute value of the magnetic field, but only changes in it**.

With careful design, the incorporation of these electronics does not introduce much noise to the already low-noise, low-T_c SQUID (typically $10^{-6}\,\Phi_0\,\mathrm{Hz}^{-1/2}$). However, using this setup as a magnetometer we note that, since the magnetic flux and field are related by an area factor, so the magnetic field noise is enhanced by this number compared to the flux noise: $B_n(f) = {}^{\Phi_n(f)}/_{A}$. The typical SQUID washer has side length of hundreds of micrometres, minimising the inductance and thus the thermal noise, but leading to a magnetic field noise too large for most magnetometer applications. To improve the effective area, a closed superconducting circuit with a sensing loop much bigger than that of the SQUID can be coupled to it as the primary sensor of the field. This allows the user to control the field catchment by varying the size, shape and position of the loop and also to configure it to be sensitive at higher order, to field gradients or otherwise. The flux transformer at its most basic is simply a pickup loop attached by leads to an input coil, which generates a coupled flux in the SQUID. A flux which threads the pickup loop, Φ_p, will generate a current in the flux transformer according to $\Phi_p = (L_p + L_i)I$, ignoring the inductance of the leads, and labelling the inductance of the loop as L_p and that of the input coil as L_i. Energised by this current, the input coil produces a flux in the SQUID loop of $\Phi_{SQ} = M_i I$, where $M_i = k\sqrt{L_i L_{SQ}}$ is the mutual

inductance between the input coil and the SQUID. The two fluxes are related by

$$\Phi_p = (L_p + L_i)\frac{\Phi_{SQ}}{M_i}. \tag{3.14}$$

This enables us to define the magnetic sensitivity, B_Φ, which is the smallest external field that will produce a response in the bare SQUID when coupled to the flux transformer:

$$B_\Phi = \Phi_0 \frac{L_p + L_i}{M_i A_p}, \tag{3.15}$$

where A_p is the relevant sensing area, which we assume here to be the area of the pickup loop. Of course, with a feedback system, we are not limited to quantised flux and can sense field changes corresponding to fractions of a flux quantum, but B_Φ remains a useful figure of merit for quantifying the sensitivity. With typical units of nT/Φ_0 (dividing both sides of Equation 3.15 by Φ_0), a lower value for B_Φ therefore implies a greater sensitivity, allowing greater resolution of the ambient magnetic field. Assuming that both the coupling constant, k, and the pick-up loop area, A_p, are independent of the inductance of the input coil, the minimum of B_Φ and thus the best sensitivity occurs when $L_p = L_i$, giving

$$B_\Phi = 2\Phi_0 \frac{L_i}{M_i A_p}. \tag{3.16}$$

The inductance of the flux transformer is therefore often designed to match that of the SQUID input coil to take advantage of the better sensitivity. This is easier said than done, and B_Φ would need to be more accurately estimated for this to be effective. Regardless, Equation 3.15 enables us to estimate B_Φ from intrinsic parameters of the SQUID system, which are commonly provided by the manufacturer. Further, with the transfer coefficient of the FLL circuit, which states the voltage response for each flux quantum, we have a straightforward connection between the voltage read out and the magnetic field. This is perhaps best estimated by calibrating the system, applying a known field and measuring the voltage response (Section 3.6.3)).

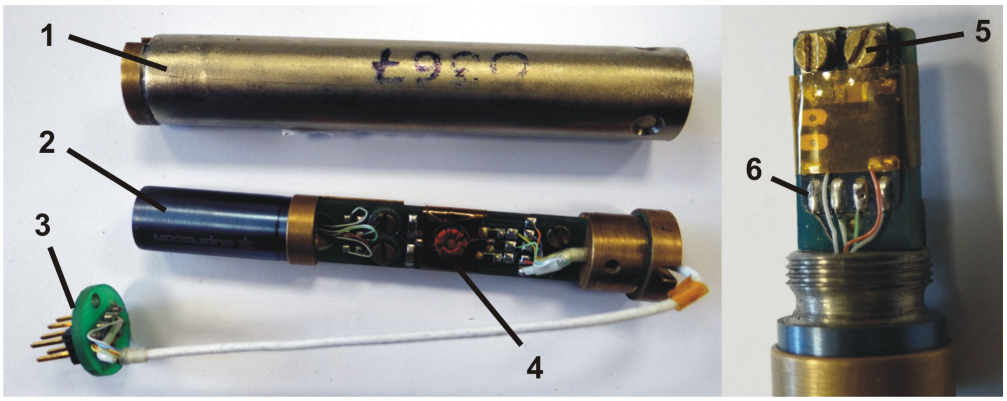

FIGURE 3.5: The CSblue SQUID with 1) Cryoperm shield; 2) niobium cap; 3) 2.54 mm pitch Samtec for connections to electronics; 4) cold transformer; 5) niobium washer, screw terminals for attachment of input; 6) solder pads for feedback, readout, heater and bias current connections to the chip encased in the glass fibre.

3.3 CryoEDM SQUIDs

The twelve SQUIDs used are of model CSblue developed by Supracon, Jena, that is designed to be suitable for magnetometer applications at temperatures of 4.2 K and lower (Figure 3.5). The SQUID loop with Nb/Al-AlO_x/Nb Josephson junctions, input coil, feedback coil and a heater are all amalgamated onto a single, integrated $4 \times 4\,\mathrm{mm}^2$ chip, which is encased in a glass fibre to protect the delicate circuitry. Screw terminals permit the attachment of a separate pickup loop, and the sensor is ready supplied with a superconducting niobium cap that surrounds the chip and excludes any field variations that are not introduced by the flux transformer. Since the SQUIDs themselves were not to be situated in a low-noise environment, they were additionally each given a Cryoperm (a material of high magnetic permeability) shield to screen the external field whilst cooling the apparatus. This is useful in minimising the noise introduced on the SQUID channels but, because of the proximity of the SQUIDs to the neutron guide, could potentially interfere with the polarisation of the neutrons. The trade-off between these two aspects is explored in Chapter 5.

The heater on the chip acts as another precaution, being simply a resistor that,

Parameter		Value
Inductance	SQUID, L_{SQ}	≈400 pH
	Input Coil, L_i	≈350 nH
	Feedback Coil, L_f	≈320 nH
Mutual Inductance to SQUID	Input Coil, M_i	10 nH
	Feedback Coil, M_f	230 pH
Coupling to SQUID	Input Coil, $^1/_{M_i}$	0.21 μA Φ_0^{-1}
	Feedback Coil, $^1/_{M_f}$	9 μA Φ_0^{-1}
SQUID Critical Current, I_{SQ}		23-26 μA
Equivalent Flux Noise		4-6 $\mu\Phi_0$ Hz$^{-1/2}$

TABLE 3.2: Typical values for SQUIDs used in CryoEDM, with parameters measured at 4.2 K with no input attached (open input), as documented in [115].

when energised, heats the SQUID and the input coil above their critical temperatures to expel any trapped flux that may have been introduced during cooling. Thermally-activated movement of vortices can cause extra flux noise, as is described in Chapter 5.

The chip is self-contained and requires only connections to a compatible set of readout and feedback electronics to complete installation. Contact pads on the chip allow for such connections, with there being four associated signals carried on these lines: the SQUID bias current, the feedback modulation, a signal to control the heater and the voltage across the SQUID. A second board contains a cold transformer to match the inductance of the SQUID circuit to that of the feedback system, improving the noise performance. Ignoring the electronics, the CSblue sensor is capable of a current noise of 1.5 pA Hz$^{-1/2}$ or better [116], although this will vary from SQUID to SQUID. The noise is expressed in terms of a more intuitive flux noise in Table 3.2, which gives a summary of the typical values of parameters associated with the SQUID. With the behaviour of these SQUIDs well documented by the manufacturer, the challenge for CryoEDM was to construct the infrastructure to let them operate optimally, given the constraints of

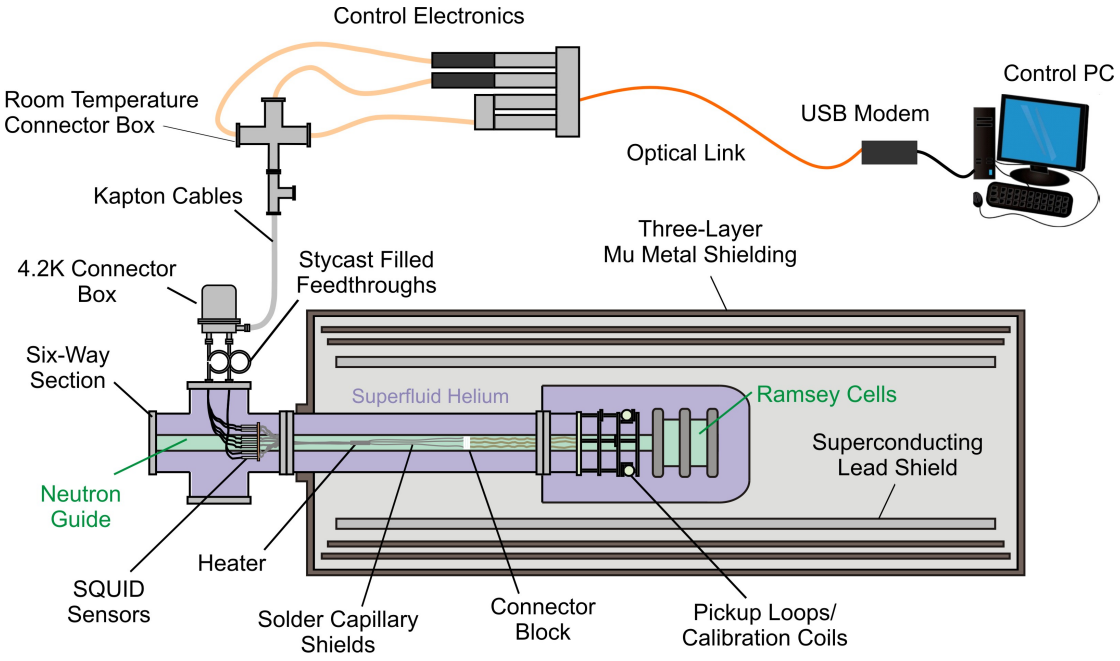

FIGURE 3.6: Schematic of the SMS for CryoEDM.

the experiment.

3.4 Overview of the SMS for CryoEDM

Figure 3.6 shows the layout of the SMS in the context of the CryoEDM components. One of the major problems to overcome was the distances involved in the circuitry. The SQUIDs were situated in the six-way section for convenience, as placing them further downstream was thought to cause potential issues in maintaining neutron polarisation (see Chapter 5), and since the readout cables could be more easily routed via one of the mouths of the six-way piece. The large separations between pickup loop and SQUID imply that shielding to protect against extraneous pickup is important. Heat sinking and mechanical stability are also requirements. The construction is, from input to output (Figure 3.6, Figure 3.7):

- **Input Circuit**
  - **Pickup loops and calibration coils** are wound and affixed respectively on a configuration of G10 formers. These are mounted to one side of the

Ramsey cells because of the danger of sparking if they are placed nearer the high voltage end. The pickup loops are of varying dimension and orientation, with the aim being to use a combination of signals to reconstruct the volume-averaged magnetic field in the neutron cells, as discussed in more detail in Chapters 6 and 7. The design of these loops is elaborated upon in Section 3.6

  - A **PTFE connector block** facilitates superconducting connections between the pickup loops and the lines connected to the input terminals of the SQUIDs. As the loops needed to be wound manually, and since some 1.5 m of wire is needed to reach the SQUID, some of which needed to be fed through superconducting shields, it was infeasible to have these connections completely intact from pickup loop to SQUID. A breakage introduces some modularity into the system, allowing the SQUIDs and pickup loop formers to be worked on and installed separately. The connector block can be used to attach lines from the two together, maintaining the superconductivity of the circuit.
  - **Superconducting capillary shields** for the twisted-wire pairs (TWPs) were necessary to restrict the spurious magnetic pickup along the input lines as they emerge out of the main shielded cryostat. The flux-removed solder becomes superconducting, protecting the twisted loops from any changes to the external field. The assembly consists of both single and triple-cored solder, carrying three wire pairs, interrupted by a lead-coated heater resistor, which can be used to remove the flux frozen in the flux transformer. The performance of these solder shields is examined in depth in Chapter 4.

- The **SQUID sensors** are mounted on copper supports placed symmetrically around the guide. They are of model CSblue as detailed in Section 3.3.

- **Readout cabling**

  - **Stycast-filled tubes** were custom built as cold feedthroughs between the six-way section and the 4.2 K connector box to handle the unconventional fittings. TWPs are fed through metal pipes of 26 cm in length bent once

over to fit into a space of 10 cm: the extra length was introduced to provide a more gradual temperature gradient, so as to reduce the heat load on the helium bath. The TWPs are encased at either end with PTFE sleeving to protect them, a mould placed over each end of the tube to form end caps, and a vacuum pump used to siphon degassed Stycast FT2850 through them. As Stycast contracts as it is cooled, pulling away from the sides of the pipe, the end caps, which overhang the edges of the tube openings, should ensure that no leak is introduced.

  - The **4.2 K connector box** is where connections between the TWPs from the Stycast-filled tubes and the **Kapton cables** are made. The connector box consists of thick copper plates that are thermally anchored to dissipate some of the heat load. The Kapton cables were designed to minimise the heat load and again stretch a distance of almost 1.5 m through stainless steel hosing, where their flexibility is a bonus (see Section 3.5). The cables terminate in the **room-temperature connector box**, where they plug into one of three specially-designed, circular PCBs that oppose each KF-50 flange on the cross piece. One of the PCBs, to which a filter board is attached, is used solely for signals to the calibration coils and heaters. All 16 lines have a 1 mH inductor and a 1 μF capacitor, together comprising a low pass filter of cut-off frequency $f = \frac{1}{2\pi\sqrt{LC}} \approx 5000$ Hz. The idea is to eliminate any high frequency components that could contribute to excess noise and interference.

- The **readout electronics** consist of the commercial, programmable flux-locked loop (PFL) units supplied by Star Cryoelectronics, and the control and ADC/-DAC (Digital-to-Analogue Converter) units produced in Oxford: all are connected to the room-temperature connector box via SCSI cables. The PFL units work much as detailed in Section 3.2, with a modulation frequency of 256 kHz. The multiplexed DAC controls signals to the calibration coils and heaters, as well as test signals administered to the SQUID input coils during the tuning process.

- The **control PC** is connected to the electronics optically via a **USB modem**, which enables portability and easy compatibility with different operating systems. The computer is equipped with **data acquisition (DAQ) software** that has been custom written by members of the Oxford group. There are a number of clients that enable management of the different parts of the SQUID system. For the SQUIDs themselves, the aptly named SquidClient allows the user to tune and prepare the SQUIDs for data-taking, whilst via DaqClient, the user can define the sampling rate and choose the record length, which determines the number of points that constitute each 'event.' Segmenting the data into events facilities easier manipulation and analysis, and a number of parameters are automatically calculated and stored for each event.

Beginning with the readout end, we will now elaborate upon a subset of these features, which saw major developments during the author's time on the experiment.

3.5 Readout cabling from 4.2 to 290 K

The main improvement in the readout was the introduction of new cabling and the room temperature connector box used to link these cables to the electronics, since the infrastructure below this, the Stycast-filled tubes and the 4.2 K connector box, was still functional and remained in place. The main priority of this cabling is to maintain a low heat load, and it was designed to supplant the woven manganin-copper cables running from liquid helium to room temperature, some of which had undergone damage in earlier phases of CryoEDM. The replacement cables are more robust and offer a low-cost alternative, with a similar level of noise performance to the original cabling. The Oxford group has a long tradition of manufacturing bespoke, cryogenic cabling for particle physics experiments, and the procedure used here invokes methods employed in similar applications [117]

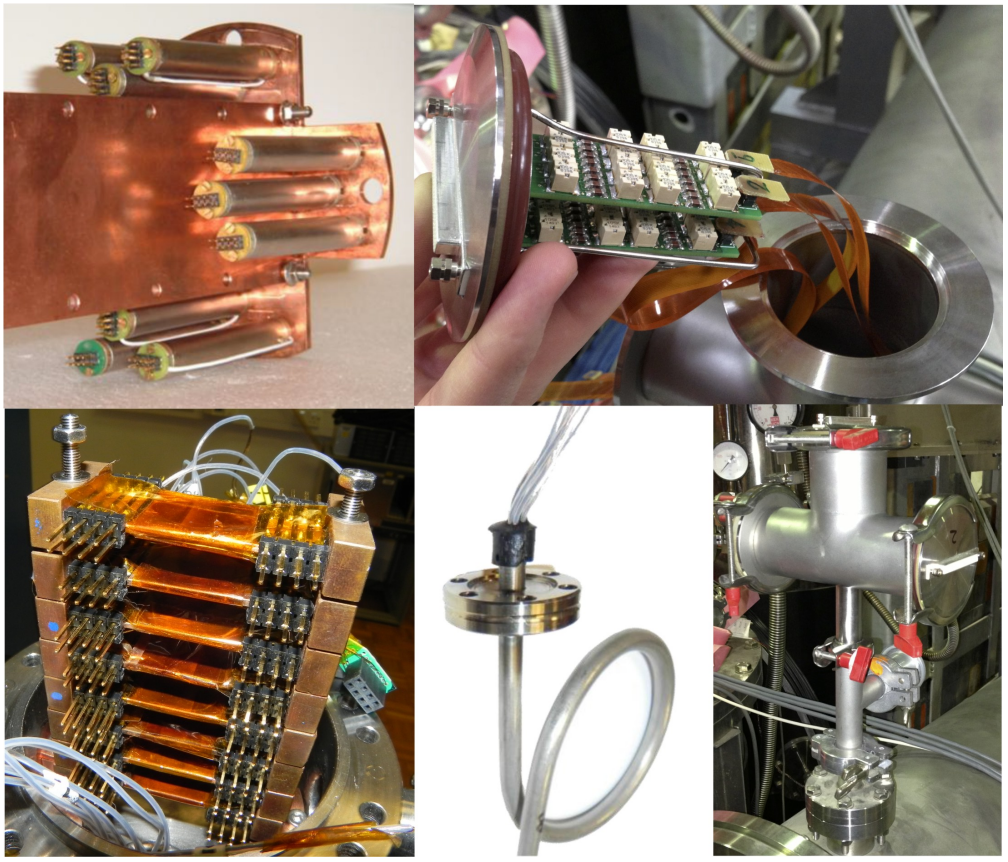

FIGURE 3.7: Various components of the SMS. Clockwise, from top left: the SQUIDs mounted on the copper guide tube; the filter boards for the calibration coils and heaters affixed on one of the three circular PCBs of the room temperature connector box; the 4.2 K connector box, with copper heat sinks; a Stycast-filled feedthrough with end cap; the room temperature connector box, mounted on the cryostat at the six-way section.

[118]. Production of these etched metal foil cables takes place entirely in-house, with the aid of the Photofabrication Unit (PFU).

The layout of the connections are drawn up in Eagle, a piece of CAD software that is commonly used for designing printed circuit boards. The final pattern will contain the blueprint of many cables within the dimensions of the UV exposure unit, allowing multiple cables to undergo the photolithographic stage of the manufacture simultaneously. The design is sent off to be made into a phototool, which acts much like a negative for a photograph (Figure 3.8). This phototool is loaded into the UV exposure unit above a Kapton, polyimide film, onto which a thin layer of stainless steel has been laminated. This is itself coated with a 'photoresist' layer,

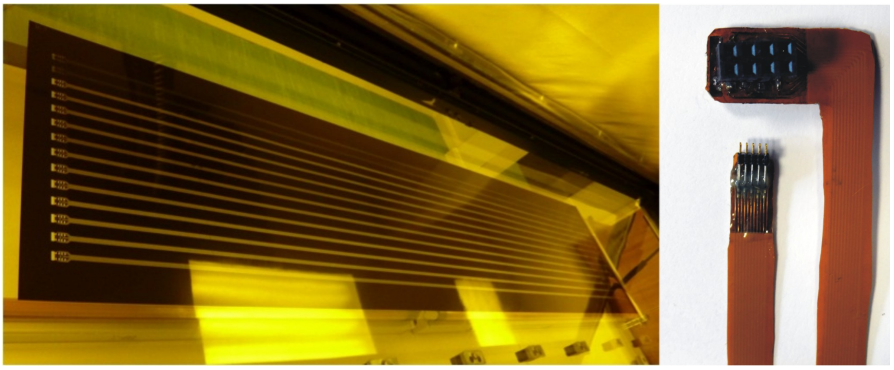

FIGURE 3.8: *Left:* The phototool for a set of straight cables that were used for SQUID tests in the Oxford laboratory. *Right:* The cables made for CryoEDM had a right-angle at one end, and horizontal, surface mount connectors at the other.

which will set when illuminated by UV light. The phototool therefore constrains this reaction to areas intended to form the tracks of the cables.

A developer solution is then applied to remove the photoresist from the unexposed foil, followed by an etchant solution to detach this foil from the polyimide substrate, leaving behind solely the conducting tracks and pads. Timing these processes is a fine art, as underexposure results in 'fuzzy' edges, resulting in varying thickness of tracks and possible connections between them that will undermine the performance of the cable, whereas overexposure can lead to undercutting of the foil that will create weaknesses if not breakages in the tracks, rendering the cable unusable. In this design, the tracks were made of identical width and spacing, such that they should require the same etching time, and sharp bends, which are prone to build up of etchant, were avoided. For instance, the right-angled connector ends were constructed with rounded corners, and the tracks were joined to the pads with a teardrop formation. Nevertheless, some breakages are inevitable, and it is necessary to check each line to make sure the connection is intact, and to clean up the track edges manually to ensure there are no bridges or shorts between them.

To strengthen the cables and protect the foil, the substrate is laminated with a second polyimide sheet to cover the middle, leaving the connector ends untouched.

The pads are first electroplated with nickel and then coated with a layer of copper to allow connectors to be soldered: an 8-pin, 2.54 mm Samtec surface mount socket is used to make the connection in the 4.2 K connector box, and a 10-pin, 1 mm Samtec horizontal mount plug for the room-temperature end.

A cable was required for each SQUID that would carry the bias current, feedback modulation, output voltage and heater signals in pairs of tracks, separated by grounding lines for crosstalk purposes. A further four cables were needed for the calibration coils and heaters, each responsible for four such devices. The end result is sixteen cables of 148 cm length, each with 11 stainless steel traces of 0.2 mm width and 0.5 mm pitch.

To estimate the heat load, we first need to establish the thermal conductivities, κ, of the materials involved. This is often quoted in the form $\kappa(T) = \alpha T^{\beta}\,\mathrm{W\,m^{-1}\,K^{-1}}$, with the parameters determined empirically. For Kapton between 5 and 300 K, this is $5.24 \times 10^{-3}\,T^{1.02}\,\mathrm{W\,m^{-1}\,K^{-1}}$ [119]. For stainless steel, as with all metals, the situation is complicated by the changing behaviour with temperature; it will also depend on the exact composition of the alloy. Fortunately, Ingleby reports resistivity measurements for this very material [117], with which we can use the Wiedemann-Franz law to approximate the thermal conductivity:

$$\kappa(T) = \frac{L_0 T}{\rho(T)}. \tag{3.17}$$

We use the theoretical estimate for the Lorenz factor of $L_0 = 2.45 \times 10^{-8}\,\mathrm{W\,\Omega\,K^{-2}}$, although its exact value will depend upon the chosen material. With a fitted $\rho(T) = (7.14 \times 10^{-7} + 3.13 \times 10^{-12}\,T^2)\,\Omega\,\mathrm{m}$ from [117], we see that until around 100 K, the constant term of ρ dominates, and that from Equation 3.17, κ in this region is roughly proportional to T. This is true at low temperatures, where heat transport is mainly via free electrons; at higher temperatures, electron scattering from other electrons, phonons or impurities becomes more important, and changes the temperature dependence of κ. For the range of $4.2 - 290$ K, we therefore use

the full expression of ρ and substitute this into Equation 3.17. The heating power, Q, can be calculated by

$$Q = \frac{A}{L} \int_{T_1}^{T_2} \kappa(T)\, \mathrm{d}T, \tag{3.18}$$

where A is the relevant cross-sectional area. With the thickness of the metal foil being $25\,\mu$m, the cross-sectional area of a track is $\approx 5 \times 10^{-9}\,\mathrm{m}^2$; likewise, each Kapton layer has a thickness of $75\,\mu$m, such that each of the ten Kapton strips between the tracks, with a width of about 0.3 mm, has an effective area of $\approx 4.5 \times 10^{-8}\,\mathrm{m}^2$, and each Kapton edge, with width around 1 mm, has an effective area of $\approx 1.5 \times 10^{-7}\,\mathrm{m}^2$. The total heat load is then given by

$$Q_{\mathrm{Total}} = Q_{\mathrm{Kapton}} + Q_{\mathrm{Steel}} = (123.8 + 45.6)\,\mu\mathrm{W} \approx \mathbf{170\,\mu W\ per\ cable.} \tag{3.19}$$

This ignores any additional heating caused by the resistive energy dissipation when a current flows along the tracks: however, since this is proportional to I^2, and any long-term current is on the order of microamps, this contribution should be negligible.

How does this compare to the manganin-copper cables? Again, the thermal properties of metals will depend upon their exact composition. Manganin is an alloy composed of supposed set percentages of copper, manganese and nickel, and thus its characteristics should not vary drastically from sample to sample. Copper, being an element, is different: any small contamination of the sample will impact greatly, providing another channel of heat dissipation via electron-impurity scattering. Figure 3.9 illustrates typical thermal conductivities for these materials, including those used for estimating the Kapton cable heat load, with a conservative estimate for copper. Although those of copper and manganin are higher than steel, as the previous cables were not built to the same specification, we cannot conclude from these alone that the heat load is necessarily reduced by using Kapton cables. The manganin-copper cables again use 11 tracks, nine of which

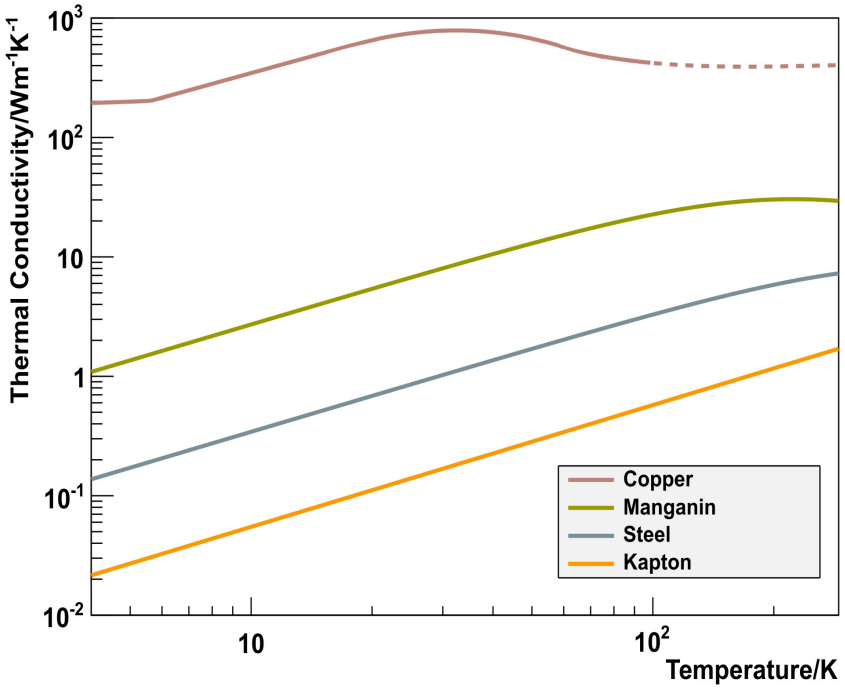

FIGURE 3.9: The variation of certain thermal conductivities over the relevant temperature range of the readout cables. That of steel and manganin have been calculated from the resistivities quoted in [117] and [120] respectively; the thermal conductivity of Kapton film is researched in [119]. The trend from copper has been adapted from [121], taking the results from industrial grade copper as a conservative estimate.

are manganin and two of which are copper, but all with the same diameter of 40 SWG$\approx 122\,\mu$m. Using the data from Figure 3.9, the total heat load for a single 2.5 m cable is ≈ 1.5 mW: moreover, if the cables were the same length as the Kapton, they would each have heat load ≈ 2.5 mW. The Kapton cables therefore have a heat load of around 7% of that of the manganin-copper cables, compared length for length. Replacing the magnanin-copper cables with the Kapton cables corresponds to a decrease in the evaporation rate of helium by $\approx 0.7\,l$ per day.

3.6 The Input Circuitry

Of the input circuitry, as outlined in Section 3.4, the pickup loops were the only components that were redesigned. However, extensive work was carried out to

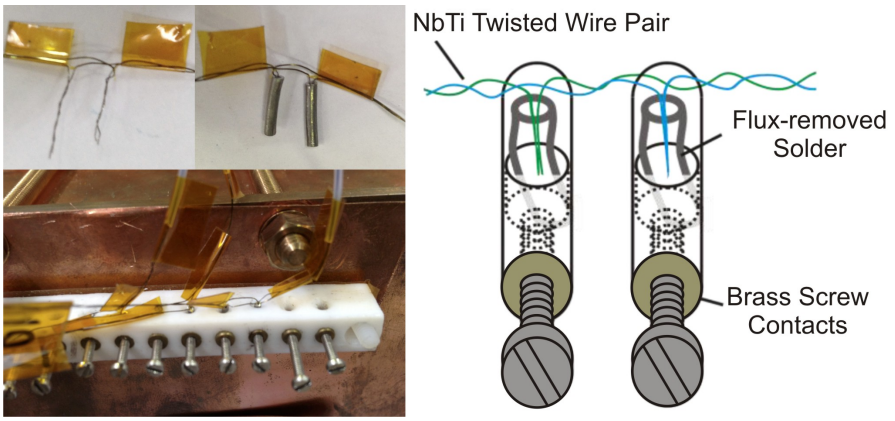

FIGURE 3.10: The photos and illustration show how superconducting contacts are made at the PTFE connector block, first enclosing the two ends of wire to be joined in fluxless solder and then compressing this to make the connection.

characterise the solder shields, as their efficiency had not yet been quantified (Chapter 4) and the PTFE connector block, as will now be discussed.

3.6.1 PTFE Connector Block

The PTFE block, as described in Section 3.4, is where connections between the pickup loops and the SQUID input lines are made. This is achieved by feeding the ends of these TWPs into a solder capillary, and clamping the solder to make the superconducting contact (Figure 3.10). The PTFE block was seen as a potential point of weakness in the circuit and the success rate of connections was tested. Out of the 20 connections made, one was seen to fail, with evidence of a resistance across the input. This resistance could be identified by applying a magnetic pulse to the input circuit, which induces a step in the coupled flux and therefore the current which flows. This current would persist indefinitely if the circuit were fully superconducting, but a finite resistance will cause it to decay exponentially according to

$$I(t) = \frac{\Phi}{L} \exp\left(-\frac{R}{L}t\right). \tag{3.20}$$

By finding the time constant $\tau = {}^{L}/_{R}$ and knowing the total inductance on the input, we could estimate the size of the parasitic resistance. Figure 3.11 shows a

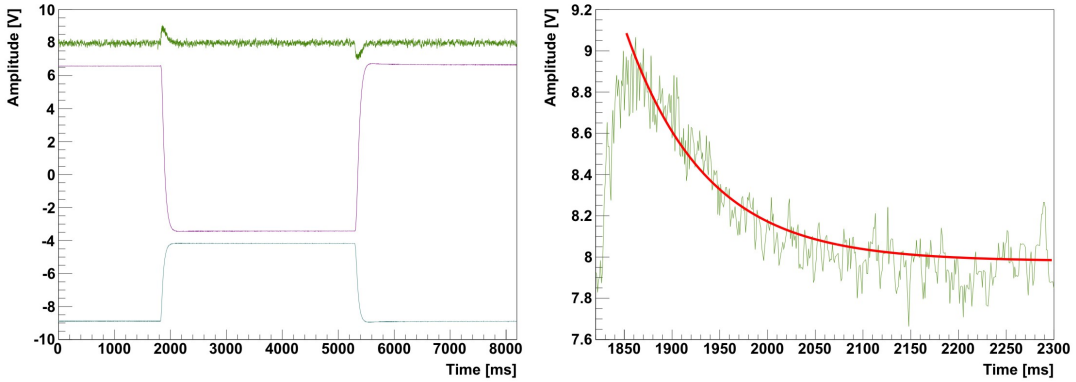

FIGURE 3.11: *Left:* The voltage response of three SQUIDs exposed to a square wave magnetic field, each with a pair of input lines connected together via the PTFE connector block. One (in green) does not track the field, and therefore shows evidence of a resistive input. *Right:* A close up of the response at the step edges, showing the decay of the voltage, and hence the current, across this resistive input.

fit to one such decay: with an input inductance of about 405 nH estimated using Equation 4.4, and a determined $\tau^{-1} = 11.8\,\mathrm{s}^{-1}$, this gives a resistance of $\approx 4.8\,\mu\Omega$. This was attributable to a resistive connection on the PTFE block, as on remaking the joins, leaving the SQUID input terminals untouched, the behaviour was seen to disappear. Of course, it is impossible to anticipate when a resistive contact is made until the system is cooled. A poor connection can be remedied on warming up the apparatus, but this is not an option for CryoEDM, where a cooldown takes several weeks and requires significant resources. The PTFE block is therefore a fallibility of the system, and the best precaution to take in preparation for a new run is to make fresh connections, as oxidisation over time can degrade the join, making the appearance of an input resistance more likely.

3.6.2 Pickup Loops

A set of 12 new pickup loop formers were constructed out of G10, an epoxy that is robust in cryogenic conditions and exhibits minimal thermal expansion under changes in temperature. Each loop is manually wound around a groove on the former with 127 μm diameter superconducting NbTi wire and held in place with

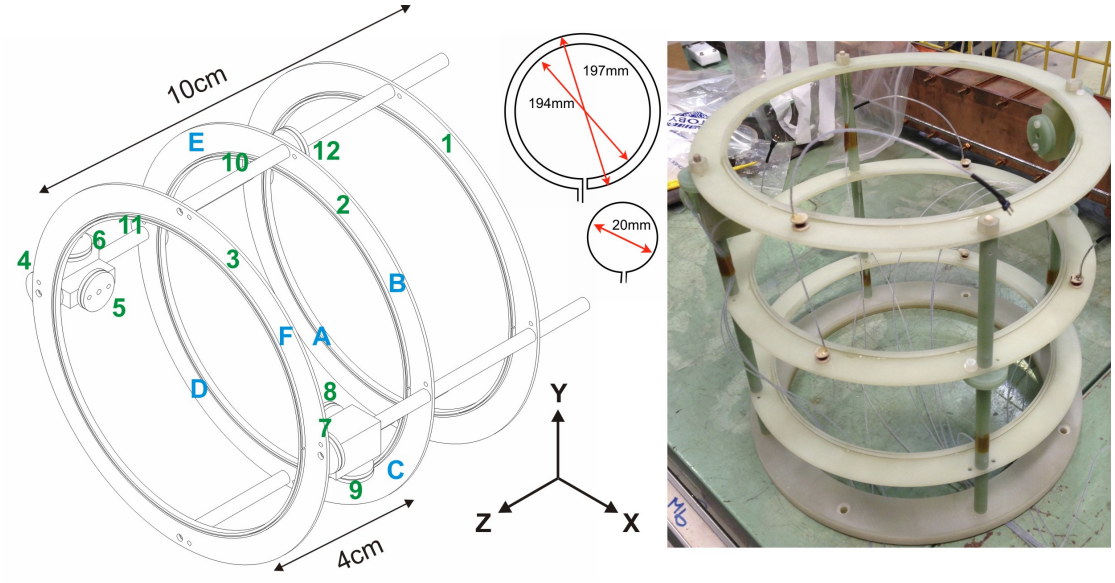

FIGURE 3.12: Schematic of the pickup loops with loops (1-12) and calibration coils (A-F) labelled, with the dimensions of the double, annular and single loops, and a photograph of the formers when put together.

Araldite. The wires must be carefully laid to prevent off-axis pickup, making sure that the loop forms a circle in a single plane. The loops are all single turn coils, but differ in shape: nine of the loops are of 20 mm diameter, whilst three are of a so-called double-loop formation with inner diameter 192 mm and outer diameter 196 mm. These latter loops consist of wire doubled back on itself as illustrated in Figure 3.12. As the pickup loop assembly is positioned ≈5 cm away from the neutral Ramsey cell centred on its axis, the configuration must be fitted around the neutron guide. The idea behind these double loops was that they circumscribe the guide and are coaxial with the cell, whilst maintaining an effective area similar to the smaller, conventional loops.

This area was chosen in light of tests conducted in the past by the group at the Laboratoire Souterraine à Bas Bruit (LSBB). Naïvely, the bigger the loop, the greater the flux which threads it and therefore the greater the induced current in the input circuit: on the other hand, the better it will act as an antenna for any RF interference. Loops of 20 mm diameter were shown to have the best signal-to-noise ratio out of the range of diameters investigated in the low-noise conditions

at LSBB [122], which should be similar to those at ILL.

Returning to Equation 3.15, we can estimate the magnetic sensitivity of these loops. In doing so, we can no longer ignore the inductance of the TWP. In fact, using Equation 4.4, the calculated 416 nH parasitic inductance of the TWP already exceeds the input coil inductance of the SQUID, meaning any attempt to optimise the pickup loop inductance and therefore size based on sensitivity alone is futile. The inductance of a 20 mm loop can be estimated as [123]

$$L_{\mathrm{PL}} = \frac{\mu_0 d}{2}\left[\ln\left(\frac{8d}{w}\right) - 2\right] \approx 65\,\mathrm{nH}, \tag{3.21}$$

with d the diameter of the loop and w that of the wire, which therefore yields a total inductance of $L_p = 481\,\mathrm{nH}$ and a sensitivity of $B_\Phi \approx 0.5\,\mathrm{nT}/\Phi_0$. However, we have used an implicit assumption in the derivation of Equation 3.15 that the flux threading the pickup loop will scale proportionally with its area. Whilst this is true of a uniform field, the flux linkage will, in general, depend on the form of the field and thus the magnetic environment. We return to Equation 3.21 and note that the logarithmic term shows relatively weak dependence on the diameter such that, over a small range, $L_p \propto d$ and therefore the area of the loop, $A \propto L_p^2$. In the case of a uniform field, substituting this back into Equation 3.15, we find that the larger the area of the loop, the better the sensitivity, with $B_\Phi \propto 1/\sqrt{A}$. To give a different example, let us consider a dipole field. In the simplest scenario, the dipole is coaxial with the loop, and is orientated in the same direction. In Chapter 6, Equation 6.7 shows that the flux which crosses the loop when the dipole is at great distance becomes $\Phi_p \approx C\frac{m}{d^2}$, where C is a constant, m is the dipole strength and d the loop diameter. Using Equation 3.14, we find that the minimum detectable dipole strength then becomes

$$m_{\mathrm{min}} = \Phi_0 \frac{(L_p + L_i)}{C M_i d^2} \propto A^{-1/2}, \tag{3.22}$$

i.e. the same dependence as with the uniform field. In contrast, for a very near-field

dipole, $m_{\min} \propto A$, suggesting smaller loops would provide better sensitivity. As the SQUIDs will be used to monitor changes in field due to the external magnetic environment, which is thought to consist of far-field dipoles, larger loops will be better suited for application in CryoEDM. The 20 mm diameter represents a compromise between increasing the magnetic sensitivity and maintaining a suitable spatial resolution and signal level with which to work with.

However, unless the fields are properly characterised, it is difficult to say with certainty whether this is the most suitable dimension, and where and in which orientation the pickup loops should be located. Six of the loops are configured as two three-axis setups, providing information in the orthogonal directions, but the majority of the loops have their planes perpendicular to the neutron cell axis since, as was discussed in Chapter 2, sensitivity to the gradients in this direction is paramount. Without knowing what the ideal pickup loop configuration should be, the setup devised here can be viewed as a test arrangement to determine which positions and orientations prove the most useful. Nevertheless, this is the configuration that will be adopted for the rest of this thesis, in particular, for simulating the SQUID signals (Chapter 6) and for exploring how to determine the fields in the neutron cells (Chapter 7).

3.6.3 Calibration Coils

A SQUID operated as described in Section 3.2 produces a voltage output, which must then be translated into an equivalent magnetic field reading in order to use the device as a magnetometer. In flux-locked loop mode, we know that the relationship between this voltage and the flux which crosses the sensing area. The calibration process involves applying known magnetic fields to the pickup loop to determine this fixed voltage response per unit of incident field (typically nanotesla), i.e. the conversion or calibration factor.

The calibration coils mounted on the pickup loop formers are energised to produce a calculable average field density across each loop. Each coil consists of 200 turns of copper wire wound around a former of 2 mm diameter and 3 mm length, giving 13 layers with an average diameter of $\sim$4 mm^2. The idea is that these are compact sources that could be approximated as dipoles at a distance, whose fields are relatively stable to deviations and are dominant over external sources during calibration. There are a number of complications, however, not least the positioning and number of these coils required to obtain an accurate calibration – McCann discusses a possible protocol [124]. However, depending on how the SQUIDs are to be used, an absolute calibration may not be strictly necessary. Whilst additional magnetic information would be useful, the application of the SQUIDs proposed in Chapter 7 will 'calibrate' the system using the neutrons, and thus this thesis will not go into the calibration process in any further detail.

Much as the pickup loops, the calibration coils will be connected via the PTFE block to long twisted-wire pairs that are fed through fluxless solder. The idea of these solder shields is to screen the lines from the ambient field, to stop the pickup of unwanted flux. These were investigated in some detail, with the results presented in the following chapter.

Chapter 4

The Characterisation of Solder Capillary Shields for the SMS

4.1 The Shielding Requirement for CryoEDM

SQUIDs are prone to increased noise levels when operating in an unclean magnetic environment, which motivates the use of Cryoperm shielding to mitigate the fields incident on the SQUID chip. In doing so, as will be seen in Chapter 5, the Cryoperm casing distorts the surrounding magnetic field which, if in the vicinity of the neutrons, could affect their polarisation. SQUIDs are also highly susceptible to electrostatic discharges, and thus they should be positioned away from the high voltage end of the experiment. Both of these factors place a design constraint that the SQUIDs should be mounted far upstream of the Ramsey cells, and thus they are secured onto the neutron guide in the six-way section. In order to have the best sensitivity to the B fields in the neutron cells, however, the pickup loops need to be placed as close to them as possible. Some $\sim$1.5 m of cable is then needed to make the connection between a pickup loop and a SQUID's input terminals, and this would need to emerge from the horizontal shields, leaving approximately 0.5 m exposed to ambient noise sources. Twisted wire pairs (TWPs) are used as,

by design, they are less prone to electromagnetic interference. However, even a perfect TWP could generate a signal that will couple to the SQUID.

To estimate the magnitude of the pickup of an ideal TWP and assess whether it is significant, we visualise it as a series of small, alternate pickup loops which detect the field in equal and opposite senses. Clearly in a non-uniform field, there will be a net flux-induced voltage. Here, we consider the simpler case of a uniform field, with an odd number of loops such that the contributions of all but that of the extra loop cancel. This sensing area can be treated as a rectangle with width equal to the wire diameter, 127 μm, and of height 3 mm. The signal read out from the SQUID, which we will label as S, will be proportional to the flux change in the SQUID loop and thus the current through the flux transformer:

$$S \propto I = \frac{\Phi}{L} = \frac{BA}{L}. \tag{4.1}$$

We can further write $S = CI$, where C (in VA^{-1}) is a conversion factor linking the current to the voltage output that depends upon the input inductance. To relate the signals generated in the TWP and the pickup loop of a single flux transformer, we note that the inductance is common to both, and that we therefore need only compare the flux threading the two elements. If the minimum field we wish to resolve is 0.1 pT using a pickup loop of 20 mm diameter, as was the sensitivity calculated using Equation 3.1, and the TWP is exposed to the estimated maximum 0.2 μT field change due to the IN15 experiment [1] the signal of the TWP referred to that of the pickup loop is

$$S_{\mathrm{TWP}} = \frac{\Phi_{\mathrm{TWP}}}{\Phi_{\mathrm{PL}}} S_{\mathrm{PL}} = \frac{0.2\,\mu\mathrm{T} \times (0.127 \times 3)\,\mathrm{mm}^2}{0.1\,\mathrm{pT} \times (\pi \times 10^2)\,\mathrm{mm}^2} S_{\mathrm{PL}} \approx 2430\, S_{\mathrm{PL}}. \tag{4.2}$$

This suggests an absolute minimum shielding factor of 2430 is necessary to reduce

[1] The IN15 experiment is a neutron spin echo spectrometer in the same experimental hall as CryoEDM. It is the biggest contributor to extraneous magnetic fields above the nominal background, generating periodic pulses that are registered at a magnitude of 0.2 μT by external fluxgates situated near the mu metal shields of CryoEDM.

the TWP signal to below the threshold we aim to be sensitive. In reality, our design requirements are more stringent than this since the field at ILL cannot be assumed to be so spatially uniform, and there will be imperfections in the TWP formation. As it is difficult to characterise a TWP, a more rigourous shielding factor objective needs to be established experimentally.

In preliminary tests in a cryostat at Oxford, bare TWPs connected to shielded SQUIDs were subjected to a uniform field generated by a pair of Helmholtz coils. The TWPs were around 13 cm long, of which 4 cm were exposed. A maximum signal size of 75.1 V was recorded when 11.5 μT was applied. The expected signal at ILL can be approximately related to this laboratory measurement as

$$S_{\mathrm{ILL}} = \frac{C_{\mathrm{ILL}}}{C_{\mathrm{Lab}}} \times \frac{\Delta B_{\mathrm{ILL}}}{\Delta B_{\mathrm{Lab}}} \times \frac{l_{\mathrm{ILL}}}{l_{\mathrm{Lab}}} \times \frac{L_{\mathrm{Lab}}}{L_{\mathrm{ILL}}} \times S_{\mathrm{Lab}}, \tag{4.3}$$

where we have used Equation 4.1. Here we have assumed, since this is an estimate, that the effective sensing area will scale will the number of bare loops and therefore the length of the exposed part of the TWP, although strictly this will only be true in a spatially uniform field, as discussed in Section 3.6.2. To estimate S_{ILL}, we first need expressions for the conversion factors and the inductances.

The inductance of the TWP is approximated as that of a return circuit of two parallel wires of equal diameter and fixed separation [125]:

$$L = \frac{\mu_0 l}{\pi}\left[\ln\left(\frac{d}{r}\right) + \frac{\mu_r}{4} - \frac{d}{l}\right], \tag{4.4}$$

where d is the separation, r the radius of the wire, μ_r its relative permeability and l the length of the TWP. We assume that the competing effects of the twisting (which would add some solenoidal component that would lower the inductance) and the greater length of wire used per metre (which would increase the inductance) would cause negligible deviations from this model. We further assume that the TWPs are tightly wound, so that $d \approx 2r$. In its superconducting state, the

permeability of the NbTi wire would be negligible, and so the $\frac{\mu_r}{4}$ term can be ignored. The estimated inductance from Equation 4.4 of a 13 cm TWP is therefore 36 nH, which together with the 350 nH of the input coil (Table 3.2) gives a total input inductance of 386 nH. Likewise, the calculated inductance of a 1.5 m stretch is 416 nH, which combined with the 350 nH of the input coil and the 65 nH of the pickup loop (Equation 3.21) yields 831 nH.

The input current cannot be measured directly, which makes C difficult to evaluate. However, the current that induces a flux quantum through the SQUID loop is given in the SQUID manufacturer's datasheet as the input coupling: $0.21\mu\text{A}\phi_0^{-1}$ (Table 3.2). Whilst this does not change significantly with different input loads, the voltage seen per flux quantum does. This could be measured by instigating an artificial flux jump, a readout artefact that corresponds to an integer number of flux quanta. This information is not available for the ILL system, so we will need to resort to an alternative way to estimate C. If we consider the circuitry of the SQUID (Figure 3.4), we note that the voltage we measure is a proportion of that supplied by the lock-in amplifier, which is behind the current that drives the feedback coil. This feedback coil provides a flux of $\Phi = MI$, where M is the effective mutual inductance. The voltage required for the feedback coil to couple one Φ_0 in the loop is therefore inversely proportional to M, assuming the effective impedances in the feedback circuit are unchanged. As this is the voltage response that we measure when the input circuit generates a Φ_0 in the SQUID loop, this is the voltage per flux quantum that we require, and we find $C \propto {}^1\!/\!_M$. The effective mutual inductance varies with the input inductance, L_i, according to

$$M_{\text{eff}} = M_0 \left(1 - \frac{L_{FB}}{L_{SQ} + L_i}\right) = 0.23 \left(1 - \frac{320\,\text{nH}}{350\,\text{nH} + L_i}\right), \tag{4.5}$$

substituting in the appropriate values. Then the ratio of the conversion factors becomes

$$\frac{C_{\text{ILL}}}{C_{\text{Lab}}} = \frac{M_{\text{Lab}}}{M_{\text{ILL}}} = \frac{1 - \frac{320}{386}}{1 - \frac{320}{831}} \approx 0.3. \tag{4.6}$$

Putting this information together, again using the value of $0.2\,\mu$T for the field at ILL, we get a calculated equivalent signal strength from the TWP of $S_{\mathrm{ILL}} \approx 6.3\,\mathrm{V}$. Given that the typical calibration factor for a SQUID at ILL (that is, the voltage equivalent for a given field sensed by the attached pick-up loop) is $3\,\mathrm{V\,nT^{-1}}$, this undesirable signal must be suppressed to below $0.1\,\mathrm{pT} \times 3\,\mathrm{V\,nT^{-1}} = 0.3\,\mathrm{mV}$. In other words, the shielding factor required is $\mathrm{SF} \approx {}^{S_{\mathrm{ILL}}}/_{0.3\,\mathrm{mV}} = {}^{6300}/_{0.3} \approx \mathbf{21000}$. This is again an approximate value, given the assumptions made and how sensitive Equation 4.3 is to any small changes in the parameters. We have also ignored the length of the wire that is within the mu metal enclosure, that will also pickup magnetic field, if at a much attenuated value. In reality, we would like to demonstrate a shielding factor at least an order of magnitude greater than this so that we can be sure the chosen shields meet the required specification.

4.1.1 Choice of Shielding

As we are interested in tracking long term changes in the field, and because high frequency noise can be filtered from the signal, we focus on the DC and not the AC shielding factor of our chosen shields. Several categories of material are commonly used for electromagnetic shielding purposes. Conductive materials, whilst useful for attenuating electric fields, are ineffective for slowly varying magnetic fields. The compensating, induced eddy currents are confined to a surface layer characterised by the skin depth, $\delta = \sqrt{\frac{\rho}{\mu\pi f}}$ (where ρ is its resistivity, μ its permeabililty and f the frequency of the incident electromagnetic radiation), so that the shielding factor improves for higher frequencies, and is comparable to the static value only when the wall thickness is much greater than δ. An alternative is ferromagnetic materials, which concentrate the field lines within them, allowing relatively few to permeate to their interior. They are typically used as broad spectrum shields, but as we have seen in Chapter 2, are difficult to characterise because of the dependence of μ on e.g. temperature, and because their effectiveness can be reduced by stress

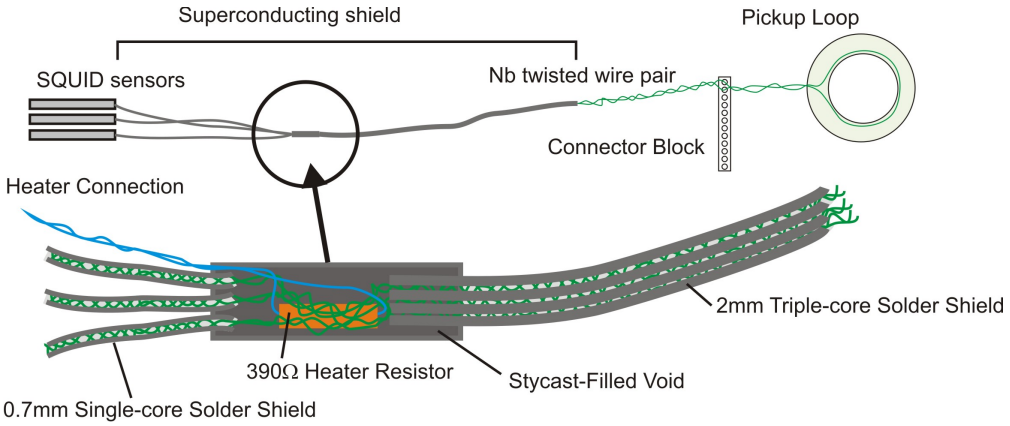

FIGURE 4.1: Schematic of a solder shield assembly, enclosing three pairs of SQUID input lines. This is the arrangement used to make the connections to the pickup loops, as seen in Figure 3.6 and detailed in Section 3.4.

and thermal cycling. As an example, a shield with a radius of 0.5 mm, a thickness of 50 μm and a relative permeability of $\mu_r = 50000$, would have a shielding factor, SF $\approx {}^{\mu_r t}/_{2r}$ (see discussion of Equation 4.7) of $O(10^3)$. Thus a ferromagnetic shield may fall short of the shielding requirement.

The most viable choice for a cryogenic experiment, and the one which is opted for here, is superconductors. Superconductors effectively show perfect diamagnetism (i.e. $\mu \sim 0$) and can offer near-complete shielding through expulsion of magnetic field from their volume, the Meissner-Ochsenfeld effect. To cause as little disturbance to the experiment as possible, the shields are designed as long capillaries to run closely around the wires. As the capillaries are cooled below their superconducting transition point, surface currents form on both the outer and inner surfaces. The first will vary to oppose subsequent changes to the applied field; the latter will persist to maintain and freeze in the flux through the centre, so long as the cylinder remains superconducting. This residual field should have no effect on a measurement as long as it is stable, since a SQUID is only sensitive to changes in field (Chapter 5).

Although there are commercial options, we have chosen to manufacture shields in-house from solder. The 0.7 mm diameter solder is heated in a bath of turpentine, dissolving the 0.5 mm diameter central core of rosin flux and leaving a hollow

lead-tin capillary intact, as described in greater detail in [126]. The shields are therefore produced at low cost, are more flexible and cause less damage to the HML enamel coating of the TWPs. It was found that lengths of up to 30 cm could be prepared in this way. Longer lengths were made from 2 mm diameter triple core solder, which could be more effectively treated. The two sections of shielding meet with a lead cylinder enclosing the join (Figure 4.1), which also encases a 390 Ω resistor used as a heater to revert the TWPs into their normal state, whose efficacy is investigated in Section 4.4.1.

There are two factors contributing to these solder shields: their cylindrical geometry, and their superconducting properties. To address the first, we can solve Laplace's equation, using the continuity equations to specify the ratio between the external and internal fields [127]. The shielding factor of an infinite hollow cylinder in a radial or transverse field is given by

$$\mathrm{SF}_{\perp}^{\infty} = \left| \frac{(\mu_r + 1)^2 (\frac{r}{R})^2 - (\mu_r - 1)^2}{4\mu_r} + 1 \right|, \tag{4.7}$$

where μ_r is the relative permeability of the material, r is the inner radius and R the outer radius of the tube. Expressing $R = r + t$, where t is the thickness of the shield, in the limit of large μ_r and $t \ll r$, the transverse shielding factor becomes $\mathrm{SF}_{\perp} \approx {}^{\mu_r t}/_{2r}$. The shielding factor for an axial or longitudinal field can be expressed as

$$\mathrm{SF}_{\parallel} = 4N\,\mathrm{SF}_{\perp} + 1, \tag{4.8}$$

where N is a geometry dependent demagnetisation factor. For a cylinder, N takes the form

$$N = \frac{\alpha \ln(\alpha + \sqrt{\alpha^2 - 1} - 1)}{(\alpha^2 - 1)^{3/2}} - \frac{1}{\alpha^2 - 1}, \tag{4.9}$$

where $\alpha = {}^{L}/_{2R}$ and L its length. $N \to 0$ as $\alpha \to \infty$, so an infinitely long cylinder cannot offer protection against longitudinal fields.

For a perfect superconductor, μ_r is zero and the predicted shielding factors from

Equations 4.7 and 4.8 are theoretically infinite: therefore, one must look for other contributions. Equation 4.7 remains a good approximation to the attenuation by a finite cylinder if the applied field is contained to within a radius of the ends. In practice, one needs to also consider any field penetration at the openings. The fields at a distance z into the shield from an open end are given by

$$H_{int\perp} \approx \frac{1}{3}\mathrm{e}^{-k_T(z/r)}H_{ext\perp}$$
$$H_{int\parallel} \approx 2.6\sqrt{\frac{L}{R}}\mathrm{e}^{-k_L(z/r)}H_{ext\parallel},$$

where $k_T = 3.83$ is the first zero of the Bessel function of the first kind, J_1 and $k_L = 2.41$ the first zero of the function J_0 [128]. It is this contribution that will dominate for a perfect superconductor. Indeed for our solder shields, the transverse field will fall off exponentially as $B(x) = e^{-x/\lambda_L}$: with, akin to the skin depth, a London penetration depth, λ_L, of $O(10\,\mathrm{nm})$, negligible field should traverse the $\sim 100\,\mu$m thick wall. In our tests, the setup is configured so that these end effects do not come into play. Despite this, we may still expect to see a signal for reasons which will now be explained.

The solder we used is composed of eutectic mixture of 60 parts tin and 40 parts lead, which has been found in one study to have a critical temperature, T_c of 7.05 K, and a critical field, B_c of 80 mT at 1.3 K [129], both above which the material would no longer be superconducting. B_c varies with temperature as

$$B_c(T) = B_c(0)\left[1 - \left(\frac{T}{T_c}\right)^2\right], \tag{4.10}$$

where $B_c(0)$ is the critical field at absolute zero. This gives $B_c(4.2) \approx 53\,\mathrm{mT}$, compared with an estimated maximum attainable field of around 4 mT from our apparatus.

However, in reality, the transition between superconducting and normal behaviours

is not so black and white for the solder capillaries. First, geometric effects can affect the field distribution so that even in a uniform field, some areas of the cylinder will experience higher fields than others. This can result in the formation of an intermediate state consisting of normal and superconducting domains. This state exists at a field $B_c' < B < B_c$ where $B_c' \approx (1 - N)B_c$, where N is the demagnetisation factor introduced in Equation 4.9. For our shields, this is negligible, with $N \approx 4 \times 10^{-5}$ for a 26 cm shield.

Second, solder belongs to the class of type-II superconductors which are distinguished by their type-I counterparts by the existence of a mixed state known as the Shubnikov phase [130]. In this regime, some magnetic field can penetrate the material, forming flux vortices on its surface: this typically allows it to retain superconducting properties at fields much higher than type-I superconductors. The onset of this behaviour is marked by a lower critical field, B_{c1}, relabelling the field at which it becomes normal, B_c, as B_{c2}. No figures could be found for B_{c1}, but it is probable that, as our maximal field is much below B_{c2}, the solder will be in the Meissner phase for all of our measurements.

The main reason why a solder capillary may be an imperfect shield is that it is unlikely to be an ideal superconductor, as we have treated it thus far. Mechanical defects or material inhomogeneities inherent in the manufacture or introduced by threading and manipulating a shield could lead to non-uniform properties along its length. Therefore to be sure that the shields meet the required standard, their shielding factors must be determined empirically if they are measurable, or at least minimum bounds set on their values.

Coils	Dimensions (mm)	No. of Turns	Separation (mm)	Central Field ($mT\,A^{-1}$)	Max. Field ($mT\,A^{-1}$)
X	35×35	300	37	7.34	9.72
Y	28×35	300	39	6.32	10.02
Z	⌀=54.5	500	48	9.74	15.46

TABLE 4.1: Specifications of the X, Y and Z coils. The dimensions given are not of the formers but the average position of the coil windings. The last column gives the maximum point field on a TWP produced by a given pair of coils.

4.2 Measuring Pickup Factors

4.2.1 Experimental Setup

An experiment was constructed to measure the magnetic pickup of TWPs both with and without the solder capillaries. An assembly of three pairs of coils aligned in mutually orthogonal directions is used to generate the test field. This was designed to be mounted on a dipstick with the SQUIDs instead of being fixed at the bottom of the cryostat. Although this latter alternative would allow for larger coils, difficulties were foreseen in estimating the positions of the TWPs relative to the coils, and hence the magnetic field strength they are exposed to. The coils were therefore dimensionally constrained to fit within the 5.7 cm diameter of the cryostat mouth. Theoretical models suggested the proposed coil geometry could generate sufficient fields, providing enough turns were included and a sufficient current could be delivered. A 2.5 mm overhang was introduced around the perimeter of the G10 coil formers to accommodate the turns of 234 μm diameter enamelled copper wire. After being wound using a mechanical coil-winding machine, GE varnish was applied at the corners to prevent the coils from unravelling. Table 4.1 lists the coils dimensions, their number of turns, and the maximum estimated fields they can produce, and Figure 4.2 illustrates the final design.

Six SQUIDs are fixed on a mounting plate attached 4 cm above the top of the coil assembly. Kapton cables lead from the SQUID outputs to the readout electronics, which are the same as those used in the experiments at ILL and LSBB. A 23 cm

twisted wire pair (TWP) made from 127 μm diameter NbTi wire is attached to each SQUID input, which is to emulate a similar section of such wires running from the pick-up loops to the SQUIDs in CryoEDM. The length is long enough to allow sufficient distance between the SQUIDs and the coils, to prevent distortive effects due to the Cryoperm shielding, but short enough to meet with the cryostat dimensions and the lengths of solder available. Each TWP was coated with a layer of Stycast to prevent unravelling of the structure with successive measurements, which would affect the size of the sensing area and therefore the signal strength. The TWPs are fed through holes in the Z coil formers that loosely hold them in place: the wires are then assumed to be running parallel to the z-axis.

In the shielded tests, the TWPs were threaded through fluxless, single-cored solder shields of 24 cm and 26 cm in length, and triple-cored shields. The end of the TWP lies a few centimetres into the shield to prevent edge effects from becoming significant. The shield permeates into the niobium end cap surrounding the SQUID chip to provide almost continuous shielding, with a $\sim$1 cm length of wire remaining exposed to allow connections to be made to the SQUID input terminals.

With all the SQUIDs fixed in place and the inputs wired, the dipstick is cooled to liquid helium temperature. After the SQUIDs have been tuned and it has been established that they are tracking magnetic field, test signals can be sent to each of the sets of coils in turn. A signal generator is used to drive a pair of coils with a 300 ms period square wave of adjustable amplitude. An edge time is chosen to be long enough such that $^{dB}/_{dT}$ does not exceed the slew rate of any of the SQUIDs. The applied voltage is scanned through several values and the signal responses on the SQUIDs recorded. A C++ macro is used to find the sizes of a selection of the voltage steps for each given channel by averaging 100 ms worth of data either side of the transition. These are again averaged to give a final value of the signal for that coil voltage or field. By measuring the combined resistance of the signal generator and the coil circuit, the voltage applied can be converted to a current, which in turn can be used to find the generated fields. With a

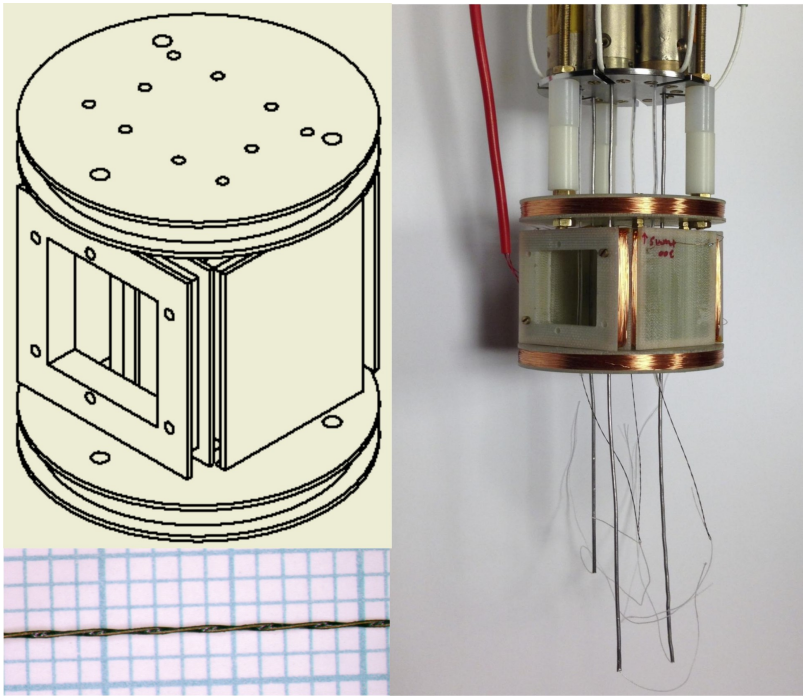

FIGURE 4.2: A photograph to demonstrate the size of the windings on the TWP (against 1 mm graph paper), a schematic of the three-axis coil assembly, and a photograph of it mounted on the dipstick. Holes in the Z coil formers allow the solder shields and bare TWPs to pass through the main field region; 4 cm of spacers keep the SQUID baseplate away from the top of the coil configuration.

maximum deliverable peak-to-peak voltage being 20 V and the internal resistance of the generator 47 Ω, the current could not exceed around 0.5 A.

To compare measurements at different applied fields, the signal responses must be normalised. The signal strength per nT of field seen by the TWP provides the definition of our pickup factor. As we are sensitive to the fields within the shields only through the TWP signal, the shielding factors are specified as the ratio of the unshielded pickup to that in the shielded case. The pickup factor is computed for each measurement, and the average for each TWP is used in subsequent analysis.

4.2.2 Converting Signals Into Pickup Factors

Theoretical calculations, like those in Section 4.1, are useful for giving ballpark estimates of the signals we may expect. However, such a treatment is inadequate in our laboratory setup as it overlooks several points, namely:

1. The orientation of the TWP loops with respect to the coils is not known, and will vary from run to run;

2. The TWPs are not perfectly wound and thus will have loop areas which vary not only between samples, but along the length of a single TWP itself;

3. The fields generated by the coils cannot be taken to be uniform in the TWP region, nor have a well defined gradient, and may have significant off-axis components;

4. The TWP may not necessary be straight, and its loop formation may change with repeated use.

If the first can be addressed with a simple rotation of the entire TWP, an overall pick-up factor can be defined which would be expected to be maintained in any measurement: $P = \sqrt{P_x^2 + P_y^2 + P_z^2}$. Since the TWP is roughly aligned with the z-axis, one can further assume that the transverse and longitudinal components should be individually conserved, allowing the definition of $P_T = \sqrt{P_x^2 + P_y^2}$ and $P_L = P_z$ respectively.

As previously discussed, imperfections in the TWPs, such as uneven twisting of the loops, kinks or loose windings make it difficult to determine the flux which threads a TWP. The figure of merit used is the integrated field along its length, which permits comparisons to be drawn between the same TWP placed in different positions, providing that its loops remain unchanged from run to run. With the origin taken to be the centre of the coil setup, the integrated fields are calculated for each of the positions of the TWPs. Those generated by the rectangular X and Y coils are found by integrating the formulae (see Appendix A) for the field at a point from $z = -0.075\,\mathrm{m}$ and $z = 0.065\,\mathrm{m}$ i.e. from the end loop of the TWP to the start of the Cryoperm shielding. The equivalent expressions for the circular Z coils are difficult to integrate, so instead the field at regular intervals along the TWP are computed and added together.

Although the predominant field generated by a pair of coils will be in the coaxial direction, given the positioning of the TWPs in the x-y plane, there may well be some significant off-axis contributions. The pickups can be related to the signals via the following matrix representation:

$$\begin{pmatrix} S_1 \\ S_2 \\ S_3 \end{pmatrix} = \begin{pmatrix} \beta_x^1 & \beta_y^1 & \beta_z^1 \\ \beta_x^2 & \beta_y^2 & \beta_z^2 \\ \beta_x^3 & \beta_y^3 & \beta_z^3 \end{pmatrix} \begin{pmatrix} P_x \\ P_y \\ P_z \end{pmatrix}$$

where $S_{1\cdots3}$ represents the signal seen in a given TWP produced by the 'X', 'Y' and 'Z' coils respectively, $\beta_{x\cdots z}^{1\cdots3}$ denotes the integrated magnetic fields $\int_{-0.075}^{0.065} B_{x\cdots z}^{1\cdots3} dz$ produced in the x, y, z directions by these three coils and $P_{x\cdots z}$ are the x, y, z pickup factors for the given TWP: the β matrices for each of the six TWP positions are listed in Table 4.2. The pickup factors may then be solved for by

multiplying each side by the inverse of the β matrix. Averaging those obtained for each run with the same configuration yields an overall pickup factor in $\mathrm{VnT^{-1}m^{-1}}$, which we use to calculate the shielding factors. Note that these cannot simply be scaled to give the estimated magnetic pickup for a longer section of TWP, since this will be dependent on the inductance as in Section 4.1.

It is difficult to account for the uncertainties due to the loop formation of the TWPs. On first inspection, all the TWPs have similar winding density (see Figure 4.2) and there is no visible difference in the size of the end loop. However, the magnetic pickup will depend upon the exact dimensions and orientations of these loops which is impossible to quantify, although unravelling of the TWP with continued use may be noticeable in a gradual increase in the average pickup factor.

Field inhomogeneities will also contribute to the uncertainty in the integrated fields. The coils cannot be placed in the Helmholtz position if the TWPs are to run between them, and thus the variation of field in this region is greater than expected. Other formations involving three or more coils were considered, since

$$1 - (-10.4,\ 6) \qquad \begin{pmatrix} 0.2908 & 0.0368 & 0.0006 \\ 0.0712 & 0.1674 & 0.0002 \\ -0.0033 & 0.0019 & -1.1563 \end{pmatrix}$$

$$2 - (0,\ -12) \qquad \begin{pmatrix} 0.1873 & 0 & 0 \\ 0 & 0.2927 & -0.0005 \\ 0 & -0.0038 & -1.1563 \end{pmatrix}$$

$$3 - (-10.4,\ -6) \qquad \begin{pmatrix} 0.2908 & -0.0368 & 0.0006 \\ -0.0712 & 0.1674 & 0.0002 \\ -0.0033 & -0.0019 & -1.1563 \end{pmatrix}$$

$$4 - (0,\ 12) \qquad \begin{pmatrix} 0.1873 & 0 & 0 \\ 0 & 0.2927 & -0.0005 \\ 0 & 0.0038 & -1.1563 \end{pmatrix}$$

$$5 - (10.4,\ -6) \qquad \begin{pmatrix} 0.2908 & 0.0368 & -0.0006 \\ 0.0712 & 0.1674 & 0.0002 \\ 0.0033 & -0.0019 & -1.1563 \end{pmatrix}$$

$$6 - (10.4,\ 6) \qquad \begin{pmatrix} 0.2908 & -0.0368 & -0.0006 \\ -0.0712 & 0.1456 & -0.0002 \\ 0.0033 & -0.0058 & -1.1563 \end{pmatrix}$$

TABLE 4.2: The integrated field β matrices for the six TWP positions, listed with their x and y coordinates in mm.

they could improve uniformity and boost the magnitude of the fields [131], but were deemed unsuitable given the complexity they would add to the construction of the setup, and the analysis of the data.. Figure 4.3 illustrates the typical field variation in the x-y plane when energising a pair of coils. Movement of the TWPs by a millimetre can cause a change in the field at the TWP position of up to 5% which in turn would affect a given pickup value by up to 5%. However, such deviations in position will be randomly distributed and will be caused not only by a shift or misalignment of the TWPs, but also from the introduction of bends in the wire that cannot be accounted for.

Since the impact of many of these factors is impossible to assess, we refrain from quoting an uncertainty on each individual measurement. The spread in measurements for a given TWP is used, as this will encompass some of these effects.

4.2.3 Unshielded Pickup Factors

Table 4.3 summarises the average pickup values found in the transverse and longitudinal directions. Many of the TWPs have standard deviations that are consistent

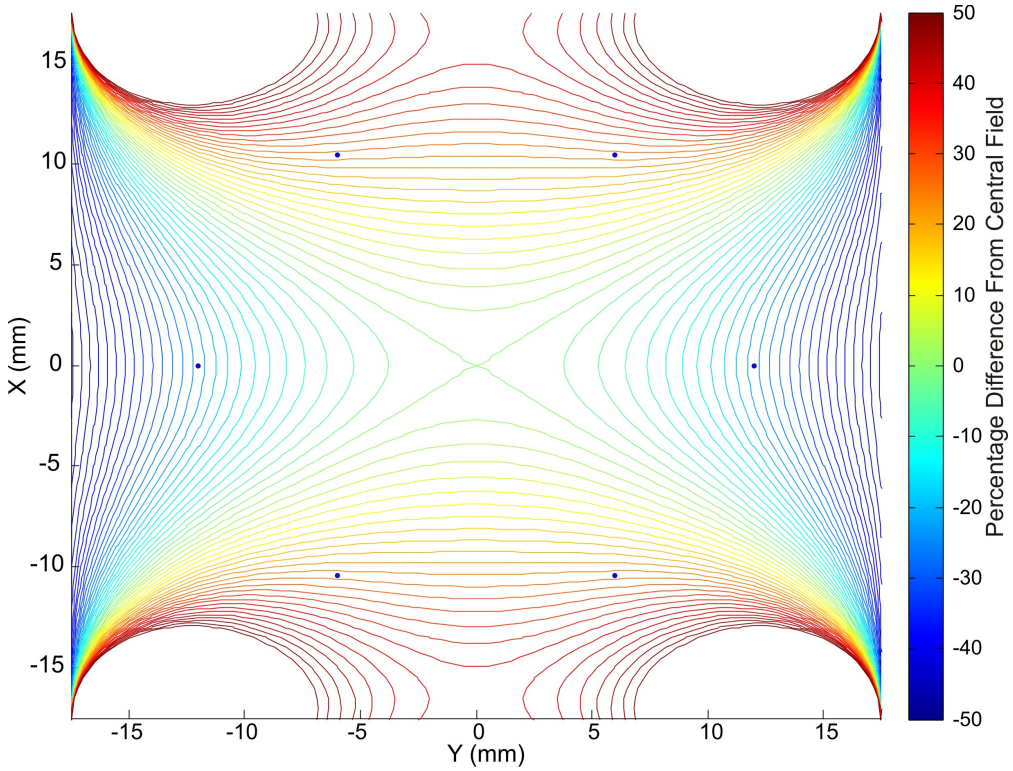

FIGURE 4.3: A contour map of the field strength of the X coils as a percentage of the central field within an x-y cross-section at z= 0, illustrating the uniformity of the field between the coils. The contours represent 2% variations, and the six TWP positions are marked.

TWP	P_T ($\mathrm{VnT^{-1}m^{-1}}$)	P_L ($\mathrm{VnT^{-1}m^{-1}}$)
1	0.300 ± 0.057	0.011 ± 0.011
2	0.247 ± 0.069	0.015 ± 0.005
3	0.080 ± 0.009	0.003 ± 0.002
4	0.300 ± 0.013	0.015 ± 0.017
5	0.104 ± 0.023	0.006 ± 0.006
6	0.143 ± 0.046	0.008 ± 0.006

TABLE 4.3: The average transverse and longitudinal pickup factors and their standard deviations for the six unshielded TWPs that were tested.

with zero longitudinal pickup, as might be expected if they truly are aligned with the z axis. Proportionally, these uncertainties are larger than those for the transverse direction, which do not exceed 30% of the average value, since this pickup is more sensitive to movement, with a small misalignment resulting in a large relative increase in the field seen.

Figure 4.4 shows the typical spread of measurements for a given TWP. Note that there is no observable consistency in the measurements of the pickup in the x or y directions alone, but that when combined as P_T they hold relatively steady. The graph can be divided into sections representing a series of runs where the TWP was untouched, which show a greater degree of similarity, e.g. the stability for runs 9-13 translates to $P_T = 0.335 \pm 0.004$. The differences between these series imply that any slight movement or change of TWP structure from rewiring the system is enough to affect the pickup factor.

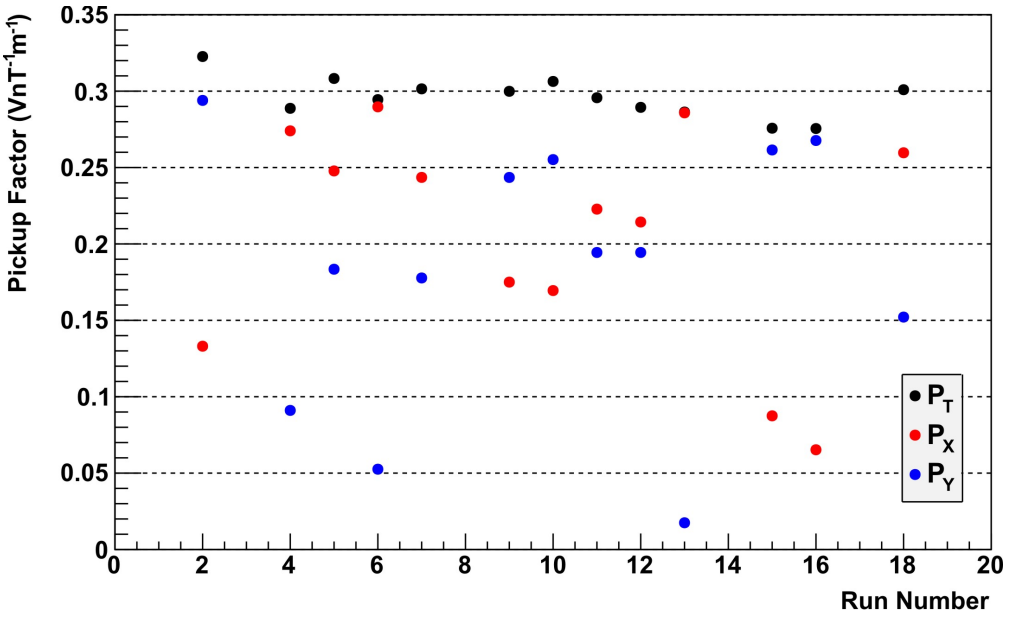

FIGURE 4.4: The average transverse pickup values for a given TWP (TWP 3) for each run. $P_T = \sqrt{P_x^2 + P_y^2}$ is particularly stable over time.

TWPs were also moved from one SQUID to another in a couple of the runs to ensure that the pickup seen is inherent to the TWP and not in part due to the SQUID it is connected to or its position between the coils. These showed variation consistent with other measurements, and so have been included in the final results. The distinctly different transverse pickup factors in Table 4.3 suggest that they are highly dependent on the formation of the TWP and amalgamating the values into one overall unshielded pickup factor would be unwarranted: the pickup values for the shielded tests must therefore also be kept separate.

Lastly, it is perhaps instructive to express these unshielded pickup factors in the form of an effective pickup area, which can be directly compared to the size of a typical ILL pickup loop. Using Equation 4.3, the equivalent area, for either the transverse or longitudinal direction, can be found as

$$A_{\text{eff}} = \frac{IL}{B_{\text{av}}} = \frac{PlL}{C}, \tag{4.11}$$

where l is the exposed length or wire, $B_{\text{av}} = \frac{1}{l}\int_l \vec{B}\cdot \mathrm{d}\vec{l}$ is the average magnetic field over the length of the TWP and using the fact that the signal, $S = CI = P\int_l \vec{B}\cdot \mathrm{d}\vec{l}$, with P the relevant pickup, all in the given direction. We use again, as in Section 4.1, that the input coupling is $0.21\mu\mathrm{A}\phi_0^{-1}$ and can measure the size of a flux jump to find the voltage per flux quantum. For an unshielded, 23 cm long TWP, of which 14 cm is exposed, we find $\sim 3.7\,\mathrm{V}\phi_0^{-1}$ (as opposed to $\sim 4.0\,\mathrm{V}\phi_0^{-1}$ for a shielded TWP and $\sim 9.9\,\mathrm{V}\phi_0^{-1}$ for a shorted input), which gives

$$A = 0.20\times 10^9\,\mathrm{VnT^{-1}m^{-1}}\times 0.14\,\mathrm{m}\times\frac{0.21\,\mu\mathrm{A}\phi_0^{-1}}{3.7\,\mathrm{V}\phi_0^{-1}}\times 414\,\mathrm{nH}\approx 6.3\times 10^{-7}\,\mathrm{m}^2,$$

or $\approx$2.7 mm^2 per metre in the transverse direction, using the average P_T from Table 4.3. Over a half metre of exposed wire, as at ILL, the effective transverse area is more or less equivalent to four times the estimated size of a single loop in the TWP. Similarly, longitudinally, we find $A \approx$0.1 mm^2 per metre, or a sixth of a TWP loop for half a metre. For an input circuit terminated with a 20 mm pickup loop, if the field were uniform and the entire circuit exposed, the TWP pickup would constitute $\sim$0.9% that of the loop for a transverse, and $\sim$0.04% for a longitudinal field.

4.3 Determining Shielding Factors

4.3.1 Shielded Pickup Values and Shielding Factors

Computing the pickup factors for the shielded tests is complicated by the fact that a signal can exist at a level comparable to the noise which cannot be identified by eye, even when running the SQUIDs in high sensitivity. By running unshielded tests in parallel, we can pinpoint where the voltage changes should occur and use the same technique to find the step in the response. In most cases, the signal change is the same sign for all upwards steps, and likewise for all downwards steps, which suggests a real signal is present. The distribution of these values is dominated by the white noise (Figure 4.5): the mean and standard error are extracted from a Gaussian fit to the data and used to qualify the size of the signal and the uncertainty on its value.

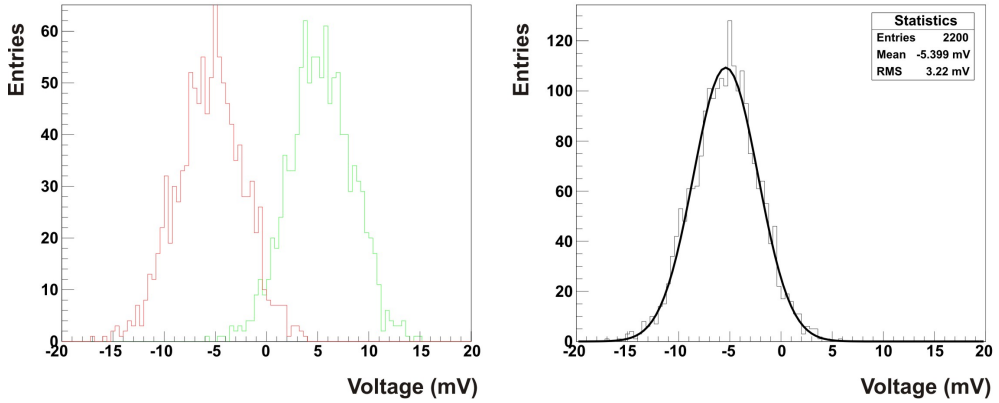

FIGURE 4.5: The voltage responses for a particular shielded channel after applying a step function of 20 V to the Z coils. There is a clear separation between the peaks of the distribution for the up (green) and down (red) steps. Combining the up and down steps sizes, a Gaussian is fitted to the histogram and the relevant statistics, the mean and standard error, are noted.

The resultant pickup values are dependent not only on the TWP, as discussed in the previous section, but also on the shield. There is no noticeable difference between the values of the 24 cm and 26 cm shield, nor between the single and triple-core samples. None of the pickup values exceed $10^{-7}\,\mathrm{VnT^{-1}m^{-1}}$, some 10^5 times smaller than the lowest unshielded value. However, even with the exact

same configuration, the shielded pickup factors are seen to vary by two orders of magnitude. Tests with shorted inputs and 1 cm long TWPs show that there is still measurable pickup within the niobium can, which limits our shielded pickup numbers. The variability of the values could therefore be explained by magnetic pickup on the exposed sections of TWPs between the end of the solder shield and the SQUID's input terminals. The shielding factors we present here are thus limits, where we use the minimum pickup in the shielded case (Table 4.4).

Shield	Type	SF_T ($\times 10^7$)	SF_L ($\times 10^7$)
1	24 cm Single-core	4.04	0.334
2	26 cm Single-core	24.0	5.35
3	26 cm Single-core	4.18	6.65
4	Triple-core	24.0	11.7
5	Triple-core	1.79	0.288

TABLE 4.4: The maximum shielding factors for each shield in the transverse and longitudinal directions.

The AC dependence of the shielding factor was not explicitly tested, but a comparison was made in the frequency domain between shielded and unshielded TWPs and SQUIDs with shorted inputs, as a measure of the intrinsic noise of the SQUID. The SQUID system was left recording overnight, and the data analysed using a fast-Fourier transform algorithm (FFT) to produce spectra of the signal size versus frequency. Figure 4.6 shows an example from one such period of data acquisition. The noise level is seen to be attenuated by the solder capillaries to the base level of the SQUID below 1 Hz, our main frequency range of interest. In the white noise region at higher frequencies, the shielded channels show signs of the 50 Hz mains noise and harmonics, as well as other frequency peaks also seen in the unshielded case. The origin of these is a mystery, since there is no noticeable pickup by the shielded TWPs on application of a sine wave at numerous frequencies, suggesting that this could be due to some mechanism of non-magnetic coupling. This may be undesirable, but is of little consequence when the focus is on long-term drift, and when the signal is time-averaged between cycles and can be digitally filtered.

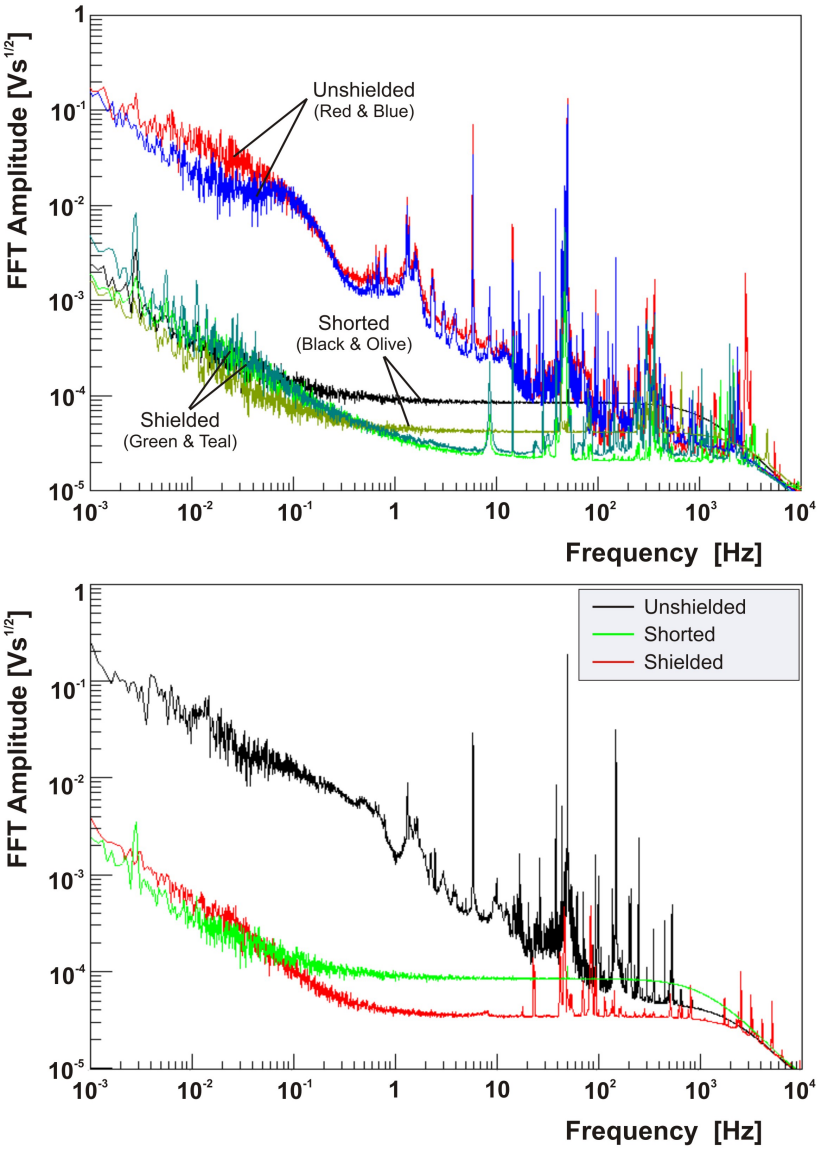

FIGURE 4.6: Overnight noise spectra, *top:* for six differently configured SQUIDs with serial numbers 0568 (red), 0020 (blue), 0562 (black), 0237 (olive), 0406 (teal) and 0073 (green); *bottom:* for the same SQUID (0562) run in three different modes. The shorted SQUIDs have a small wire between the input terminals, with negligible inductance, and thus its noise spectrum is approximates the intrinsic SQUID noise.

As the baseline noise is influenced by the inductances within the SQUID, to make a proper comparison, a single sensor was operated in each of the three configurations on separate nights in more or less the same magnetic environment (Figure 4.6). The difference between the white noise level for the shielded and shorted cases is then obvious, and can be put down to the difference in the inductances of

these inputs. As this is dependent on the square root of the temperature, we would expect this to be further reduced in the actual experiment by a factor of $\sqrt{4.2/0.5} = 2.9$, where 4.2 K corresponds to the liquid helium temperature of our laboratory experiments, and 0.5 K to the operating temperature at ILL.

4.3.2 Repairing Solder Shields

Movement of the neutron guide will likely cause wear and tear of the shields. To test the impact of breakages in the shields on the magnetic pickup, the TWPs were fed through two shorter lengths of shield of 13.5 cm and 12 cm held in place with Kapton tape. In subsequent measurements, these were bound repeatedly with lead tape or soldered together with fresh solder to simulate a mended shield. The purpose of these measurements was to investigate whether the shielding factor could be sufficiently restored via repairs. Therefore no study was made into the effects of the size of the breakage, its position on the shield, or its performance in different applied fields. However, [132] suggests that the impact of the gap is greatest if it is nearer the centre of the shield, where end effects cannot moderate the breakage, as is the case in our investigations, and thus we might envisage these tests as a worst case scenario.

The effect of an annular defect for our shield geometry was examined using analytical models created in an open source, finite-element (FE) software, called Finite Element Method Magnetics (FEMM). As this is a 2D solver [2], the problem was defined axisymmetrically about the central axis of the solder cylinder. The solder is represented as a perfect superconductor cooled in zero-field conditions, that is with null magnetic vector potential at each of its edges, and then exposed to a uniform, transverse field. Gap sizes from 50 μm to 750 μm are examined, and the shielding factor is defined as the ratio of the integrated field across a 0.5 cm axial

[2] Another FE programme called OPERA will be introduced in Chapter 5 that can be used for 3D problems, but as we have only a magnetostatic licence, it cannot be used to model the vector behaviour of the superconductor.

length centred on the defect, with the broken shield and with no shield. This assumes that the relative field attenuation will be dominated by the effect of the gap rather than the intact areas of the shield. Figure 4.7 illustrates the variation of the shielding factor and the field at the centre point of the gap with gap width, and how these two variables are inversely related. Incidentally, that pattern observed when applying a uniform field in the longitudinal direction was found to be identical, as the field lines are similarly drawn into the defect.

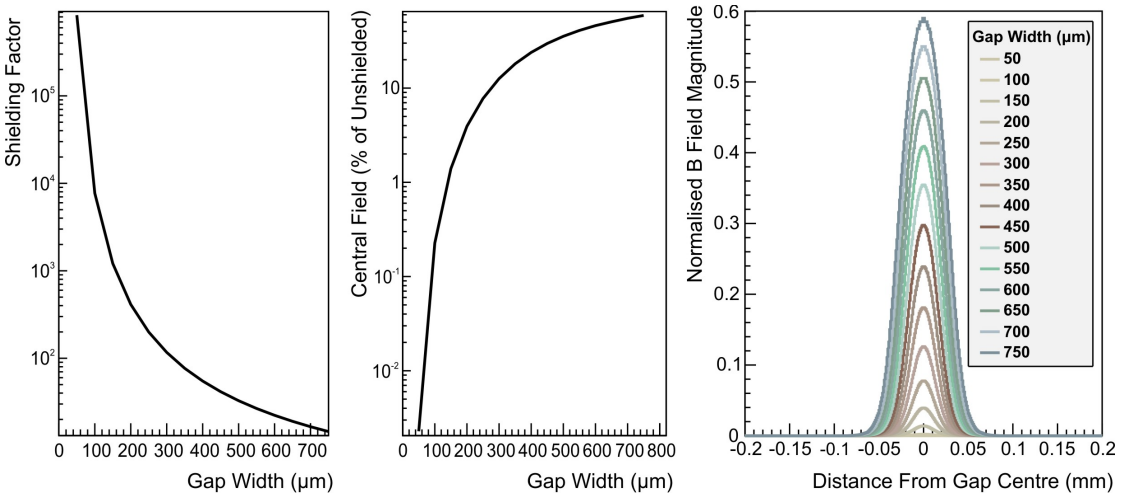

FIGURE 4.7: The variation of shielding factor, central field in the gap and field profile across the gap that the TWP would be exposed to for varying gap widths.

The average gap size of the broken shields that were experimentally tested was around 250 μm, and therefore one might expect a shielding factor of just above 100 from Figure 4.7. This is inconsistent with the numbers we obtain in Table 4.5, which fall an order of magnitude lower. This suggests that the model is lacking, as it assumes the breakage to be perfectly perpendicular to the solder axis, and that the applied field is spatially uniform across the gap when in reality, this will depend on where it is relative to the centre of the coils.

Whilst this discrepancy can be reconciled if we consider that we have modelled the perfect scenario, that associated with the repaired shields cannot. Clearly, the shielding factor drops when the tape is of width comparable to the gap, and is well-separated from the solder, but it was impossible to recreate a shielding factor

		Run 1		Run 2	
Shield	Status	SF_T	SF_L	SF_T	SF_L
6	Broken	12.5	24.2	6.24	9.24
	Repaired	855	13.0	405	34.3
	Twice repaired	2.32×10^5	2.50×10^5	7.72×10^3	8.45×10^3
7	Broken	10.8	19.1	17.1	7.56
	Repaired	366	23.5	—	—
	Twice repaired	5.64×10^5	2.30×10^5	2.85×10^3	465
8	Broken	4.67	21.0	6.68	3.61
	Repaired	445	49.7	84.7	3.47
	Twice repaired	662	193	1.42×10^4	4.03×10^3

TABLE 4.5: The maximum shielding factors for each broken shield through stages of repair.

as poor as that recorded for any repaired shield. Although it is difficult to model the tape, we can be sure that it is least longer than the gap width (the tape we use has a nominal width of 3 cm), is itself around 25-50 μm thick, and lies within a few hundred microns of the surface of the solder, separated only by a thin layer of adhesive. There are few parameters to vary, and therefore the difference between experiment and theory implies that the lead tape is of much poorer quality than expected, perhaps with impurities or holes which do not provide protection against the incoming field. On the other hand, these inadequacies might explain why the shielding factor is generally so much better for the shields which have been properly repaired. The pattern of field penetration observed in the simulations is analogous to the diffraction of waves through gaps and around corners (Figure 4.8). This means that, unless the isolated defects in the layers of tape match up or at the very least overlap, little field can infiltrate into the centre of the breakage.

At any rate, Table 4.5 shows that the shielding factor **can** be restored through the application of lead tape, although to varying degrees of success. These repairs were not quite as efficient as those from re-soldering the joints, with a single breakage mended in this way producing a sample with a shielding factor indistinguishable from the unbroken shield, with $SF_T = 3.84 \times 10^6$ and $SF_L = 2.36 \times 10^6$. However, this method is undesirable, as it would risk deforming the shield and causing heat damage to TWP encased within. The conclusion is therefore that for major

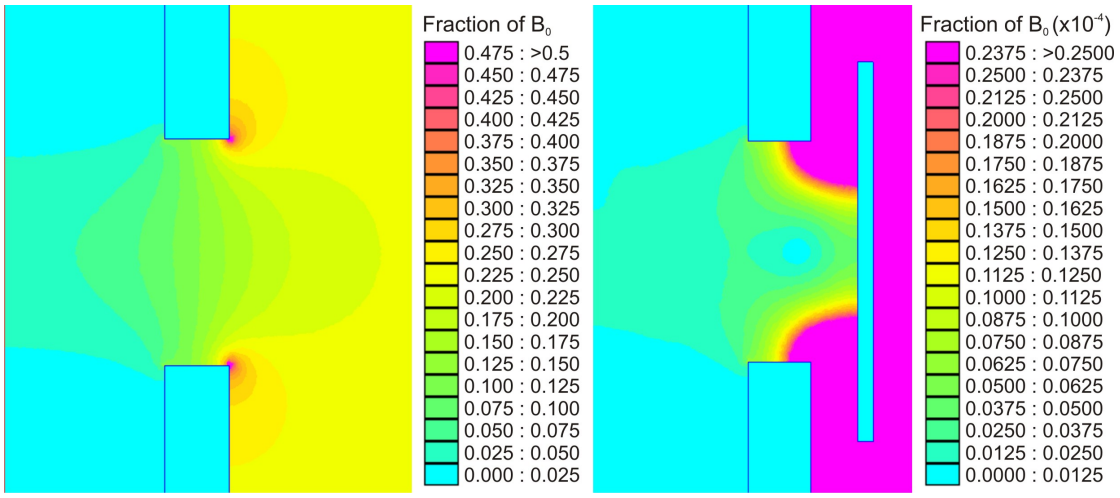

FIGURE 4.8: Field variation when a transverse field of magnitude B_0 is applied to solder with the defect, a gap of 300 μm, (left), and with the 'repaired' defect, with tape of 25 μm thickness in front of the gap (right).

repairs, with outright breakages in the capillary, re-soldering is a necessity to ensure adequate shielding. But for minor blemishes, the use of lead tape should suffice as a stop-gap to maintain and prevent and further degradation in the shielding.

4.4 Heaters

Heaters were built into the input lines as a means of nulling any supercurrent flowing in them, by taking the TWPs out of the superconducting state. This function is necessary given that the pickup loops are close to the Ramsey cells, such that a current through them could generate a magnetic field that could be seen by the neutrons. In particular, processes during the tuning of the SQUID could induce a significant current in the input circuit, and therefore the general protocol is to quench this current once the SQUID is operational. This section will first look at the power that must be supplied to the heater in order to achieve this functionality, and secondly, in line with previous investigations, the magnetic pickup at this point in the input circuit.

4.4.1 Heater Power

The aim of the following tests was to establish that the heaters, supplied with a maximum power of ≈250 mW (or 10 V from the control electronics) can indeed heat the input lines out of their superconducting state. This was previously found not to be the case, and therefore prior to these tests, the heaters were potted in Stycast to improve the thermal contacts between the TWPs and the resistor. The power required could be determined by monitoring the SQUID channels connected to the heater as a known magnetic wave form is applied to the exposed end of the TWP – this could be achieved with the same Helmholtz coil setup used in the shielding factor tests. At the point when the input lines become normal-conducting, the SQUID channels will no longer track the field: the power supplied to the heater is gradually ramped up until this point is reached (Figure 4.9). The time taken for the TWP to react is minimal: on the order of seconds for an abrupt change in power supplied to the heater. These powers vary between TWPs on a given heater (Table 4.6), possibly based on the extent of the thermal coupling between the line and the heater, but all are well in range of 250 mW.

It should be noted that a signal sent to the heater will itself cause a magnetic field that will be seen by the SQUIDs. This is not so much a problem if the heater is powered above the threshold for reversion to the normal-conducting state, such that additional flux will not be trapped in the input circuit. The greater problem is the possibility of extraneous magnetic pickup on the heater lines whilst the SQUIDs are operational. A possible future solution is to use optical fibres as a means of supplying power to the heater to avoid an extra current source near the SQUID lines, which could couple in unwanted flux.

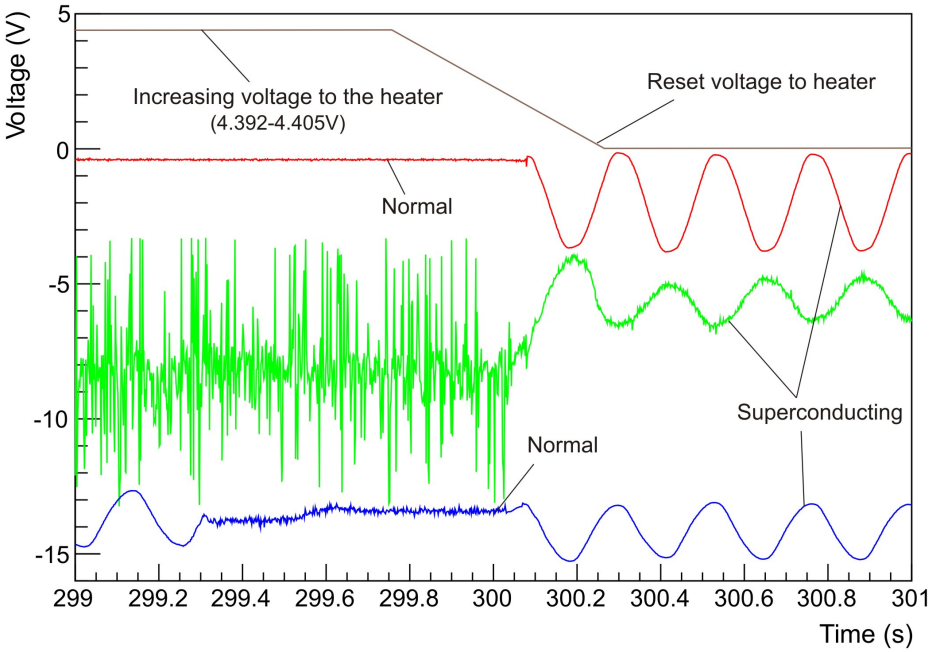

FIGURE 4.9: The response of three SQUID channels (red, blue green) as their associated heater (brown) is supplied with a 6 V ramp over a 300 s period, tantamount to a voltage resolution of 0.02 Vs^{-1}. The blue transition point is identifiable at ≈299.3 s, where it ceases to track the sinusoidal magnetic field. The green signal never becomes fully normal, but becomes worse at following the field, suggesting part of the shielding may go normal first. The reversion to the normal-conducting state is almost instantaneous as the power delivered to the heater is reset.

	Power (mW)		
	Channel 1	Channel 2	Channel 3
Heater 1	34.22±0.78	36.61±0.57	45.74±0.64
Heater 2	26.03±0.48	25.81±0.48	30.44±0.52
Heater 3	30.00±0.73	44.97±0.89	29.94±0.73

TABLE 4.6: The powers needed to revert the heater TWPs back from the superconducting to the normal state.

4.4.2 Magnetic Pickup

The heater resistor is introduced at a join between triple and single-cored sections of solder shielding, as was illustrated in Figure 4.1, and encapsulated in a lead cylinder to maintain the superconductivity of the shielding. The entire assembly was tested to examine its shielding efficiency, with each of the three heater channels wired to a SQUID input. The TWP ends which emerge the other side of the heater were terminated with a PTFE connector block, as described in Chapter 3. This was encased in lead tape and kept at a minimum of 8 cm from the bottom of the

coil assembly to prevent extraneous magnetic pickup. The 5.2 mm diameter lead cylinder which houses the resistor is mounted in the centre of the coils, since this is the component of interest. It is impossible to quantify the exact field the heater lines are exposed to and thus to determine a shielding factor: but as a means of comparing the heater setup, relative pickup values are defined with respect to the central field between the given pair of coils. Table 4.7 documents these values – at the lower end, the pickups may be consistent with those observed for the shielded TWPs. However, the numbers in general point to additional pickup and moreover, the spread in results suggests these are heavily contingent on the positioning of the heater and potentially its cooldown history. Moving the heater away from the coils did in general result in smaller signals for the same applied field, but there is no way of estimating how much field the heater configuration is subject to to standardise the pickup values, nor definitively pinpoint the source of the pickup – this is a major limitation of the setup. Having said this, whilst comparison between runs may therefore be complicated, measurements within a run, where the heater remains fixed in position, can be interpreted with greater confidence.

		Heater 1 TWP		
Test		A	B	C
1	P_T	3.6×10^{-8}	8.2×10^{-9}	2.2×10^{-7}
	P_L	3.6×10^{-7}	2.5×10^{-7}	2.0×10^{-6}
2	P_T	—	1.2×10^{-8}	2.3×10^{-7}
	P_L	—	1.3×10^{-7}	2.0×10^{-6}
3	P_T	8.6×10^{-6}	$9.\times10^{-6}$	3.0×10^{-6}
	P_L	1.5×10^{-4}	1.6×10^{-4}	6.3×10^{-5}
4	P_T	1.1×10^{-7}	4.7×10^{-7}	4.5×10^{-7}
	P_L	—	4.6×10^{-6}	7.8×10^{-6}

Heater	TWP	P_T	P_L
2	A	9.3×10^{-8}	4.0×10^{-7}
	B	1.6×10^{-6}	1.7×10^{-6}
	C	2.4×10^{-8}	4.7×10^{-8}
3	A	1.1×10^{-5}	6.4×10^{-5}
	B	2.0×10^{-7}	3.4×10^{-6}
	C	3.9×10^{-6}	4.3×10^{-5}

TABLE 4.7: Transverse (P_T) and longitudinal (P_L) pickup values in VnT^{-1} for the three heater setups tested, for each of their three connections (TWPs A,B,C). The results imply that there is extra pickup above the level of a shielded TWP, and are suggestive that this is strongly dependent on the exact location and orientation of the heater component.

Figure 4.10 illustrates how the magnitude of the signal does not scale proportionally with the applied field. The response is characterised by a region where

it increases with an approximate power law, followed by a section exhibiting an almost linear dependence on the field. This behaviour could be as a result of leakage of field through joins or some breakdown in the superconducting properties at points on the assembly, which may be explained by taking a closer look at the properties of a type-II superconductor. A layer of surface current is induced as an external field is applied, which acts in opposition to preserve the inner superconducting material in the Meissner state. In a larger field, further current-carrying shells are added such that the flux-front separating the two states, the penetrated and Meissner phases, migrates deeper into the superconductor. Eventually in a high enough field, the entire material carries the critical current density, J_c, the field reaches the interior and its level increases linearly with the applied field. There are many critical state models which describe the progression of field penetration in hard type-II superconductors, i.e. the relationship between J_c and the fields involved [133].

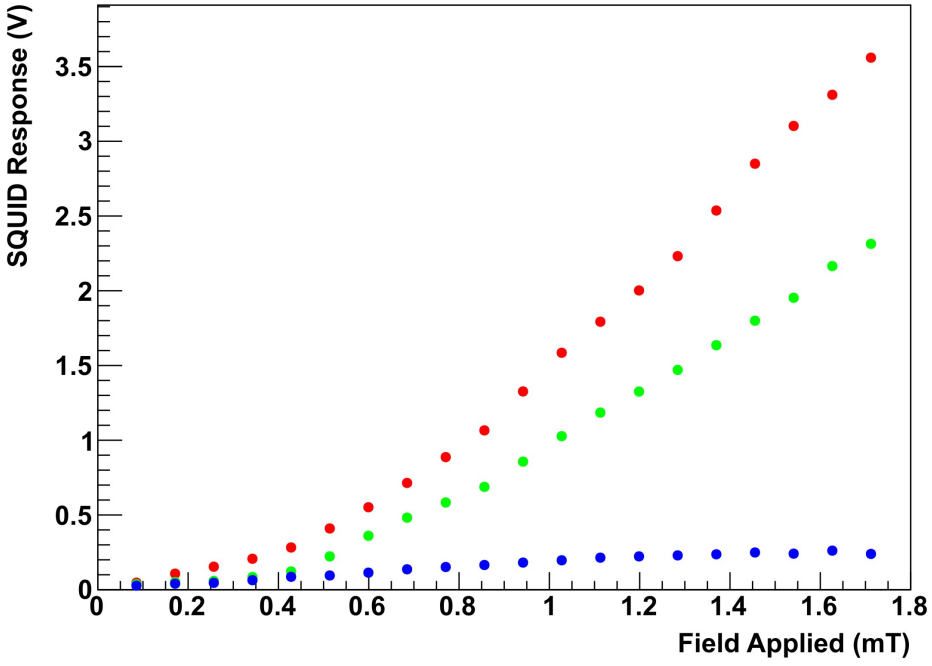

FIGURE 4.10: The measured signal sizes with field generated by the Y axis coils. Between 0-0.8 mT, the magnitudes ramp until they vary linearly > 0.8 mT. This could be indicative of the loss of perfect superconducting behaviour. The SQUIDs used are 0020 (red), 0237 (green) and 0562 (blue).

This behaviour is most likely to onset where the local field is highest: in this case,

a position of potential weakness is where the single-cored shields enter the lead cylinder. An OPERA model was made to investigate the field distribution at this junction, modelling the superconductor as a medium of negligible permeability [3]. They show that, in the cases of transverse field, the field in the vicinity of the shields as they enter the heater can build up to around 3 mT, or ~ 1.5 times the nominal external value (Figure 4.11). The field concentration may be greater if the shields are in fact closer together, and as the lead heater will become superconducting before the solder, any flux exclusion from this element may enhance the field at the shield positions. If there were a thin layer or a hole in the solder covering the join between shield and heater, the field accumulation could be enough to cause some of the solder to become saturated and begin to let in field.

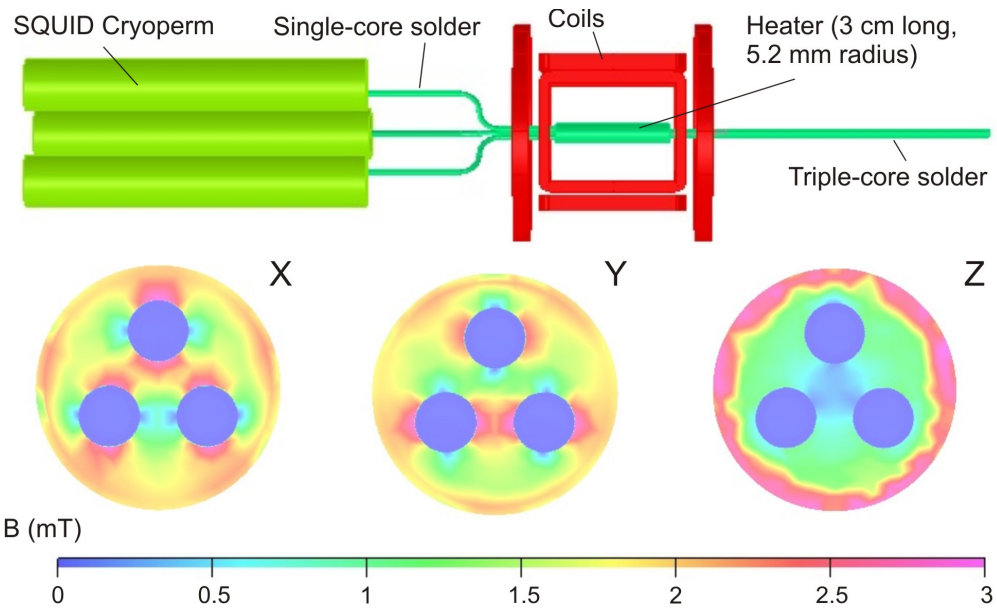

FIGURE 4.11: The full heater setup modelled in OPERA. The straight sections of the single-cored shields and the heater cylinder were analysed, giving the three plots. They are X-Y cross-sections of the field distribution at the join with the X,Y, and Z coils powered maximally respectively.

[3] As we have noted previously, the OPERA module we have available is not a suitable choice for modelling superconductors, since it is a magnetostatic package. However, by setting $\mu \approx 0$, the result should approximately simulate a superconductor in finite-field cooling conditions. Of course, we would rather it were zero-field cooled, since the field is applied after cooling in the experiment. However, since we are only interested in the fields outside the assembly, the field distribution on the periphery of the solder should remain a good approximation.

In conclusion, the values in Table 4.7 indicate that there is extra pickup due to the inclusion of the heater, which at its highest, would mean that the shielding factor is in danger of falling short of the 2.1×10^4 requirement that was derived at the beginning of this chapter. If the hypothesis that has been detailed in this section is correct, it would imply that the heater configuration would benefit from some changes to minimise any weak points where the superconducting medium could be breached. Rather than a complete redesign, this could be simply the addition of extra overlapping layers of superconductor around the joins, much in the same vein as in the discussions of the application of lead tape in Section 4.3.2.

Chapter 5

The Effects of Cryoperm: SQUID Sensitivity vs Spin Depolarisation

The inclusion of a Cryoperm shield for each SQUID, as we shall see in this chapter, bears consequences not only on the noise level of the SQUID, but also on the field environment in its vicinity. Previous chapters have made mention of both of these effects: in Chapter 3, Cryoperm was introduced as a means of mitigating the field seen by the SQUID as it is cooled (the cooling field), whereas in doing so, as remarked in Chapter 4, due to their proximity to the neutron guide, the shields could also affect the field seen by the neutrons. As the field a SQUID is cooled in is linked to a source of noise that it could register, a large magnitude could significantly undermine its resolution; as the ability for a neutron to maintain its polarisation is related to the homogeneity of the field it samples, even a small field distortion could have an adverse effect.

This chapter will assess the impact Cryoperm has on both factors in two parts: the first will introduce the concept of $^1/_f$ noise and investigate the extent to which Cryoperm can mediate its causes; the second will, through simulation, determine how detrimental Cryoperm is to the neutron polarisation, and explore ways to

lessen its effects. The ultimate aim is to suggest possible resolutions to the trade-off in the context of CryoEDM.

5.1 The Reduction of Low Frequency, $^1/_f$ Noise

The SQUIDs were supplied optionally with Cryoperm as a means of suppressing the field during cooling, and attenuating large changes in field whilst they are operational. The main risk of running the SQUIDs without this shielding in a magnetically variable environment is that they will become sensitive to extra sources of noise that could limit their performance. In particular, we will see that this affects the amount of low frequency, $^1/_f$ noise that is coupled into the SQUID.

$^1/_f$ or pink noise is widely found in physical systems, not least in electronic devices. In many cases, this flicker noise is attributable to known condensed-matter phenomena and can be limited with careful choice of fabrication techniques and operating conditions. For our long-term measurements, the presence of such noise could be problematic as it will not only limit the low frequency resolution of the SQUIDs, but could manifest itself as drift in their signals. This section gives a review of potential $^1/_f$ noise sources, and investigates whether the primary cause of this noise on our SQUIDs is indeed linked to the cooling field. In doing so, it addresses whether Cryoperm shielding can help maintain a lower noise level necessary for operations at ILL.

5.1.1 The Origins of $^1/_f$ Noise in SQUIDs

Koch *et al* were the first to explore the underlying causes of $^1/_f$ noise in SQUIDs [134]: they found critical current noise present in the Josephson junctions through the trapping and release of electrons in the tunnel barriers, a noise source that is

removed nowadays by suitable bias reversal schemes, and a second flux noise component that could not be explained at the time. The discovery of high-temperature superconductors prompted more detailed studies into the vortex state, which uncovered the processes of vortex motion and hopping. These lead to a noise contribution that can be severely reduced in the manufacture of the SQUID washer, most notably in the addition of pinning centres and the reduction of the linewidth, which inhibits the formation of vortices. Despite these improvements, which more or less eliminate the noise due to vortices in low-T_c SQUIDs operated in low temperatures and low fields, there emerged the mystery of why $^1/_f$ noise was still evident in the millikelvin range. Attempts to understand this extant noise have since seen a resurgence, driven by the development of superconducting integrated circuits for quantum information processing. One leading hypothesis ascribes this noise to the hopping of electrons from one defect site to another [135]. Each trap has an associated preferred spin direction in which the electron is locked, and thus not only the difference in position but the change in the orientation of the electron's magnetic moment results in a different flux coupled to the SQUID loop. Yet another theory posits that the noise arises from the random reversal of spins at the interface of the SQUID washer and its substrate [136]. In spite of noted correlations between temperature and geometry [137], none of these explanations have as yet been explicitly verified.

Whilst research at the cutting edge of the field focusses on the SQUID as a single entity, many practical applications use SQUIDs in conjunction with other devices or circuitry, which may couple in further sources of noise. In the case of the ILL SQUIDs, the SQUID itself will have been optimised for low noise performance by the manufacturer, but the input configuration is user-defined and not pre-characterised. Therefore, even if we find we are insensitive to the intrinsic SQUID washer noise, it is likely that some other form of noise would be picked up by the flux transformer. A major candidate for such extraneous noise will be vortex hopping in the superconducting elements of the SQUID setup, such as the niobium

can (see Figure 3.5) or the input circuit itself.

Movement of a single vortex from one pinning site to another, much like many of the phenomena mentioned above, can be envisaged as a two-level system (TLS), giving rise to random telegraph noise through transitions between the two states [138]. Using the Wiener-Khinchin theorem [139], it can be demonstrated that its spectrum takes the form of a Lorentzian:

$$S(\omega)_{\mathrm{TLS}} \propto \frac{\tau}{1+\omega^2\tau^2}, \tag{5.1}$$

combining the lifetimes of the two states, τ_1 and τ_2, as an effective rate for the process, $1/\tau = 1/\tau_1 + 1/\tau_2$. Since the fundamental process is thermally activated, the lifetimes of the two states may be expressed as:

$$\tau_{1,2} = \tau_0\, e^{E_{1,2}/kT}, \tag{5.2}$$

where E_n is the activation energy, or the barrier height of the given quantum well, and τ_0 is the characteristic attempt time of the process. If we assume that the energy distribution of the states is uniform, it can be shown that a set of such TLSs can generate a $1/f$ spectrum. Thus with the density of states given by $g(\epsilon) = \frac{1}{E_2 - E_1}$, with E_2 and E_1 now being cut-off energies, expressible from 5.2 in the form of $E = \ln(\tau/\tau_0)\, kT$, the statistical weight $g(\tau)$ can be written as:

$$g(\tau) = \frac{d\epsilon}{d\tau} g(\epsilon) = \frac{\tau^{-1}}{\ln(\tau_2/\tau_1)}. \tag{5.3}$$

Therefore the total power density of the TLSs is given by:

$$S(\omega)_{\mathrm{Total}} = \int_{\tau_1}^{\tau_2} g(\tau) S(\omega)_{\mathrm{TLS}}\, d\tau \tag{5.4}$$

$$\propto \frac{1}{\ln(\tau_2/\tau_1)} \int_{\tau_1}^{\tau_2} \frac{1}{1+\omega^2\tau^2}\, d\tau \tag{5.5}$$

$$\Rightarrow S(f)_{\mathrm{Total}} \propto \frac{1}{f}, \tag{5.6}$$

for f in the range $(2\pi\tau_2)^{-1} \ll f \ll (2\pi\tau_1)^{-1}$: a summation of the contributions of many TLSs therefore yields a $^1/_f$ noise source.

Since the density of vortices is commensurate with the number of TLSs, we should find that the $^1/_f$ amplitude increases with the cooling field. If it is indeed the dominant source of low frequency noise, not only should this correlation with field be apparent, but we might expect to see a significant difference in the noise properties when the SQUIDs are run with and without Cryoperm.

5.1.2 Preparatory Investigations in Different Cooling Fields

The following subsection details experiments undertaken to characterise the noise properties of our SQUIDs in varying magnetic conditions. The SQUIDs were run with the niobium can and, unless otherwise stated, with shorted inputs. This emulates the exposed wire that would be present if a solder-shielded TWP were attached to the inputs, and allows us to see the effects of a flux transformer, whilst being relatively unresponsive to changing fields and external interference. To ensure the SQUIDs were subject to a static cooling field, a solenoid was mounted around them on the dipstick, rather than affixed within the dewar/cryostat.

Initials tests were conducted directly inside a helium dewar, which placed tight dimensional constraints on the assembly. A solenoid, with winding density of 30000 turns per metre, was wound around an open-ended cylindrical lead shield closely wedged around the SQUIDs. Only three SQUIDs were used, since this was the occupancy of the available baseplate that could comfortably fit through the dewar mouth. The results of these preliminary tests are not fully recorded here for brevity, but they revealed a number of issues:

- The system was particularly sensitive to oscillations, specifically any small movement of the dipstick. The lead and solenoid were not firmly affixed, and so differences in the contractions of the materials at cryogenic temperatures could

create space for them to move. Motion of the lead relative to the SQUIDs, due to its trapped flux, would result in a change in the field seen by the SQUIDs.

- Higher frequency noise components were apparent. These seemed to have been due to the Taconis effect [140], and on blocking the open end of the dipstick, were found to disappear.

- The optimal bias current would change with the applied cooling field, generally, but not always, decreasing with increasing field. This phenomenon is evident in all three sets of runs and is discussed in more detail in the succeeding analysis. As long as the SQUIDs are retuned in each run, it was hoped that this would not cause significant problems in interpreting the results.

- Reheating the SQUIDs does not significantly impact on the noise spectrum. As the SQUIDs would cool in the same field frozen in by the niobium cap, this suggests that either retrapping of flux occurs at the same sites on the washer, coupling in more or less the same flux noise, or that vortex hopping in the SQUID chip circuitry is overshadowed by stronger noise sources.

- Noise levels when the field was present throughout the measurement differed little from those when the field was removed after cooling. Field removed estimates are less accurate as, although the SQUIDs retain a remnant magnetisation from exposure to the field, this decays slowly over time. They are commonly used, however, in circumstances where it is difficult to maintain a stable field for the duration of the measurement. In subsequent tests without a lead shield and with higher solenoid currents, field removed measurements were preferred.

The data suggested a marked difference between the SQUIDs run with Cryoperm and those without. However, the limited statistics and the high variability of the data imply that the tests are better served in more regulated conditions. This therefore signalled the transition to cryostat-based investigations, where the wider mouth of the cryostat would let six SQUIDs be tested at once.

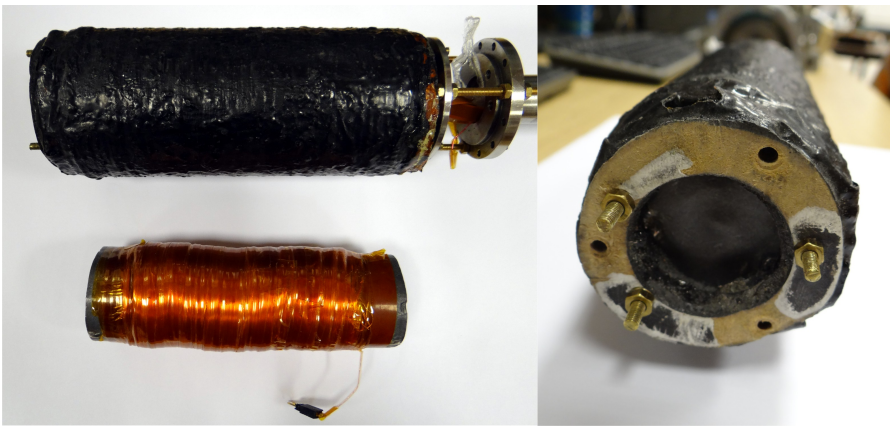

FIGURE 5.1: *Left*: the two solenoids for the Cryoperm investigations, the one below being that for the preliminary tests. *Right:* The new solenoid can generate larger fields, can enclose more SQUIDs and has improved mechanical stability, with a closed end allowing it to be properly affixed to the dipstick.

In the next iterations of tests, whose results are analysed in this chapter, a new solenoid was constructed with a winding density of 42000 turns per metre (Figure 5.1), such that when energised with a 1 mA current, it generates a field comparable to the Earth's field ($\sim 50\,\mu$T). This was again set up to apply a field parallel to the SQUID chips, having been coated in Stycast to provide mechanical stability and protect the enamel coating of the wire, and with a closed end with holes to allow it to be affixed to the brass rods of the dipstick, thus providing some field stability.

One set of tests focussed on higher currents, and therefore fields, in order to accentuate any difference the Cryoperm could introduce. Field removed measurements were performed as the increased heat load from the solenoid makes the helium boil-off rate more critical. For another set of measurements at lower fields more typical of ILL, the field was retained throughout each cycle. A new lead shield was made and firmly fastened between the solenoid and SQUIDs. This was to protect against any field changes caused by current drift in the solenoid, such that they do not mask true sources of pink noise at smaller fields. To judge the relative levels of noise, a noise spectrum is generated for each measurement. The findings and analysis process are now discussed.

5.1.3 Analysing the Noise Spectra

The amplitude spectral density (ASD) of the voltage signals read out from the SQUIDs provides information about the noise power distribution. Much like in Chapter 4, a Fast Fourier transform (FFT) algorithm is used to estimate the ASD. The intrinsic ASD of the SQUID should have two major components - one which is invariant with the frequency (the so-called white noise) and one which is roughly inversely proportional. The latter pink noise, though generally referred to as $1/f$, can have an exponent that can take non-integer values when expressed in the form $f^{-\alpha}$. If the ASD is taken to be the square root of the power spectral density, the $1/f$ noise would be expected to manifest as a line with slope of around -0.5 in log coordinates. This is potentially the product of multiple $1/f$ sources, but it is later shown that it may be possible to isolate the component related to the magnetic field by monitoring the change in the exponent as the field is varied.

In order to properly characterise the low frequency range, it was found that at least half an hour of data-taking was desirable. The signals were sampled every $100\,\mu$s and the data sectioned into events, each with 16384 points. According to the Nyquist theorem, this lets us reconstruct frequency components up to 5 kHz without aliasing. However, having a higher sampling rate is not necessarily beneficial when data-taking for a long duration, since the more measurement points, the larger the file size and the greater the processing power required to run the FFT algorithm. For most of the ASD plots illustrated, approximately 1024 events are stitched together to provide around 1024×16384 FFT samples. This corresponds to a bin resolution of around 6×10^{-4} Hz via the formula f_s/N, where f_s is the sampling frequency, and N is the number of FFT points. If our sole goal is to look at the low frequency range of the spectrum, then the same frequency resolution can be achieved with a lower sampling rate and fewer data points. However, since it is also useful to characterise any variation in the white noise level, the chosen

parameters serve as a compromise, allowing accurate spectral analysis in a bandwidth spanning 6 to 7 decades between approximately 10^{-3} Hz to 10^{3} Hz, whilst not overwhelming the FFT procedure.

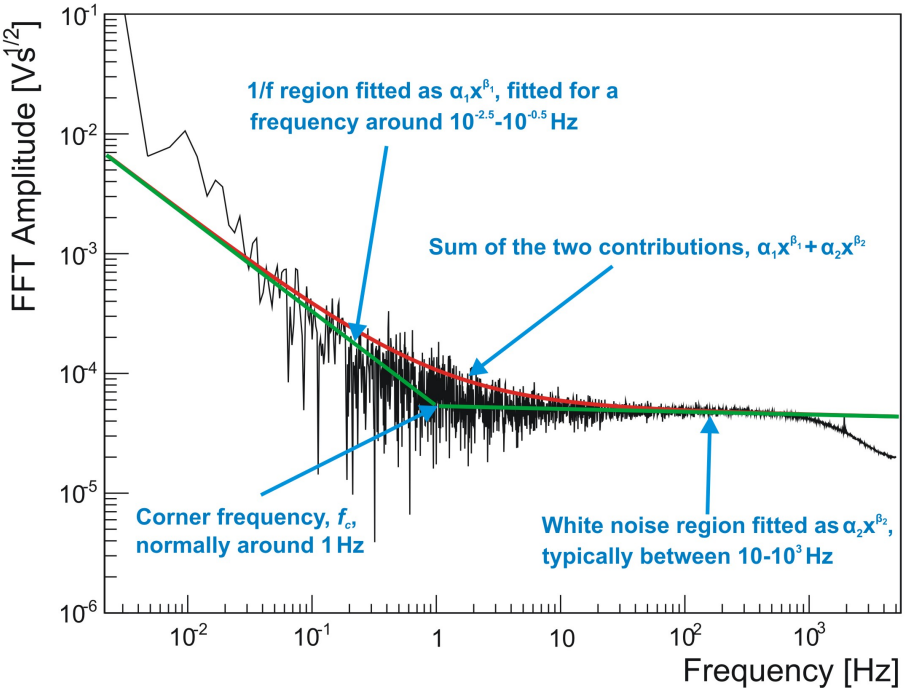

FIGURE 5.2: Example fit of a noise spectrum for a given SQUID channel.

The proposed fit to the spectrum is parameterised as $\alpha_1 x^{\beta_1} + \alpha_2 x^{\beta_2}$, with the first term representing the $1/f$ component, and the second, the white noise. As the ROOT fitting algorithms cannot distinguish between two terms with identical expressions, the spectrum was split into two parts to be separately fitted, assuming each to be dominated by a single colour of noise. Instead of fitting αx^{β}, which would require an iterative, non-linear approach, the data were linearised such that a linear regression fit can be performed, as per the visualisation in Figure 5.2. Although the ASD is smoothed and windowed, the bin resolution clearly impacts on its quality in the lower frequencies. A robust fitting algorithm is therefore adopted, which disregards the outliers in the data. Here, we invoke ROOT's built-in TLinearFitter class, whose regression method is an implementation of the Fast Least Trimmed Squares (FLTS) algorithm developed by Rousseuw and Van Driessen [141]. The fitting procedure remains a least squares minimisation, but

the process is only applied to a subset of data, chosen as a result of a series of 'concentrations' which iteratively select the data with the lowest residuals.

It is worth contemplating what observations would be anticipated if there is a source of $1/f$ noise linked with the magnitude of the magnetic field. First, the amplitude of the total $1/f$ noise found from the fit, α_1, should increase with the applied field, if there is an additional noise component caused by the field. Not so intuitive is how the exponent, β_1, should change. To see how the power might vary with the field, consider the following. We have a sum of noise contributions to which we fit such that $\sum_i a_i x^{b_i} \approx Ax^B$, and on adding a further noise component, $\sum_i a_i x^{b_i} + Cx^D \approx Gx^H$. The question we wish to answer is how the new exponent, H, is related to the old parameter, B. As the fit was calculated in logarithmic coordinates, it makes sense to rewrite the equations as

$$\ln\left(\sum_i a_i x^{b_i}\right) \approx \ln A + B\ln x$$

$$\text{and} \quad \ln\left(\sum_i a_i x^{b_i} + Cx^D\right) \approx \ln G + H\ln x.$$

Substituting in the approximation from the first expression, the LHS of the second relation can be rewritten as

$$\ln\left(\sum_i a_i x^{b_i}\right) + \ln\left(1 + \frac{Cx^D}{\sum_i a_i x^{b_i}}\right) \approx \ln A + B\ln x + \ln\left(1 + \frac{C}{A}x^{D-B}\right).$$

Then, differentiating both sides and multiplying through by x, we have

$$B + \frac{D-B}{\frac{A}{C}x^{B-D} + 1} \approx H.$$

We would expect that the amplitude of the extra noise component would be much smaller than the sum of all the original contributions, and that the exponent will not differ greatly from that of the previous fit. Choosing an x value within the range of the fit, $10^{-3} < x < 10^{-1}$, under the assumptions that $C \ll A$ and

$|B - D| < 1$, we can ignore the 1 in the denominator such that

$$B + \frac{C}{A}(D - B)x^{D-B} \approx H.$$

From this relation, we can see that, no matter which x is chosen in the required frequency range, if D is greater than the exponent estimated from the sum of all other noise contributions, B, the new estimate of the exponent, H, will be greater than B, and vice versa. This means that, if on increasing the applied field, the same noise source with exponent D is amplified, we would expect the exponent determined by the fit to monotonically increase/decrease depending on whether D is bigger/smaller than that found from all other contributions.

5.1.4 Results

The results of these investigations provide evidence that the Cryoperm shields reduce the level of $^1/_f$ noise and moreover, that the additional noise seen on the unshielded SQUIDs is magnetic in nature. In all figures, the field applied by the solenoid is quoted instead of the current that is supplied to it, despite this being the parameter we work with. This is for ease of interpretation in the context of a noise source related to the field, but it must be remembered that it refers to the field at the centre of the solenoid and thus seen by an unshielded, but not a shielded SQUID.

First, a direct comparison of SQUIDs with and without Cryoperm run in parallel shows a clear distinction between the two (Figure 5.3). Whilst the noise level of the shielded SQUIDs remains relatively insensitive to the applied field, that of the unshielded SQUIDs is higher and increases with the field magnitude. There is also a difference in the exponent – the separation implies that the extra source(s) of

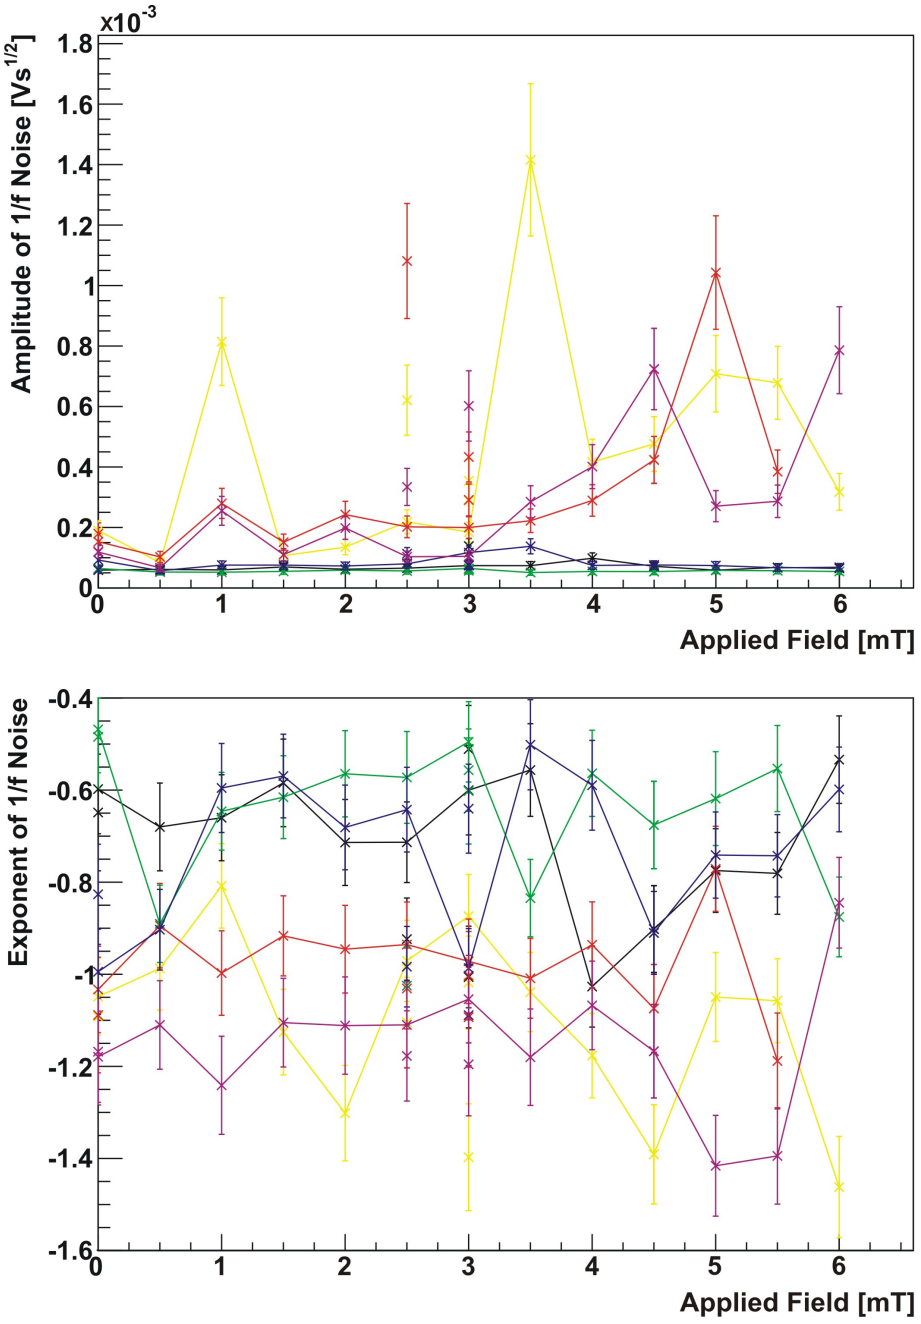

FIGURE 5.3: *Above:* Amplitude of the $1/f$ noise estimated from the fit versus the applied field for six SQUIDs run simultaneously: three shielded with Cryoperm (black, blue and green - SQUIDs 0411, 0857 and 0410 respectively), and three unshielded (yellow, red and pink - 0367, 0854, 0855); *Below:* the exponent of the $1/f$ noise with applied field for the same six SQUIDs. Both graphs show a difference between the shielded and unshielded SQUIDs.

noise has a more negative component(s) than the average constituent of the non-magnetic noise, although a decrease in exponent with applied field is imperceptible for the unshielded SQUIDs due to the limited statistics.

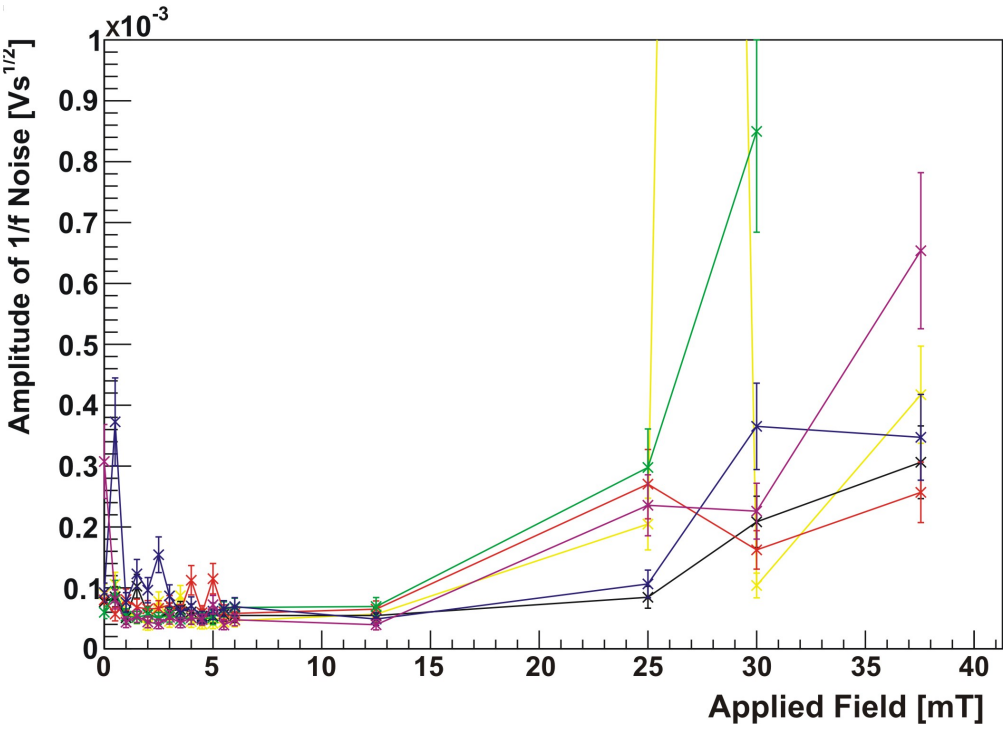

FIGURE 5.4: The amplitude of the $^1/_f$ noise for six shielded SQUIDs in increasing applied field. This demonstrates a breakdown in the shielding of the Cryoperm at the higher end, with an increased noise level due to the magnetically-related noise. The SQUIDs are the same as those used in Figure 5.3, represented by the same colours.

Second, the fields were pushed further in the case of solely shielded SQUIDs to see if there is the onset of the same theorised noise (Figure 5.4). Indeed the $^1/_f$ amplitude is seen to exceed the nominal value at around 15 mT, and generally increases thereafter. It can be shown that this is consistent with the unshielded results with an OPERA simulation. This suggests that the field at the position of a SQUID chip with Cryoperm is equivalent to that without when the applied field is ∼40 times greater. This is plausible on comparing the respective levels of noise.

Lastly, an unshielded SQUID operated with no connections on its input demonstrates a much lower noise floor in comparison, which does not vary with the cooling field. This is an indication that there are noise sources introduced by the flux transformer, some of which are field dependent.

We can infer from all three observations that there is a noise contribution linked

with the magnetic field, and that the Cryoperm helps in reducing its impact. However, it is apparent that there are irregularities in the data due to the inherent variability of the cooling process. On occasion, it was suspected that the liquid helium level was low and a SQUID would then sit above the fluid, as with the repeated points at 2.5 mT on Figure 5.3. This raises its white noise level, as it is fundamentally linked to temperature (see Section 5.1.5), but also the $^1/_f$ if a responsible mechanism were thermally activated.

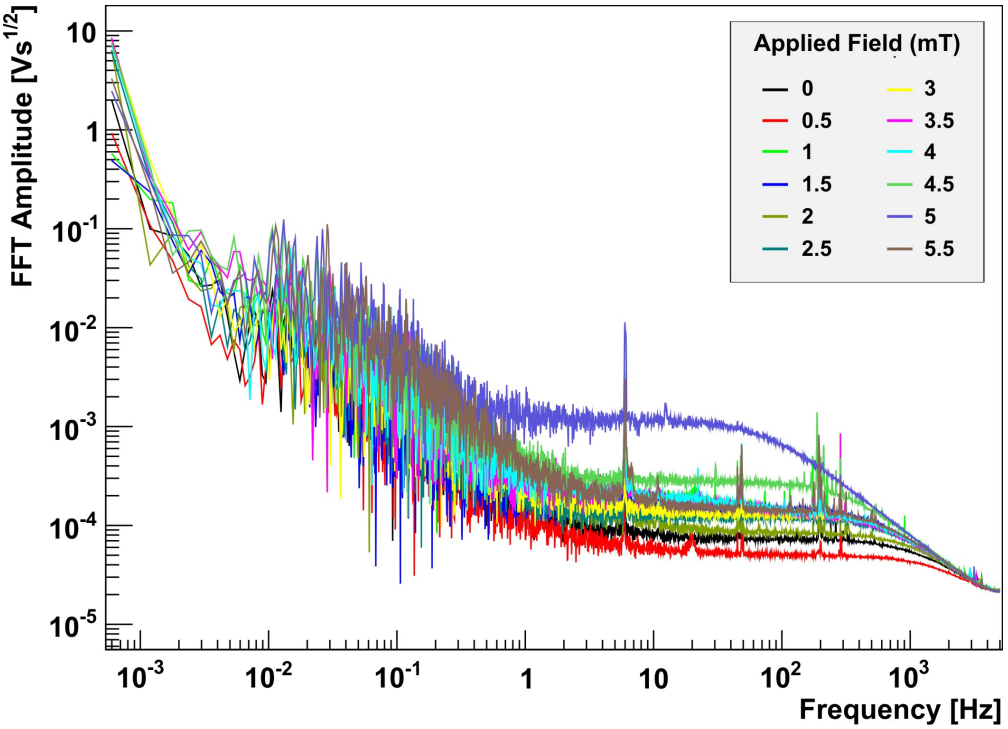

FIGURE 5.5: The noise spectra for the same, unshielded SQUID (0367) in different cooling fields, showing an unmistakeable difference in the white noise levels, which is not seen for shielded SQUIDs until the highest of fields. It is not invariably the case that the greater the field, the higher the noise floor; an increase is often accompanied by a change in bias current, which suggests a modification of the intrinsic behaviour of the SQUID in different field environments.

It was also noticed that the bias current for this and other SQUIDs would vary from one cooldown to another, which appears to manifest in changing white noise levels in the noise spectra, as illustrated in Figure 5.5. Rapid cooling of the apparatus can create internal field gradients and asymmetries that can alter the bias current. Further, vortices may be prone to form on the electrodes of the Josephson Junctions, as they represent leading edges perpendicular to the field. These

may give rise to local magnetic field variations, which disturb the coherence of the phase across the junction and thus suppress the critical current. The screening parameter, $\beta_L = {}^{2LI_0}/_{\phi_0}$, is driven away from its ideal value of unity, which could result in hysteretic behaviour and a higher intrinsic noise of the SQUID. This could culminate in a higher white noise level than from a rise in temperature alone.

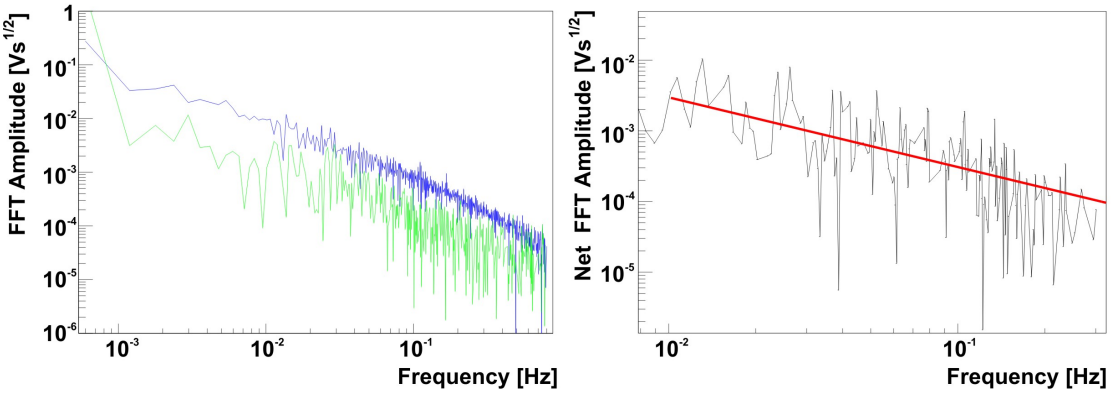

FIGURE 5.6: *Left:* The $^1/_f$ components for an unshielded SQUID at the Earth's field (green) and at 120 mA (blue), after the according white noise level has been subtracted from each spectrum. *Right:* The difference between the two gives the supposed noise source arising from the increased magnetic field, which can be fitted to.

The variability of the results may also be explained if there are no well-defined pinning sites, such that vortices can by chance form more inwardly in the input circuitry, and couple in more flux. Due to budget limitations and the time constraints imposed by the finite lifetime of the helium, the focus was on getting a wide range of measurements and repetitions were only possible for a select number. All of these factors may prevent us from drawing hard quantitative conclusions, but do not undermine the general outcome of this research, which is that there is benefit in using Cryoperm shields.

As a final examination, we re-analyse the spectra, subtracting the spectrum with no solenoid field for each SQUID, having shifted the former so that their white noise baselines match (Figure 5.6). In theory, this should leave us with just the $^1/_f$ noise related to the magnetic field, since this is the only explicit difference

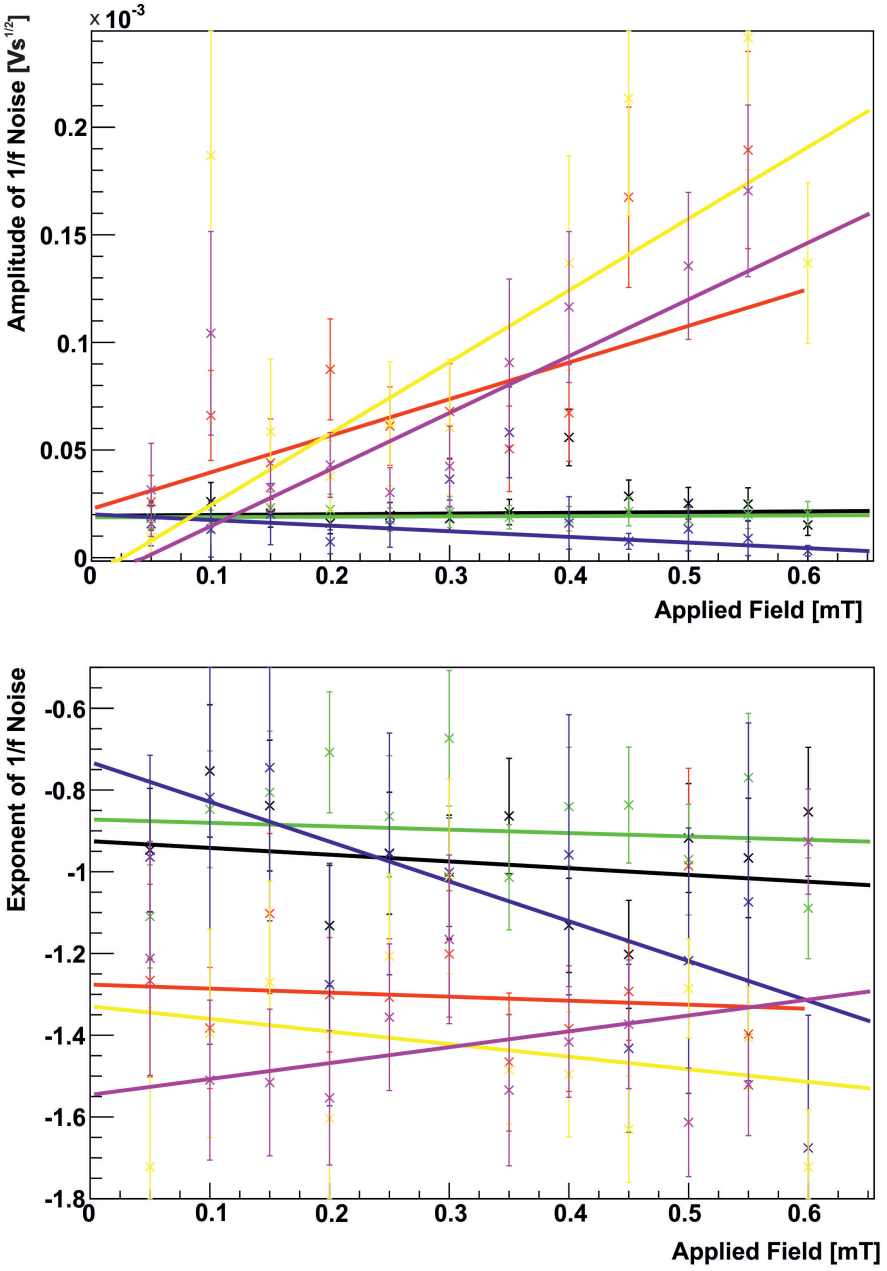

FIGURE 5.7: The amplitude (above) and the exponent (below) of the net $1/f$ noise resulting from subtracting the baseline spectrum taken at the Earth's field for each of the SQUIDs from the corresponding spectra at each measured current, where black, blue and green are shielded SQUIDs; red, yellow and pink are unshielded.

between the two: i.e. returning to the prescription in Section 5.1.3, this eliminates the common contributions $\sum_i a_i x^{b_i}$ and leaves behind Cx^D. Fitting in log coordinates, the amplitude should show a linear dependence with field and the exponent should remain constant. The variation between cooldowns will of course

affect these measurements since, for example, we have remarked that the processes which impact on the white noise level could also influence the $1/f$ noise. However, the underlying trends should be apparent with sufficient data, and will provide clarification.

Best fit lines have been used for the data presented in Figure 5.7 to better identify the underlying trends. It is evident that the $1/f$ amplitude remains more or less static for the shielded SQUIDs whereas it is seen to increase for the unshielded; the exponent is not quite so clear cut, with those of the unshielded SQUIDs being in general lower than that of the shielded. These results, particularly for the amplitude, provide further clarification that there is a difference between SQUIDs run with and without Cryoperm, and that the latter are more susceptible to a noise whose effect grows with the magnitude of the external field. In addition, it is also apparent, from the spread of the data points and their errors, that the unshielded SQUIDs are more sensitive to the cooling conditions and changes in their environment. This ultimately makes their performance less predictable, and is another aspect to bear in mind if considering running the ILL SQUIDs without Cryoperm.

5.1.5 Consequences for CryoEDM

Investigations have shown a SQUID with Cryoperm to have superior noise properties than one without and, in particular, be more robust to potential noise deriving from the cooling field. It remains to be seen whether this difference is significant in the context of operations for CryoEDM.

The magnitude of the field in the vicinity of the SQUIDs will be around 300 μT, as justified in Section 5.2.2, on the lower end of the fields examined. The difference between the shielded and unshielded configurations is then not as pronounced but may still prove critical. Figure 5.8 shows the fitted power spectral densities of the same SQUID run in three different configurations at this field level, and Table 5.1

illustrates the noise voltages. These are found by integrating these curves for the relevant bandwidth and square rooting the result, having accounted for the white noise contribution in the $1/300$ to 1 Hz range. The $1/f$ contribution is then approximately six times higher for the unshielded SQUID than for the shielded SQUID. Assuming the difference in the white noise floor is again attributable to the variability of the cooling conditions, we find that the total noise voltage is ≈1.5 times greater for the SQUID without Cryoperm.

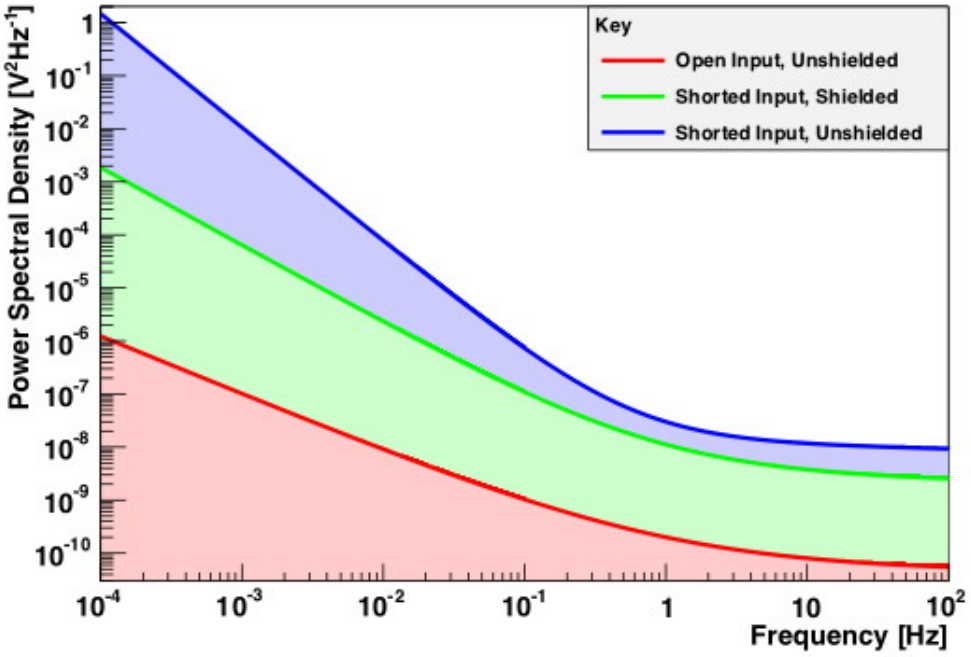

FIGURE 5.8: Noise power comparison for the same SQUID (0367) operated in three different configurations, having been cooled in a 300 μT field.

Thus, in this instance, the resolution is 1.5 times poorer for an unshielded SQUID. However, the same cannot be said for a SQUID in CryoEDM since the operating conditions will differ: namely, there is a reduction in temperature from 4.2 to 0.5 K, and an increase in input inductance from 350 to 831 nH, using the values for the 1.5 m TWP and a 20 mm diameter pickup loop calculated in Chapters 4 and 3 respectively. So long as the $1/f$ component changes by a similar amount to the white noise, then this gap in resolution will not widen.

SQUID Configuration	$1/f$ Noise Voltage/mV ($1/300$ to 1 Hz)	White Noise Voltage/mV (1 to 10^3 Hz)
Unshielded, Open Input	0.018	0.225
Shielded, Shorted Input	0.256	1.516
Unshielded, Shorted Input	1.472	2.923

TABLE 5.1: Noise voltages of each colour of noise in the given frequency ranges for the three SQUID configurations illustrated in Figure 5.8, $1/300$ Hz being the sampling frequency of the Ramsey resonance measurement.

The input inductance should affect the mutual inductance between the SQUID and feedback loop, M_{fb}, such that the voltage we measure goes as

$$V_{fb} \propto \frac{1}{M_{fb}} \propto \left(1 - \frac{320\,\text{nH}}{L_i}\right)^{-1},$$

since $\Phi_{SQ} = M_{fb} I_{fb}$. To verify this, the results for the shorted input are cross-checked with those for the open input. With nothing attached to its input terminals, the latter has an infinite effective input inductance; in the case of the shorted SQUID, L_i is essentially the inductance of the input coil, 350 nH from Table 3.2. When transformed from an open to a shorted input, the results would scale from 0.018 to 0.210 mV for the $1/f$, and 0.225 to 2.625 mV for the white noise – these are in greater agreement with the shielded SQUID, which implies that the extra noise for the unshielded SQUID is introduced by the flux transformer.

With $\Phi_{SQ} = M_i I_i$, and assuming M_i to remain unchanged, the signal due to this source of noise can be written as

$$V_{fb} \propto \left(1 - \frac{320\,\text{nH}}{L_i}\right)^{-1} \frac{\Phi_i}{L_i}. \tag{5.7}$$

The difficulty is in expressing Φ_i: pursuing the idea that this is due to the thermally-activated hopping of vortices,

$$\Phi_i \propto e^{-E/kT},$$

where E is the activation energy of a pinning site – [142] finds this to be on the

order of a few meV for niobium thin films, which would mean that a transition from 4.2 to 0.5 K would effectively eliminate this source of noise and therefore suggests the use of Cryoperm would not be beneficial. However, there are a number of factors which influence the activation energy – the size of the potential barrier itself can depend on the ambient temperature and magnetic field. Moreover, the pinning energy will be affected by the quality of the material, and thus could be different, for instance, for the niobium cap. If the primary source of noise is vortices on the flux transformer, a different material entirely, we must also consider that its added length will increase the number of vortices, and that much of it cannot be screened by the Cryoperm. Lastly, at some crossover temperature, the movement of vortices via quantum tunnelling will become the dominant mechanism [143] [144], which might mean that the change in temperature does not reduce the noise as dramatically. There is simply not enough information about the nature of the source(s) to accurately predict the noise voltage expected in this scenario. But can more be said of the performance with and without Cryoperm?

We return to look at the white, Johnson noise, whose power spectral density in terms of voltage variance is

$$V^2 = 4k_B T R \Delta f, \tag{5.8}$$

where k_B is the Boltzmann constant, T the temperature, Δf is the relevant bandwidth and R the resistance of the circuit in question. Then,

$$V_{fb} \propto \left(1 - \frac{320\,\text{nH}}{L_i}\right)^{-1} \sqrt{T} \tag{5.9}$$

leads to a noise voltage in CryoEDM of 5% that of our measurements. For the reduction in our hypothesised source of $^1/_f$ noise to be comparable, the activation energy would need to be two orders of magnitude smaller than estimated in [142], which seems implausible. Therefore this noise source will account for a much smaller proportion of the total noise at 0.5 K than at 4.2 K, and the difference between the shielded and unshielded SQUIDs will be much less than a factor of

1.5. Having said this, it is dangerous to suggest that the Cryoperm is inessential on the basis of this theory alone, particularly if its temperature dependence has not been determined.

5.2 The Depolarisation of Neutrons in Transport

Loss of polarisation results in a lower visibility of the resonance curve, which in turn diminishes the statistical sensitivity of the experiment. Polarisation during neutron transport is maintained by ensuring a gradual decrease in field along the $\approx$ 6 m of neutron guide, from 1 mT at the source entrance to 5 μT in the Ramsey cells, such that the neutron spins may follow it adiabatically. This transition is managed by a set of eight holding field coils, but any superconducting or ferromagnetic materials in close proximity to the neutron guide will cause local departures in the field, which have the potential to enhance any depolarisation. This is a particular concern with the SQUIDs if they are operated with Cryoperm and mounted where proposed in the six-way section. This section of this chapter aims to assess the field disturbance caused by the Cryoperm and determine the extent of possible depolarisation that could result.

5.2.1 Quantifying the Spin Depolarisation

The major criterion that must be met to preserve the relative direction of the neutron spin is that the magnetic field which it samples changes at a rate slower than the frequency with which its magnetic moment precesses (the Larmor frequency, $\omega_L = \gamma B$). If there are no magnetic sources in the guide tube, and the fields and neutron motion can be taken to be primarily along its central axis, the change in field as seen by a neutron may be expressed as $\omega_B = \frac{1}{B}|\frac{\delta B}{\delta t}| = \frac{1}{B}|\frac{\delta B_z}{\delta z}|v_n$, where B is the magnitude of the field and v_n is the speed of the neutron. This adiabatic

condition is then expressible as a parameter

$$k = \frac{\gamma B^2}{v_n |^{dB}/_{dz}|} : \tag{5.10}$$

if k is much greater than unity, the neutron spin should be able to track the changes in field, and no significant depolarisation should occur. This is a useful concept for describing a single point or in a region with a uniform field gradient, but not for one with a spatially variable magnetic environment.

Here we adopt a numerical approach, and introduce the longitudinal and transverse relaxation times, T_1 and T_2 respectively, which characterise the decay of a neutron's polarisation to its thermodynamic equilibrium value. T_1 can essentially be thought of as the decay constant for the projection of the neutron's spin vector in the initial polarisation direction: that is,

$$P_z(t) = P_z(0) e^{-t/T_1},$$

if the magnetic field is taken to be aligned in the z-direction, and $P_z = \vec{P} \cdot \hat{z}$ is the projection of the polarisation vector in the magnetic field direction. T_2 may be similarly expressed, substituting P_z by the projection of the polarisation vector in the x-y plane, P_{xy}. These relaxation times are dependent on factors other than magnetic field inhomogeneities, such as neutron-surface interactions, T_1 being more affected by certain mechanisms than T_2 and vice versa. The focus here, however, is on depolarisation due to variations in magnetic field along the guide tube and therefore T_1 is the parameter that is examined, as ideally the field and the neutron spin vectors are initially aligned in the same direction along the guide tube axis. Therefore, to establish whether or not the Cryoperm has a significant impact on the polarisation, the T_1 times for scenarios with and without the SQUIDs can be compared. To achieve this, both the magnetic environment and the neutrons' response to it must be simulated.

The latter is done by a Monte Carlo (MC) type program, developed by Prof P. G. Harris, which can simulate the movement of the neutrons and behaviour of their spins when supplied with a defined field map. The neutrons are taken to be moving ballistically, and start with spins aligned with the local magnetic field, their evolution governed by the equation:

$$\vec{\sigma}(t) = \left(\vec{\sigma_0} - \frac{(\vec{\sigma_0}\cdot\vec{B})\vec{B}}{B^2}\right)\cos(\omega t) + \frac{\vec{\sigma_0}\times\vec{B}}{B}\sin(\omega t) + \frac{(\vec{\sigma_0}\cdot\vec{B})\vec{B}}{B^2},$$

with $\omega = \gamma B$ the Larmor precession frequency of a neutron and σ_0 its starting spin vector. This is the solution to the Bloch equation of motion, $\dot{\vec{\sigma}} = \gamma\vec{\sigma}\times\vec{B}$, subject to a constant, uniform magnetic field. The simulation tracks the net polarisation of the neutrons, α, with time, where α is maximal initially when the spins are aligned with the local field. T_1 is defined by $\alpha(t) = \alpha(0)\mathrm{e}^{-t/T_1}$ and is therefore found by fitting to the curve of $\ln(\alpha)$ against time. The process is iterated multiple times to obtain an average for T_1, which we use to compare the different configurations.

The simulation takes as inputs the probability of specular reflection from the walls and the maximum UCN velocity, which defines the initial Maxwellian velocity distribution – for CryoEDM, these were measured as 0.84 for the beryllium guide tube and $6.95\,\mathrm{ms}^{-1}$ respectively. The program does not account for any temporal variation in the magnetic fields, and assumes that other relaxation mechanisms, such as spin-spin interactions and wall relaxation are negligible. This is suitable for our purpose, since this affords a straightforward comparison between the scenarios with and without Cryoperm that is not masked by non-magnetic effects, and is solely reliant on changing field with position and not time. It also models the volume as a neutron bottle in which the neutrons move randomly, rather than predominantly in the downstream direction. An 80 cm section of guide tube centred on the SQUIDs was chosen with this in mind, such that it is long enough that the likelihood a neutron will spend a disproportionate amount of time in the Cryoperm region is small, but conversely that is short enough that its effects are

still measurable. Nevertheless, it must be realised that the T_1 times generated by the MC are not meant to be taken at face value, but rather serve to provide a relative scale for comparison between different magnetic environments.

5.2.2 Modelling the Magnetic Environment

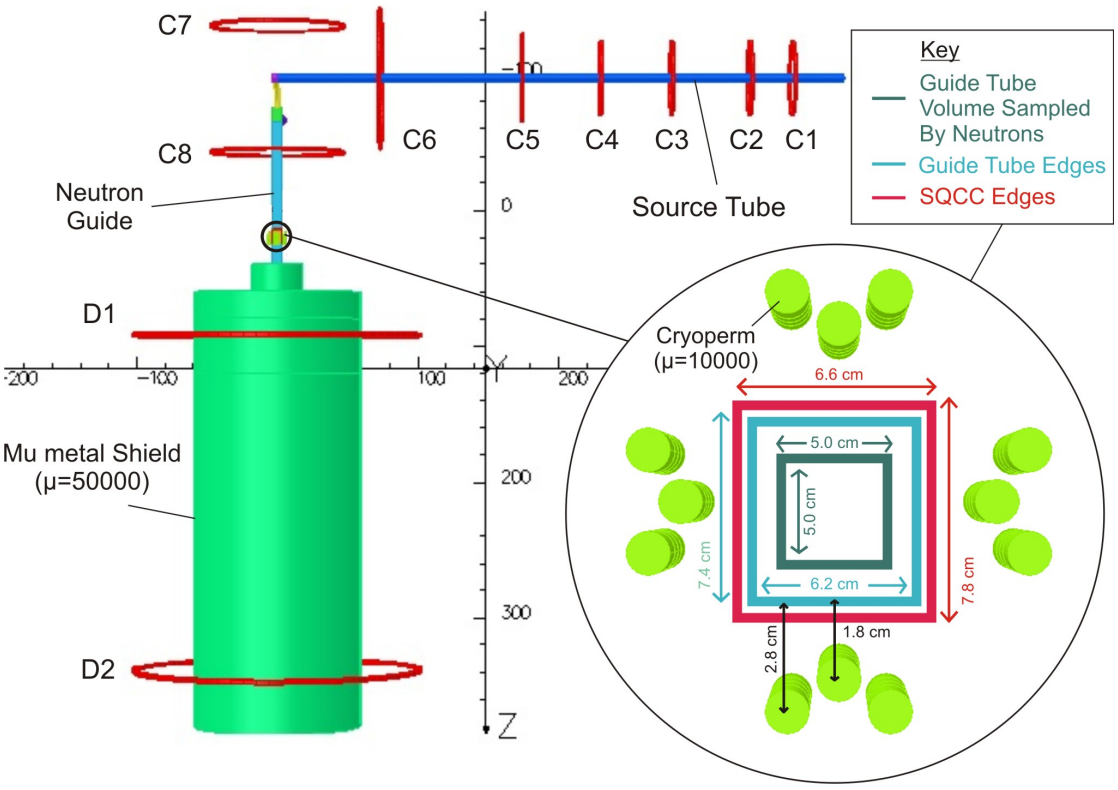

FIGURE 5.9: The full model, as viewed in the −Y direction. The model includes both the source section and transfer region for completeness, with the eight holding field coils, C1-8, and the activation compensation coils, D1-2. The detail shows a cross-section of the guide tube at the six-way section, with the 12 SQUIDs and SQUID Compensation Coil (SQCC).

OPERA, a finite element analysis package, was used to construct a magnetostatic model of the part of the experiment of interest. This builds on Katsika's model [145], and uses the ambient field measurements taken by her at ILL: $B_x = 9.4\,\mu$T, $B_y = -32.2\,\mu$T and $B_z = -22.8\,\mu$T. The design includes the solenoid and Solenoid

End Correction Coil (SECC), that are hidden under the mu metal layers in Figure 5.9, and the SQUID Compensation Coil (SQCC) which is discussed in Section 5.2.4. Most importantly, it features the holding field coils, C1-8, which are used to maintain a smooth change in field down to the 5 μT in the Ramsey cells. These coils are wired as two sets of four such that a single current is common to C1-4, and another to C5-8. The currents are always directed upstream so as not conflict with the Earth's field, especially around the transfer region where there are two 90° bends.

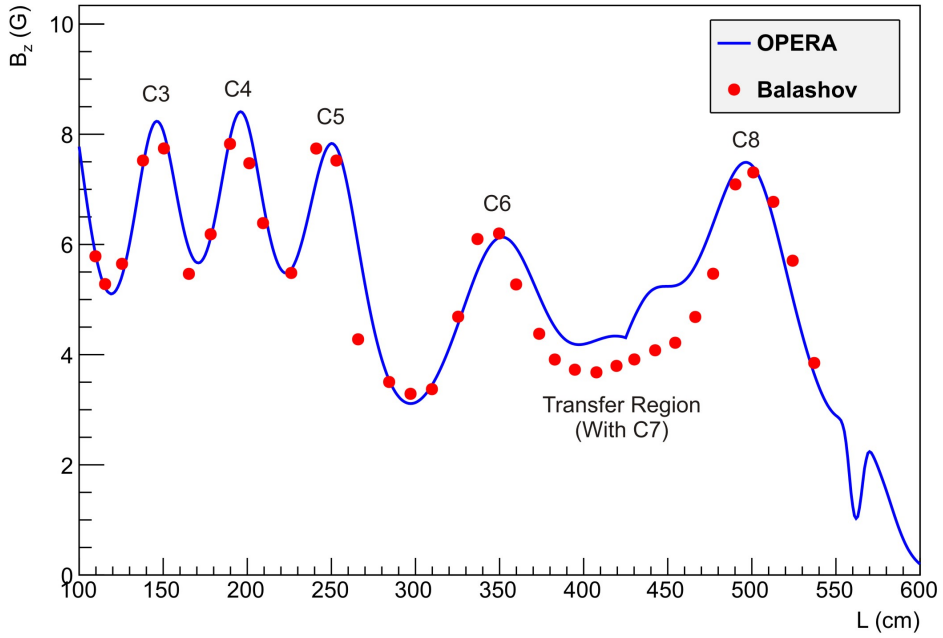

FIGURE 5.10: Comparison of the OPERA generated data and Balashov's measurements taken onsite [146], for the field directed down the guide tube. This enables us to calculate the CURD values to be used in the simulation, corresponding to a C1-4 current of 6.2 A, and a C5-8 current of 20 A. There is an overgeneration of the field in the transfer region, perhaps due to some misalignment in the ILL measurements or misjudgement of orientation when extracting values from OPERA.

As the coils are modelled as single conductors of various thickness the conversion between experimental and simulated current is not straightforward. OPERA uses Current Density (CURD) as its parameter of choice, that is the current in the coil divided by its cross-sectional area. To determine the conversion factors between current and CURD, the CURDs were adjusted so that the fields produced in the

simulation matched those from a past run at ILL (Figure 5.10). For the most part, these currents will be kept the same in our simulations as they only weakly affect the Cryoperm region – they are listed in Table 5.2. On the other hand, the C8 coil has a significant impact, and so its current is varied between 25 A, the maximum that could be applied, and 10 A, an approximate lower limit if the field is to vary smoothly. The optimum currents for the solenoid and SECC were pre-determined by studies carried out by Pendlebury and Grozier as 10 mA and 0.974 A respectively [147], and these must be energised in our simulations, else the field levels fall close to zero and the neutrons depolarise on sampling the upstream end of the bottle. Instead of iterating the procedure for every configuration of

	Element(s)							
	C1-4	C5	C6	C7-8	D1-2	SQCC	SECC	Solenoid
Current (A)	6.2	*10-25*			0	*0-0.02*	0.01	0.974

TABLE 5.2: The currents chosen for the coils in our model. The values in italics are those which are varied in this analysis, with the SQCC being a feature that is explored in Section 5.2.4. A DC offset in the active compensation coils, D1-2, negligibly affects the Cryoperm region, and so they are not energised.

currents, the process can be streamlined by analysing each element separately and superimposing the results to produce the overall field distribution. As each coil is fixed in position, its field contribution for a given current can be found by scaling the field at unit current accordingly. Thus the simulations need only be run once for each set of coils and the Earth's field, with the mu metal and with and without the Cryoperm.

5.2.3 The Effect of the Cryoperm

A comparison between the scenarios with and without Cryoperm shows it to have a detrimental effect on the T_1 time, with a factor of 5-10 decrease over the full range of currents. In the case of no SQUIDs, the T_1 peaks at a C8 current of around 4 A, outside our usable scope, then steadily decreases with increasing current (see

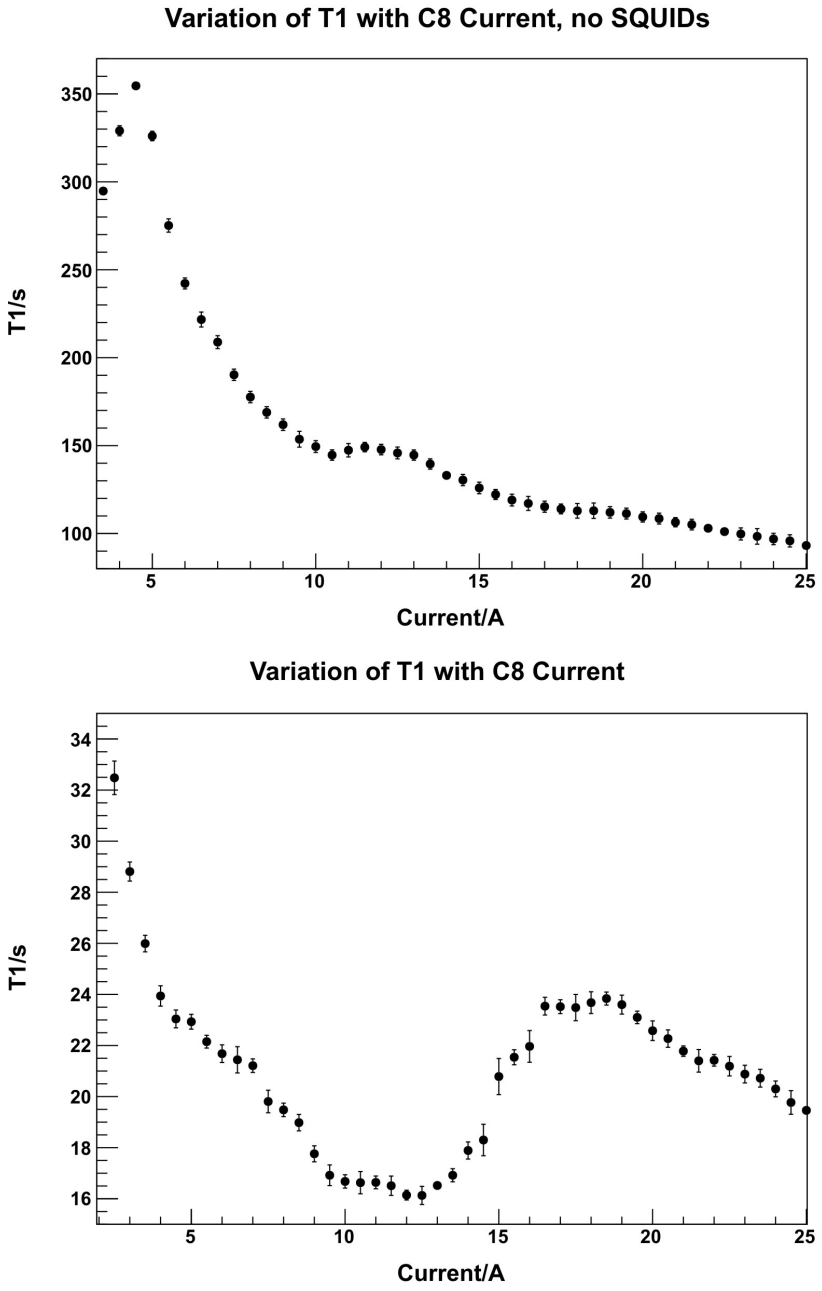

FIGURE 5.11: A comparison between the T_1 times with no Cryoperm (above) and with Cryoperm (below). With no other coils or SQUIDs present, a C8 current of approximately 4.5A is optimal. Introducing the SQUIDs has a considerable negative impact on the T_1 times in our current range, with values dropping by an order of magnitude from the scenario with no SQUIDs.

Figure 5.11). Including the SQUIDs introduces some significant structure, such that T_1 does not monotonically fall but is seen to recover slightly at the higher end of the current range. This behaviour can be reconciled by examining the field pattern and considering what happens to k. The Cryoperm causes a dip in

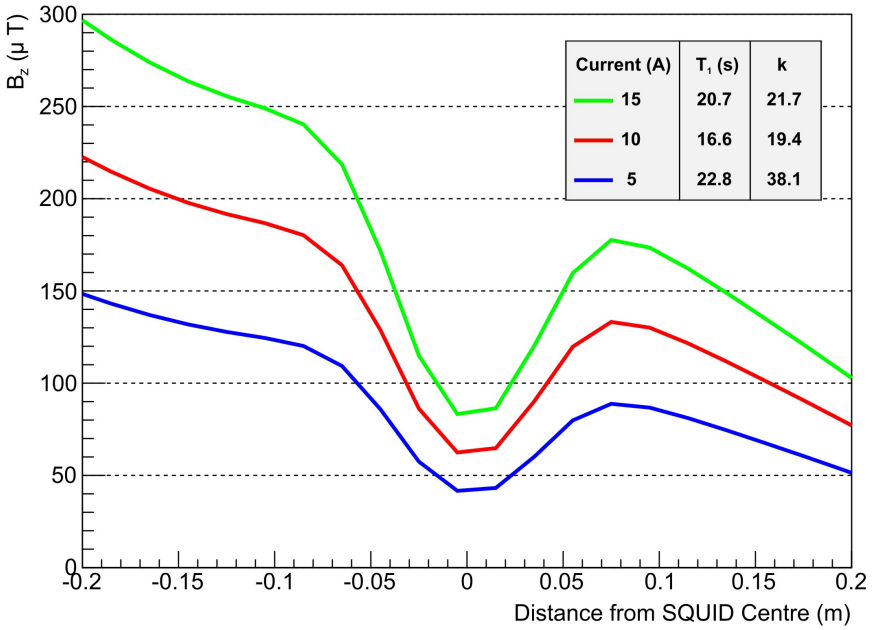

FIGURE 5.12: Comparing the field profiles, the T_1 times and the k values for C8 currents 15 A, 10 A and 5 A. There is no indication from the field profiles as to why the T_1 time should be lowest for 10 A, but the lower value of k suggests that we are around the minimum in the trade off between the field's gradient and its magnitude.

the field, whose relative size grows with the C8 current; however, the bottom of the dip is then at a higher field. The trend in k, and therefore in T_1, can be explained by the interplay between the field magnitude and field gradients within this depression, referring to Equation 5.10 (Figure 5.12).

In any case, T_1 drops to a nominal value of around 20-30 s in the presence of the Cryoperm, and whilst these are not accurate values, this is not acceptable given the filling time of the neutron cells is 60 s. As there is little other choice for the positions of the SQUIDs under the design of CryoEDM, the only choices for the experiment are to run the SQUIDs without Cryoperm or to mitigate its effects. Having discussed the possibility of the former previously in this chapter, we will now look at an implementation for the latter.

5.2.4 The SQUID Compensation Coil (SQCC)

5.2.4.1 Characterising the SQCC

The SQCC is a rectangular solenoid wound on the neutron guide with dimensions 7.8 cm x 6.6 cm x 10 cm, as integrated in the OPERA model. It was constructed as a temporary measure to counteract the supposed effect of the Cryoperm by generating a field to compensate the depression. It is centred 1.5 cm relative to the midpoint of the SQUIDs, due to the positioning of bolts on the neutron guide.

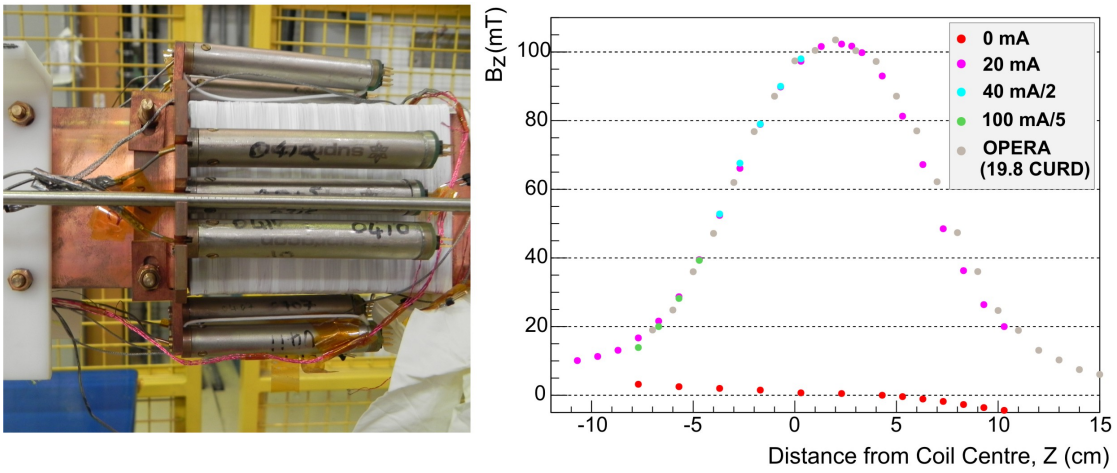

FIGURE 5.13: *Left:* A photo of the SQCC, a rectangular solenoid wound around the guide tube. *Right:* The SQCC was energised at a range of currents and the results normalised to check that they are consistent. The ambient field was also measured, to which the asymmetry in the results can be attributed. The CURD of the coil as modelled in OPERA was adjusted so that its field matched the measurements.

An axial fluxgate was used to measure the magnetic field generated along the central (or z) axis of the coil for various supplied currents (Figure 5.13). There is a slight asymmetry in the ambient field, thought not to be inherent in the guide tube but attributable to magnetic materials in the vicinity of the work area at the time these measurements were made. This was checked by measuring the field with no current, which indeed demonstrated a skewed residual field. Likewise, to check the field scales linearly with the current, a number of different currents were used (Figure 5.13), which were found to be consistent with one another. A current

to CURD conversion factor of 1.066 CURD = 1 mA could then be established for use in the OPERA model.

Results from studies without the SQCC, as were illustrated in Figure 5.14, do not advise on suitable currents for C8, and therefore both it, and the SQCC, are varied over their full ranges. The maximum current that can be applied to the SQCC when cold is 20 mA, limited by the internal resistance of the power supply. Figure 5.14 illustrates the T_1 times, having examined currents in the C8 and SQCC coils in intervals of 0.5 A and 1 mA respectively. It shows T_1 to be greatest (83.7 s) for a C8 current of 13 A and a SQCC current of 20 mA, and this is further determined to be a stable maximum, with deviations of ±0.25 A in C8 and ±0.5 mA in the SQCC, leading to a variation in the T_1 of at most ∼4%. The SQCC can therefore restore the T_1 time to two-thirds the value without the Cryoperm (Table 5.3), which is certainly beneficial, but begs the question of whether the design, and therefore the T_1 time, could be improved.

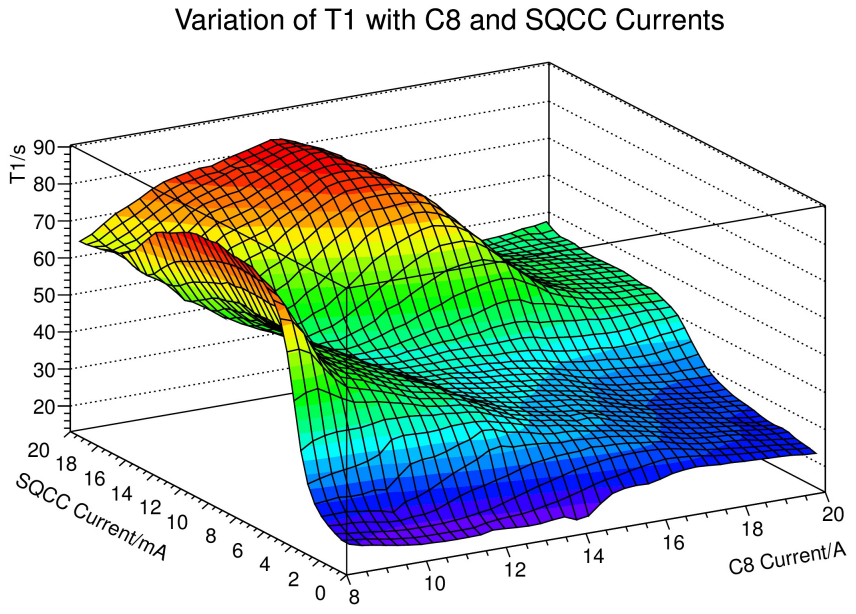

FIGURE 5.14: Illustration of the dependence of T_1 on C8 and SQCC currents. The plateau suggests that the maximum at C8=13 A, SQCC=20 mA is stable, and moreover, it remains that maximum even at higher SQCC currents up to 30 mA beyond our range of interest (not illustrated).

Without SQUIDs	With SQUIDs	With SQUIDs and SQCC
121.2 s	16.4 s	82.4 s

TABLE 5.3: T_1 times for the optimal current configuration, illustrating the compensating effect of the SQCC.

5.2.4.2 Improved Designs for the SQCC

1. An Analytical Fit for a Rectangular Solenoid

The present SQCC has been shown to be, in theory, beneficial, but it can only do so much: even accounting for it being off-centre, a single rectangular coil cannot produce a field that well matches the Cryoperm dip (Figure 5.15). If we are free to choose a new design, but are still limited to coils of rectangular profile that wrap around the guide tube, a symmetrical configuration about the centre of the SQUIDs with three (or more) coils should prove better. This should be able to compensate both the central depression and the brows either side. We assume that we are not otherwise restricted, in that we can offset the SQUIDs or file down bolts to allow continuous winding of the coils if need be.

With many more parameters to consider, i.e. the length, current and offset of each of the coils, instead of a scattergun approach in OPERA, an improved prescription to define them is to identify the desired field pattern, i.e. the inverse of the Cryoperm dip, and fit to it. This ignores any distortion of the SQCC field itself by the Cryoperm or mu metal, which will affect the quality of fit. Nevertheless, this should provide a useful starting point from which we can make small adjustments.

The approach requires an analytical formulation for the field from a rectangular solenoid at any point within it. In the axial direction, we derive this using the

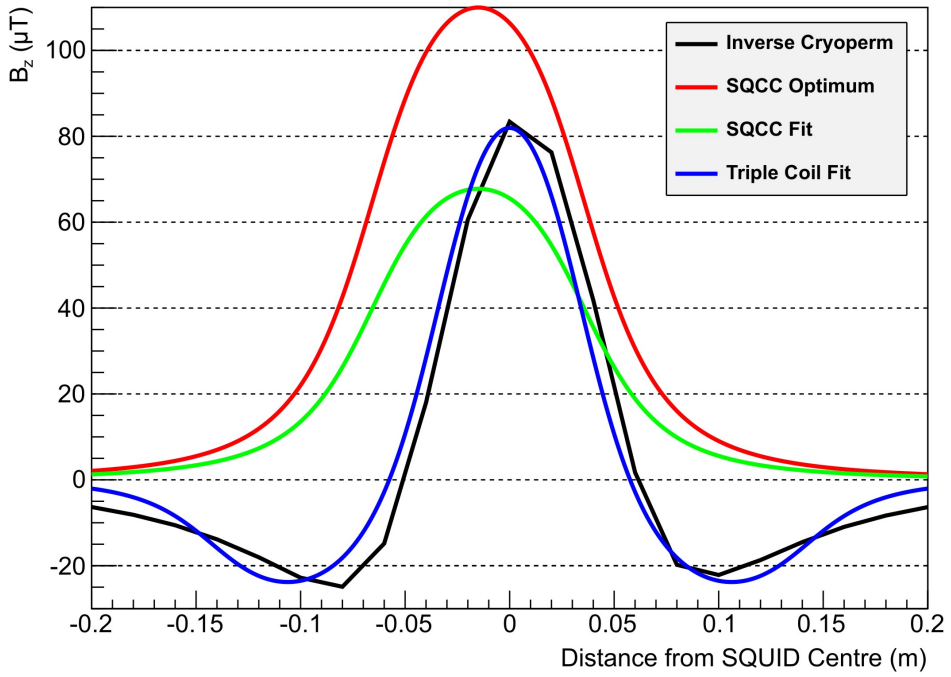

FIGURE 5.15: Various fits to the Cryoperm dip using different configurations of rectangular solenoid. The analytical fit for the SQCC is much lower than the OPERA optimum, which suggests that overcompensation is better. The fit can be much improved with an arbitrary coil geometry, here with a symmetrical assembly of three coils centred on the SQUIDs.

framework from Section 4.2.2 as

$$\frac{\mu_0 In}{4\pi}\left[\arctan\left(\frac{(W/2+y)(l/2-z)}{(L/2+x)\sqrt{(L/2+x)^2+(W/2+y)^2+(l/2-z)^2}}\right)-\right.$$
$$\arctan\left(\frac{-(W/2+y)(l/2+z)}{(L/2+x)\sqrt{(L/2+x)^2+(W/2+y)^2+(l/2+z)^2}}\right)+$$
$$\arctan\left(\frac{(W/2-y)(l/2-z)}{(L/2+x)\sqrt{(L/2+x)^2+(W/2+y)^2+(l/2-z)^2}}\right)-$$
$$\left.\arctan\left(\frac{-(W/2-y)(l/2+z)}{(L/2+x)\sqrt{(L/2+x)^2+(W/2+y)^2+(l/2+z)^2}}\right)\right]+\ldots$$

labelling the relative displacement of the point from the solenoid centre in the axial direction as z and those transverse as x and y, and with the variables as set out in Appendix A (L and W are the lengths of the sides of the coil, l is the length of the solenoid, I the current and n is the density of the windings (i.e. number

of coils/l)): the ellipsis signifies there are 16 terms in all, considering all other permutations of the variables (L/2+x), (L/2−x), (W/2+y), (W/2−y). Down the central axis of the coil, this simplifies to

$$\frac{\mu_0 In}{\pi}\left[\arctan\left(\frac{2L(l/2-z)}{W\sqrt{L^2+W^2+4(l/2-z)^2}}\right)-\arctan\left(\frac{-2L(l/2+z)}{W\sqrt{L^2+W^2+4(l/2+z)^2}}\right)\right.$$
$$\left.+\arctan\left(\frac{2W(l/2-z)}{L\sqrt{L^2+W^2+4(l/2-z)^2}}\right)-\arctan\left(\frac{-2W(l/2+z)}{L\sqrt{L^2+W^2+4(l/2+z)^2}}\right)\right].$$

This expression may prove more useful since the neutrons should spend longer near the centre of the guide tube, and therefore there might be benefit in placing greater emphasis on correcting the field in that region. Both formulae are used to fit three coil configurations to the inverse of the Cryoperm distortion, which consists of a middle coil and two flank coils either side. For the latter off-centre coils, we require the transformation $z \to z \pm a$ where a represents its displacement from the centre of the SQUIDs. The free parameters in the fit are then I, l for each of the coils, and a for the flank the coils, assuming the same winding density as the single coil of 56.8/cm.

	Central Coil		Flank Coils			
	Length (cm)	Current (mA)	Length (cm)	Current (mA)	SQUID Centre Separation (cm)	T_1 (s)
Version 1 (Full Profile)	6.4	21.8	10.8	5.5	8.6	56.6
Version 2 (Central Axis)	7.5	30.3	11.2	7.5	8.8	63.6

TABLE 5.4: Triple configurations resulting from fits to the inverse of the Cryoperm dip using the full field profile and just the field down the middle of the guide tube. It is assumed that the winding density is the same as for the single coil, that the two flank coils have the same specification because of the symmetry of the field distribution, and that these coils can overlap.

The parameters and T_1 times for the two fits are listed in Table 5.4, which indeed shows that fitting to the field profile down the central axis is the better option. But both T_1 times as they are fall short of that we have found for the present SQCC of 83.7 s. As was previously implied, the method does not deal with any departures from this ideal field due to the presence of the Cryoperm. Further, as

the fields from rectangular coils will never be exactly compensatory, whilst they may be closely matched to the Cryoperm down the centre, there will be gradients introduced in other parts of the guide tube cross-section. The result of the single coil SQCC illustrated in Figure 5.15 suggests that overcompensation is useful, for it could ensure that the field from the SQCC is dominant over the Cryoperm distortion such that the field profile is relatively smooth across the cross-section.

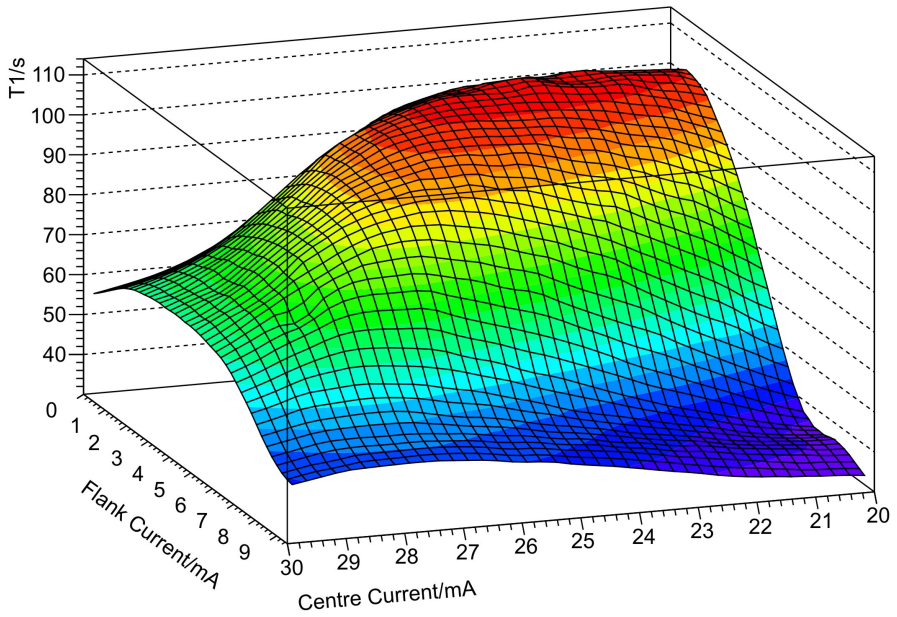

FIGURE 5.16: Illustration of the dependence of T_1 on currents of the flank and middle coils of the proposed triple coil SQCC. The 23 mA exceeds the current that can be supplied by the power supply, but this can be reconciled by increasing the winding density of this central coil.

So as insinuated before, we can take the better of the two configurations and refine it through an OPERA model to take account of these factors. Fixing the lengths and positions of these coils as found from the fit, the currents in them are varied, much in the same vein as with the present SQCC and C8 coils in Section 5.2.4. The T_1 times are illustrated in Figure 5.16, demonstrating a maximum of 106 s, around $^7/_8$ the time without Cryoperm, at a middle coil current of 23 mA and a flank coil current of 2 mA.

2. An Analytical Fit Using the Laplace Equation

The triple coil arrangement offers a significant improvement, but one could imagine

doing better with yet more coils, iterating the process of finding the difference between the generated field and the ideal field profile, fitting to this difference and once again tweaking the parameters with respect to an OPERA simulation. This could be continued ad infinitum to recover and even surpass the T_1 time without Cryoperm, but represents a somewhat inelegant solution that furthermore is particularly vulnerable to any small changes in the magnetic environment.

Here we present an alternative method to arrive at a new design using the Laplace Equation for the magnetic scalar potential, assuming that the windings, whilst fixed to the outside of the guide tube, can take any shape. For any magnetic potential distribution specified on the boundary of a sourceless area, Laplace's equation enables us to uniquely solve for the potentials and thus the fields in the region. This means that, for any field configuration, there is an exclusive pattern of equipotentials were we to take a random cross-section at the edge of this volume. Conversely, by artificially holding the potentials at all points on the periphery, we can set the fields within the interior. This is achieved by winding the coils along the equipotentials, and is in fact implicitly used in the design of a solenoid, where all coils are perpendicular to the field front, and in the more complicated designs of the $\cos\theta$ and spherical coils.

In our instance, having simulated the fields numerically, we can simply find the scalar magnetic potential values on the proposed coil surface. The required equipotentials can be found from the difference in potential between the simulations with and without the SQUIDs. In the first approximation shown in Figure 5.17, which illustrates the coil windings if we are to use a single-layered coil, the wiring of the closed equipotential loops presents a complication, since any connections between other windings would have a component of orientation parallel to the coil axis, which will generate a field in the transverse direction. The technical construction of such a coil, let alone any such design imperfections, is a challenge. However, the idea of using precision-engineered coils is a powerful one and has been demonstrated to great potential by members of the SNS collaboration [148].

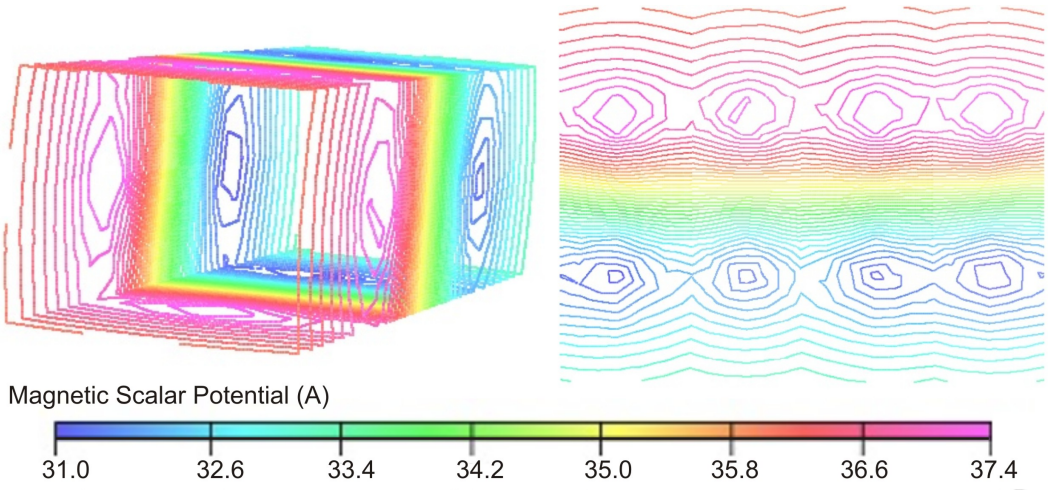

FIGURE 5.17: The 'ideal' coil windings as derived by solving Laplace's equation, for a coil with the same dimensions as the original SQCC (7.8 cm x 6.6 cm x 10 cm,), but centred on the SQUIDs. Note the effect of the Cryoperm in bunching the fields in the centre, due to the concentration of field by the material. The windings cannot be realised in closed equipotential loops at its ends.

5.3 Summary

This chapter detailed investigations into the Cryoperm shielding from two different persepectives: one concerning its benefit in screening out the external field during cooling, and another regarding its potential to disrupt the field pattern in the neutron guide as it does so. In the former, we saw that there is extra noise on the SQUID channel if it is run without Cryoperm, and found evidence that this was related to the cooling field. In the expected field level at ILL of 300 μT, these measurements demonstrate a factor of 1.5 difference in total noise voltage between the SQUIDs operated with and without Cryoperm. This discrepancy was reasoned to decrease with temperature if the noise source is thermally-activated. However, this remains a theory, and the proposed dependence of the noise on the temperature and on the inductance of the input were not verified to enable us to make a more accurate prediction of its manifestation in the signals of SQUIDs run in CryoEDM.

Simulations of the magnetic elements of CryoEDM in OPERA showed the Cryoperm to cause a depression in the field along the guide tube. This field profile was in

turn analysed using a Monte Carlo simulation, which tracks the neutron polarisations as they move throughout the volume. The distortion in field was seen to have a profound impact on the lifetime of the neutron polarisations, reducing T_1 by up to a factor of 10 for the cases examined. Compensating for the effect is possible by producing a field to oppose and fill in this depression. Different designs for a SQUID compensation coil were examined that could potentially restore the T_1 times: but equally, it must itself be properly characeterised to ensure that the field it produces does not induce inhomogeneities that could further depolarisation.

It is difficult to estimate the tolerable noise voltage or, as we mentioned, its expected level without additional information as to its nature – Chapter 6 will also discuss this. If the difference in the noise voltage between the unshielded and shielded SQUIDs is predicted to fall then, as it stands, consideration of both of these lines of thought lead us to the conclusion that the SQUIDs should be run without Cryoperm. If this decision is made than added precautions should be taken. For instance, a set of coils could be installed around the SQUIDs that are operated in feedback with external fluxgates, generating a field to cancel out the ambient field as the SQUIDs are cooled. The coils are solely energised during the cooling process so as not to disturb the neutrons throughout data-taking. The niobium cap could be equipped with a heater to release any trapped flux if the SQUID is seen to be operating in a particularly volatile state. These are all ideas that could be incorporated into CryoEDM, but add further complexity. In a future cryogenic project, the simplest solution is to use Cryoperm but, having considered its interplay with the rest of the apparatus, position the SQUIDs where they may still operate and be easily accessed, but cause little disturbance to the operation of the experiment.

Chapter 6

The Simulation of the SQUID and Neutron Magnetic Signals

An aim of this thesis is to illustrate how the SQUID system could be used to determine the magnetic fields within the neutron cells. To demonstrate how a 0.1 pT resolution could be consistently achieved, a statistically large data set of magnetometry measurements is required in order to comprehensively test the scheme suggested in Chapter 7. Here we document the development of software and algorithms which allow us to simulate temporally variant magnetic data.

This chapter begins by outlining each of the components that are used in the model of the ILL setup, in which the magnetic environment is treated as randomly distributed dipole sources. A major focus is the shielding, since the SQUID system has been devised considering the field variations in the vicinity of the pickup loops to be at the most on the order of nanotesla. This was dealt with via a numerical approach to solving the Generalised Laplace Equation, yielding the magnetic potentials in the volume of interest. This chapter details this framework and methods used to improve it, and how the magnetometry signals are found from this potential map. Lastly, it summarises the outcome of these simulations, ahead of the usage of their data in Chapter 7.

6.1 The Components of the Model

A magnetic representation of the CryoEDM setup was devised in an object-orientated architecture, whose central components are the sensing areas and magnetic sources (Figure 6.1). This allows the user to piecewise construct any configuration of these elements, and thus the scope is not limited to the 12-SQUID system of CryoEDM that we focus on for much of this chapter. In addition, the model incorporates the mu metal and superconducting shields. This allows it to simulate the attenuation of the fields by these components. Each of these elements is now discussed in turn.

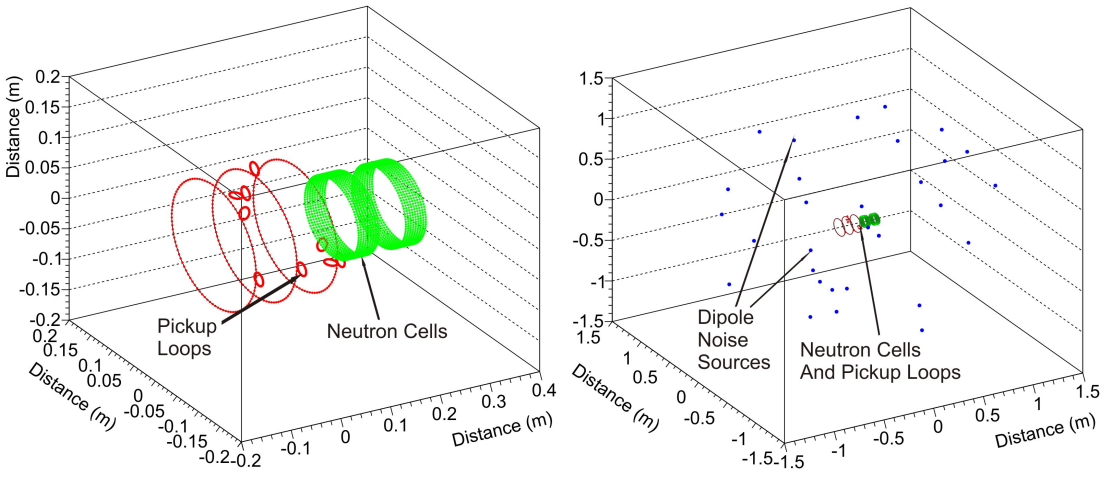

FIGURE 6.1: An example of the positions of simulated dipoles with respect to the pickup loop (red) and neutron cells (green). The dipoles are randomly distributed, whilst the loop ensemble is as in Chapter 3.

6.1.1 The Magnetic Background

The magnetic environment is simulated as a series of sources external to the shielding. These should dominate the fluctuations in the signal, provided that any internal sources have been well-characterised, as would be the strict requirement in any nEDM experiment, such that their contributions can in principle be identified

and accounted for. As the sources will generally be at large distances external to the mu metal shields, they are approximated as point-like dipoles, with any quadrupole (whose contribution to the field falls as $1/r^4$) or higher terms negligible in comparison.

Each dipole is modelled as a stationary source with a temporally variant signal consisting of a collection of sinusoidal waveforms, $B_0 \sum_n A_n \cos(\omega_n t + \theta_n)$, where the θ_n are arbitrary phases, and the $A_n = 1/\omega_n$ are the amplitudes of each contribution, normalised such that B_0 then represents the magnitude of the dipole. In the frequency domain, this produces peaks in the $1/f$ part of the SQUID spectrum, motivated by the observation that SQUIDs run as magnetometers typically show an enhanced $1/f$ component. Each dipole contributes to a data point a value of

$$\left[B_0 \sum_n \frac{1}{\omega_n^2} \sin\left(\omega_n t + \theta_n\right) \right]_{t_1}^{t_2} \tag{6.1}$$

multiplied by the relevant transfer function, where $t_2 - t_1$ is equivalent to the length of a measurement, be it 1 ms for a SQUID, or 300 s for the neutron cells.

The dipoles are each given a random orientation, position and magnitude to reflect the fact that the environment in the ILL experimental hall has not been mapped. However, they are assumed stationary as over the course of data-taking as the majority of these sources are expected to remain fixed or oscillate with small loci relative to their distance from the pickup loops. Such sources may include current loops or magnetic objects, which should give transient signals due in part to fluctuating conditions in the laboratory. Having said this, there will be isolated sources that will move significantly with respect to the apparatus during measurements. Those which are readily identified, such as the overhead crane, could be combatted by excluding any data during which they are active and recalibrating the magnetometry system. Other less conspicuous sources, such as vibrating components within the experiment's cryogenics or electronics, are more difficult to address. To show that these can, to some extent, be reconciled within the chosen framework,

we first begin with the general expression for a dipole field at a point:

$$\vec{B}(\vec{r}) = \frac{\mu_0}{4\pi}\left(\frac{3\vec{r}(\vec{m}\cdot\vec{r})}{r^5} - \frac{\vec{m}}{r^3}\right), \tag{6.2}$$

where $\vec{m}$ is the dipole moment (the unit orientation vector of the dipole multiplied by its magnitude) and $\vec{r}$ is the relative position vector between the dipole and the measurement point, r being its magnitude. If the given source were to oscillate in the direction of $\vec{r}$, then its position vector will vary as $\vec{r} \to \vec{r}(1 + A\cos(\omega t))$, where A is an arbitrary magnitude and ω is the frequency with which it oscillates. Substituting this into Equation 6.2, this gives the field magnitude, $B(\vec{r}) = \vec{B}(\vec{r})\cdot\hat{n}$, at this point as

$$B(\vec{r}, t) = \frac{B_0(\vec{r})}{(1 + A\cos(\omega t))^3}, \tag{6.3}$$

where B_0 is the original value for the stationary dipole at $\vec{r}$. The question we have is whether this is expressible in terms of our time-varying stationary dipole functions, i.e. able to be decomposed as a Fourier series with real coefficients, wherein each term can be attributed to a frequency on one or more dipoles. The expression $(1 + A\cos(\omega t))^{-3}$ can be written as a convergent sequence of cosine terms of increasing frequency provided that $A < 1$, such that the denominator is always positive and non-zero, the graph never asymptotes, and there are no discontinuities in the periodic function. An object moving at a constant speed could be envisaged as such a source with a low frequency, much greater than the length of a measurement cycle.

What if the dipole were to oscillate in a direction perpendicular to the relative position vector, i.e. with $\vec{r} \to \vec{r} + A(\vec{m}\times\vec{r})\cos(\omega t)$ (or $\vec{r} \to \vec{r} + A\vec{r}\times(\vec{m}\times\vec{r})\cos(\omega t)$)? This gives

$$B(\vec{r}, t) = \frac{B_0(\vec{r}) + C\cos(\omega t) - D\cos^2(\omega t)}{(1 + A^2\cos^2(\omega t))^{5/2}}, \tag{6.4}$$

where $C = \hat{n}\cdot(\vec{m}\times\vec{r})(\vec{m}\cdot\vec{r})A$ and $D = (\hat{n}\cdot\vec{m})r^2A^2$. With a similar argument, the denominator can be shown to be a Fourier series, which yields another series when multiplied by the numerator.

Lastly, with components parallel and perpendicular to $\vec{r}$, with say $\vec{r} \to \vec{r}(1 + A\cos(\omega t)) + B(\vec{m} \times \vec{r})\cos(\omega t) + C\vec{r} \times (\vec{m} \times \vec{r})\cos(\omega t)$, we have

$$B(\vec{r}, t) = \frac{B_0(\vec{r}) + F\cos(\omega t) + G\cos^2(\omega t)}{(1 + 2A\cos(\omega t) + (A^2 + B^2 + C^2)\cos^2(\omega t))^{5/2}}. \tag{6.5}$$

Again, the issue is the denominator, from which we infer that the inequality $(A - 1)^2 + B^2 + C^2 > 1$ must hold if this can be represented as a Fourier series.

A pickup loop does not provide a point measurement of the field, so it is difficult to quantify exactly how much of an effect source movement has on the fields registered by the loop ensemble. As an example, Figure 6.2 shows the percentage change in flux through a loop when a source initially centred on-axis is moved parallel and perpendicular to this axis. Off-axis perturbations decrease catchment of field perpendicular to the plane of the loop, whilst increasing that parallel to it. The effect is therefore most pronounced for a dipole which is oriented perpendicular to the loop normal, where the initial flux is zero by symmetry: however, this is also the case where the least field is coupled into the loop at a perturbed position. At a metre away, the closest the source could be in our model, a source aligned with the axis would have to move by 5.8 cm to produce a 1% difference in the field. These findings suggest that, especially if they are isolated sources, whose contributions are swamped by others, the likely impact moving sources would have on the integrated flux could well be negligible.

However, the model does not account for a more complicated motion involving, for instance, rotation of the dipole moment, and one can put forward many such counterexamples that cannot easily, or are impossible to, be expressed as contributions of stationary sources. If the goal is to generate as accurate a representation of the magnetic environment as possible and include every possible motion, appearance and disappearance of these sources, the model devised here is certainly lacking. But, it is certainly capable of generating a feasible magnetic background, albeit limited in the movement of its sources. Moreover, as has been argued, with little

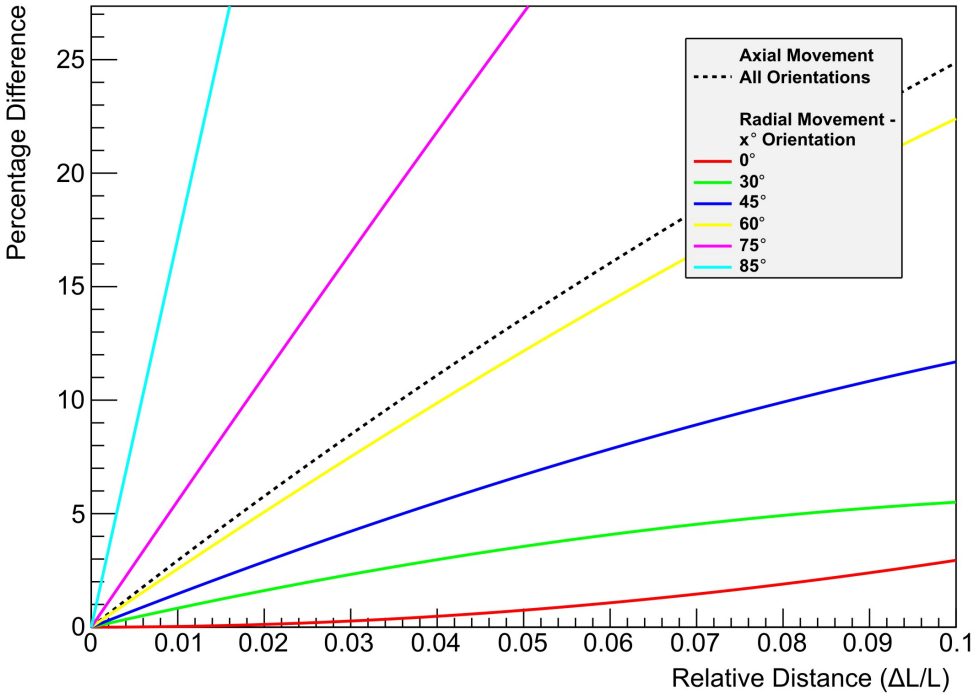

FIGURE 6.2: The percentage change in flux coupled into a pickup loop for a dipole initially centred on the loop axis, and moved in the axial or radial directions.

but speculation concerning the ILL environment, such a static background could well predominate in most measurements, if major disturbances can be properly noted using, for example, the external fluxgates. For an exposition of a potential scheme to extract the neutron cell signal, this model should therefore be sufficient to generate data with which to test it.

6.1.2 The Pickup Loops

At the heart of the CryoEDM model lie the pickup loops (Figure 6.1), with dimensions and locations as specified in Chapter 3. The first step to finding the SQUID signal associated with each loop is to determine the flux that threads the given sensing area at any moment. For each pair of loop and source, a **transfer function** is defined as the flux traversing the loop for unit magnitude on the source. At a given instant in time, the signal from the loop can be found by multiplying

each of its transfer functions by the present magnitude of the corresponding source and summing them together.

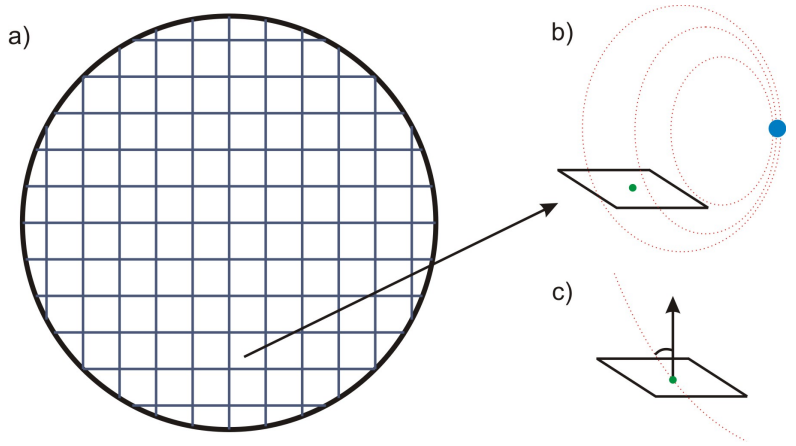

FIGURE 6.3: a) Each pickup loop is sectioned into a grid with elements of equal area, except at the edges, so that there is no positional bias in calculating the flux across the entire loop; b) the field at the centre point of the grid element is found and c), resolved in the direction of the normal to the pickup. This is multiplied by the area of the element, and the contributions of each element summed to give the total flux across the pickup loop. The accuracy can be improved by making the grid finer.

The integration of the field over the pickup loop cannot be readily solved for general dipole and loop positions and orientations. Therefore, to determine the transfer function, the chosen pickup loop is sectioned into a grid with elements of equal area (see Figure 6.3), to ensure there is no directional bias in the calculation of the flux. The field across each smaller area is then assumed to be uniform and estimated as that at its centre point. The flux through this element is taken as the field in the direction of the normal to the pickup loop at this point multiplied by its area, that is $\vec{B}(\vec{r}) \cdot \hat{n}$. This treatment applies to all areas, even those which are cut off by the loop and approximated as triangles and trapezia. The sum total of these fluxes yields an estimation of the total flux which traverses the loop.

The inaccuracy due to the discretisation can always be reduced by increasing the number of grid elements. But to get a feel for its size and how many elements are reasonable, we can examine a simple scenario that can be reduced to an analytical expression. In the case of the source on-axis, the integrand becomes, using

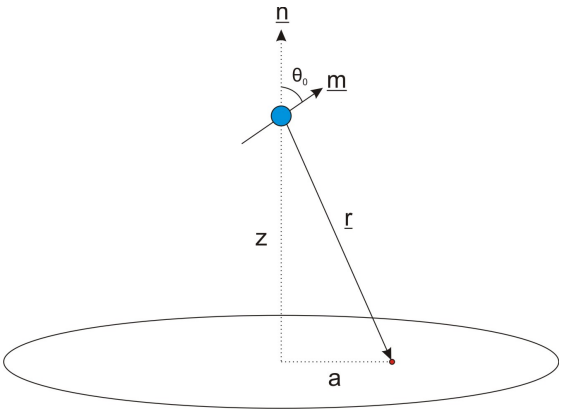

FIGURE 6.4: A diagram relating a sample point on the pickup loop (red) to a dipole source on-axis (blue). The relative position vector ($\vec{r}$), normal vector of the pickup loop ($\vec{n}$) and the dipole moment of the source ($\vec{m}$) are labelled, as are the relative displacements in the axial (z) and radial (a) directions.

Equation 6.2 and referring to Figure 6.4,

$$B_0 \frac{(2z^2 - a^2)\cos\theta_0 \pm 3az\sin\theta_0}{(a^2+z^2)^{5/2}}, \tag{6.6}$$

where $B_0 = {}^{\mu_0|\vec{m}|}/_{4\pi}$, and the second term in the numerator is either constructive or destructive depending on the position of the sample point (i.e. whether it lies left or right of the axis in Figure 6.4). This antisymmetric part is cancelled in an integral over the full loop i.e. over 2π, giving the flux $\vec{B}\cdot\vec{A} = \int \vec{B}\cdot\hat{n}\,\mathrm{d}A$ for a loop of radius R as

$$B_0 \int_0^{2\pi}\int_0^{R} \frac{(2z^2 - a^2)\cos\theta_0}{(a^2+z^2)^{5/2}}\, a\mathrm{d}a\mathrm{d}\theta = B_0 \frac{2\pi R^2 \cos\theta_0}{(R^2+z^2)^{3/2}}. \tag{6.7}$$

A comparison between this real value and the discretised version, which divides the grid into 2000 segments radially, gives a 0.013% difference at all axial distances of concern for the simulation, for all orientations of the dipole. Although this error is likely bigger for a dipole off-centre, it should still be negligible with this number of grid elements compared to the noise sources we will deliberately introduce to perturb the signal and test the interpolation method.

One of these sources will be the intrinsic noise of the SQUID, which was examined in some depth in Section 5.1.5. This predicted how the two colours of noise,

white and pink, would transform from the laboratory conditions to those at ILL. However, in the case of the $^1/_f$ noise, its temperature dependence and origin are unknown, making an estimate of its effect on the CryoEDM SQUID signals difficult. We therefore take a different approach here and ask what is the largest noise voltage that could be tolerated. This will depend on how the SQUID signals are used. Chapter 7 will present an idea of combining the signals to yield an approximation to the neutron cell fields. If all 12 signals are used and they are weighted equally, then naïvely, to reach a resolution of 0.1 pT, the precision of each must be $\sqrt{12} \times 0.1\,\mathrm{pT}$. If the calibration factor is around 3 V/nT, then the noise voltage must be less than $\approx 1\,\mathrm{mV}$ – this is the estimate that will be used for much of Chapter 7, which will then also look to see how the performance is affected by the level of noise. To add this to the simulated signals, we recognise it as an RMS voltage and approximate the contribution as a Gaussian with a probability density function of

$$P(x) = \frac{1}{V_{\mathrm{RMS}}\sqrt{2\pi}} \exp\left(-\frac{x^2}{2V_{\mathrm{RMS}}^2}\right). \tag{6.8}$$

Another possible cause of departures from the true signal are quantum artefacts known as resets. These occur when the voltage signal approaches the extremes of the range of the feedback system at around $\pm 10\,\mathrm{V}$. To ensure the feedback remains functional, its output is temporarily zeroed, from which a new working point is established. This will differ from the original output by a multiple integer of flux quanta, and thus a voltage discontinuity manifests as the feedback is restored.

In a slowly varying magnetic environment, it could well be that there are no more than one or two resets in tens of hours of data-taking, which can therefore easily be avoided by excluding data around each artefact. However, since the resets are uncorrelated, there could be many such regions where one of the twelve SQUIDs is affected. Moreover, the chosen scheme in Chapter 7 may require recalibration after a reset. Resets should therefore be corrected to maximise the longest continuous stretch of data – McCann outlines the protocol in detail [124].

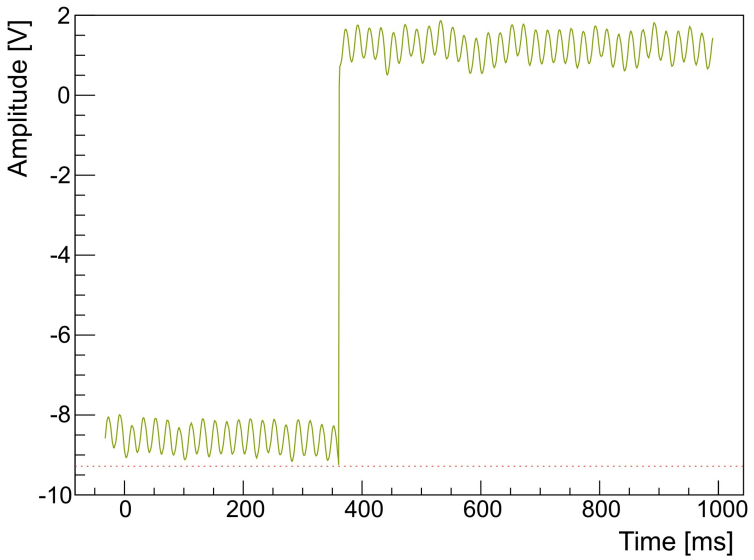

FIGURE 6.5: A reset on a SQUID channel, where the signal is offset by an integer number of flux quanta as the voltage hits the limit of the feedback range (about ±9.3 V, illustrated by the dotted red line).

The size of a reset on a given SQUID channel is always the same if its configuration does not change. Therefore, resets can be manually invoked and then measured. The sizes should be primarily affected by the white noise, thus demonstrating a Gaussian distribution. With enough measurements, this should in theory allow near-perfect correction of a reset, and they are therefore not included in our model.

6.1.3 The Ramsey Cells

The two Ramsey cells are taken as cylinders of 125 mm diameter and 45 mm height: they are positioned with a separation of 50 mm between the two cells, with the start of the first cell some 70.5 mm from the nearest pickup loop. The concept of a transfer function is not useful in this instance, since the neutrons are sensitive to the volume-averaged field, $\langle B \rangle$. Instead, having similarly divided the cell into a grid but in three dimensions, we sample the field at the centre of each element to yield a field map for each source. These are scaled by each dipole signal, superimposed and averaged over to give $\langle B \rangle$ for each 300 s free precession period.

The neutron signal is broadened by a number of effects, such as depolarisation over the course of a measurement cycle and losses at the detector end. These factors affect the overall neutron counts and, as we have seen in Chapter 2, can be amalgamated to give the visibility of the Ramsey curve, α. To calculate the statistical noise, we first return to Equation 2.9

$$N_{u \text{ or } d} = \bar{N}(1 \mp \alpha \cos(\lambda T)),$$

with λ being the difference between the applied AC frequency, ω, and the neutron Larmor resonance frequency, $\omega_0 = \gamma B$. The neutron counts therefore give a measure of the magnetic field and it is its associated statistical uncertainty which is the principal source of any uncertainty in B. If the magnetic field has drifted since the previous measurement, then the present resonant frequency will be different from the previous estimation by a small perturbation, $\delta\omega_0$. The difference between the up and down counts, $\Delta N = N_u - N_d$, for a given measurement would then be

$$\Delta N = -2\bar{N}\alpha \cos(\lambda T) = -2\bar{N}\alpha \cos(\omega - \omega_0 + \delta\omega_0)T, \tag{6.9}$$

where $\bar{N}$ is the average occupation of either the up or down state across the central Ramsey fringe, assuming this to be the same for both states. Using double-angle formulae,

$$\begin{aligned}\Delta N &= 2\bar{N}\alpha[\sin(\omega - \omega_0)T \sin(\delta\omega_0)T - \cos(\omega - \omega_0)T \cos(\delta\omega_0)T] \\ \delta N &\approx 2\bar{N}\alpha\delta\omega_0 T,\end{aligned}$$

if we are at one of the working points, where $\omega - \omega_0 \approx \pi/2$, and the difference between the counts for the two states is small, δN. Rearranging, this gives:

$$\delta B = \frac{\delta N}{2\bar{N}\alpha\gamma T},$$

where we have used $\delta\omega_0 = \gamma\delta B$. $2\bar{N}$ is simply the total number of neutrons

detected in this meausrement, i.e. $N = N_u + N_d$. Therefore the uncertainty in B, $\sigma(B) = \sigma(\delta B)$, can be expressed as:

$$\sigma(B) = \frac{1}{\alpha\gamma T}\sigma\left(\frac{\delta N}{N}\right).$$

The uncertainty on the right hand side can be simplified as follows:

$$\begin{aligned}\sigma\left(\frac{\delta N}{N}\right) &= \left(\frac{\delta N}{N}\right)\left[\left(\frac{\sigma(\delta N)^2}{\delta N}\right)^2 + \left(\frac{\sigma(N)}{N}\right)^2\right]^{1/2} \\ &\approx \left(\frac{\delta N}{N}\right)\left(\frac{\sigma(\delta N)}{\delta N}\right) = \frac{1}{\sqrt{N}},\end{aligned}$$

because $\sigma(\delta N) = \sigma(N) = \sqrt{N}$, as they are both governed by counting statistics, and $\delta N \ll N$. This gives a final expression for the uncertainty of the B-field as:

$$\sigma(B) = \frac{1}{\alpha\gamma T\sqrt{N}}. \tag{6.10}$$

This affects every magnetic field measurement made using the neutrons themselves as the magnetometer. For the aimed for values in CryoEDM of $\alpha = 0.9$, $T = 300\,\mathrm{s}$ and $N = 7 \times 10^5$, this leads to $\sigma(B) \approx 24.2\,\mathrm{fT}$. As with the intrinsic SQUID noise, this is added using a random generator based on a Gaussian distribution.

6.1.4 The Mu Metal & Superconducting Shields

All neutron EDM experiments will have some magnetic shielding in place to attenuate external fields that will detract from the desired magnetic field in the neutron cells. In the case of CryoEDM, this constitutes a concentric closed cylinder of high permeability material, which primarily reduces fields transverse to the cells, and a superconducting solenoid with accompanying shield, which generates and freezes in a stable, static field of $5\,\mu\mathrm{T}$ in the direction of the cylinder axis. This shielding must be modelled since, whilst predominantly axial fields can be enforced through careful choices of dipole positions and orientations, the field pattern would be

quite different from a random configuration whose transverse components have been reduced. There is no known analytical method for calculating the field at an arbitrary point within a cylindrical shield due to an arbitrary source outside it. Hence we use a numerical approach, discretising the volume of interest into a set of sample points on a 3D grid, solving Laplace's equation for the magnetic scalar potential, Φ, within it, taking the gradient at each point to determine the fields.

Whilst an accurate representation of the CryoEDM system is not necessary, given this is a hypothetical pursuit, the model here remains as faithful to the experimental configuration as possible. The mu metal shield, according to Grozier [147], consists of three layers each of thickness 1.6 mm. The dimensions are problematic since they would require incredibly fine grid spacing to resolve, which in turn implies a long execution time that makes the program infeasible to run multiple times. The transverse shielding factor is arguably more important than the axial to maintain field homogeneity in the neutron cells. For a multi-layered shield, it can be approximated by

$$S^T = \prod_{i=1}^{n-1} S_n^T S_i^T \left(1 - \frac{r_i}{r_{i+1}}\right)^2, \tag{6.11}$$

where S_i^T is the transverse shielding factor for a single layer, and r_i is its radius [83]. Equation 4.7 indicates the approximation $S_i^T \approx {}^{\mu_r t_i}/_{2r_i}$. Therefore, the layers of $\mu_r = 50000$ with combined shielding factor of ~ 7000 may be matched by a single shield of thickness 8 cm and $\mu_r = 10000$ spanning the radii of the tri-layer shield, i.e. with inner radius 0.509 m and outer radius 0.589 m, as it is in our model.

Lastly, we also include the superconducting shield and solenoid as a single open cylinder of length 2.68 m, radius 0.36 m and thickness 1.6 cm, which is roughly as is documented by Grozier [147]. This is simulated with $\mu_r = 10^{-8} \approx 0$ as a first approximation to the flux exclusion from the bulk superconductor. As detailed in Chapter 4, to accurately model its behaviour, one would need a vector formulation

and a discretisation of the London equations, which is beyond the scope of this initial work. The implications are discussed further at the end of this chapter.

6.2 A Finite Difference Solver for Magnetic Fields

A bespoke piece of software was developed to determine the attenuation of the dipole fields by the shielding. This is important, since the capability of the SQUIDs to follow the fields in the Ramsey cells, under the scheme outlined in Chapter 7, will be affected by the size, homogeneity and prevailing direction of the fields in their vicinity. The resultant fields from each dipole are calculated separately, enabling us to scale and superimpose them as we wish to create a complex magnetic environment. A bank of many hundreds of dipoles is required to ensure diversity in the field maps and truly test the scheme in Chapter 7.

The reasons for developing a tailor-made code are twofold. First, it can be integrated with the software for the magnetic model, and can be readily modified to incorporate future additions or changes. Second, the problem at hand is not easily accommodated in commercial programs. Take, for example, OPERA – it does not allow the simulation of point sources, and therefore a dipole could be modelled as a coil at large distance. It requires a minimum amount of user interaction, first to set the parameters of this 'dipole,' which complicates randomising its properties, and then, in separate instances, to start the solver and analyse the solution. OPERA is designed for pre-defined, 'real-world' problems, and is not efficient for a Monte Carlo style process, where we require several hundreds of simulations.

The framework used here is based on the Finite Difference Method (FDM). This is similar to the Finite Element Method (FEM) (used in e.g. OPERA and COMSOL), but is somewhat easier to understand intuitively and program. The shield geometry can be realised in cylindrical coordinates without the flexible mesh of an FEM and the problem does not involve multiple inhomogeneities, wherein there

is a convergence penalty when implemented in an FEM scheme. The formulation detailed in this section has been written from scratch, adapting some of the concepts found in [149], [150] and [151] since the incorporation of many of its features, such as the cylindrical grid and shields, is not readily reported in the literature.

6.2.1 The Algorithm

The scheme presented here ultimately solves the Laplace Equation for the magnetic scalar potential in a given region, having defined the potential at its boundaries. In cylindrical coordinates, the Laplace Equation becomes

$$\nabla^2\Phi(r,\theta,z) = \frac{1}{r}\frac{\partial}{\partial r}\left(r\frac{\partial\Phi}{\partial r}\right) + \frac{1}{r^2}\frac{\partial^2\Phi}{\partial\theta^2} + \frac{\partial^2\Phi}{\partial z^2} = \nabla\cdot\mathbf{M}, \tag{6.12}$$

where $\mathbf{M}$ is the magnetisation. The volume is discretised as a 3D grid with spacings Δr, $\Delta\theta$ and Δz in each respective orthogonal direction. It is assumed that there are no sources of permanent magnetisation throughout it (i.e. that the right-hand side of Equation 6.12 is zero) and it is defined with diameter and height 2 m greater than that of the shields to form a background, where the potentials on each face of the volume should be unperturbed by the presence of the shield. At each point on an outer face, the potential due to the dipole source with magnetic moment $\vec{m}$ is specified as

$$\Phi = \frac{\vec{m}\cdot\vec{r}}{4\pi r^3},$$

which provide the Dirichlet boundary conditions for the problem.

Since the magnetic scalar potential defines a smooth function across this volume, we note that the value of the potential at any position is correlated with those at adjacent points. Knowing the potential at each edge of the grid, it should therefore be possible to interpolate a value at any point within the grid. Here we invoke the FDM, with centred difference approximations to the partial derivatives

in Equation 6.12 obtained from first principles. The first derivative is given by

$$f'(x_0) \approx \frac{f(x_0+h) - f(x_0-h)}{2h},$$

from which we can say

$$f''(x_0) \approx \frac{f'(x_0+h/2) - f'(x_0-h/2)}{h} = \frac{f(x_0+h) - 2f(x_0) + f(x_0-h)}{h^2} + O(h^2).$$

These can be used to write a first-order approximation for Equation 6.12 as

$$\nabla^2\Phi(r,\theta,z) = \frac{\Phi_{i+1,j,k} - 2\Phi_{i,j,k} + \Phi_{i-1,j,k}}{(\Delta r)^2} + \frac{1}{r_i}\frac{\Phi_{i+1,j,k} - \Phi_{i-1,j,k}}{2\Delta r} + \frac{1}{r_i^2}\frac{\Phi_{i,j+1,k} - 2\Phi_{i,j,k} + \Phi_{i,j+1,k}}{(\Delta\theta)^2} + \frac{\Phi_{i,j,k+1} - 2\Phi_{i,j,k} + \Phi_{i,j,k-1}}{(\Delta z)^2}, \quad (6.13)$$

where i, j, k represent indices for the r, θ, z coordinates respectively. Rearranging this formula, we find that, under this approximation, the potential at a given point in the mesh is related to that of its seven nearest neighbours:

$$\Phi_{i,j,k} = \frac{a_r}{r_i}\left(r_{i-1/2}\Phi_{i-1,j,k} + r_{i+1/2}\Phi_{i+1,j,k}\right) + a_t\left(\Phi_{i,j-1,k} + \Phi_{i,j+1,k}\right) + a_z\left(\Phi_{i,j,k-1} + \Phi_{i,j,k+1}\right), \quad (6.14)$$

where $r_{i\pm1/2} = r_i \pm 1/2\Delta r$,

$$a_r = \frac{r_i^2(\Delta\theta)^2(\Delta z)^2}{2(r_i^2(\Delta\theta)^2(\Delta z)^2 + (\Delta r)^2(\Delta z)^2 + r_i^2(\Delta\theta)^2(\Delta r)^2)},$$

$$a_t = \frac{(\Delta r)^2(\Delta z)^2}{2(r_i^2(\Delta\theta)^2(\Delta z)^2 + (\Delta r)^2(\Delta z)^2 + r_i^2(\Delta\theta)^2(\Delta r)^2)},$$

$$a_z = \frac{r_i^2(\Delta\theta)^2(\Delta r)^2}{2(r_i^2(\Delta\theta)^2(\Delta z)^2 + (\Delta r)^2(\Delta z)^2 + r_i^2(\Delta\theta)^2(\Delta r)^2)}.$$

The computational molecule in this 3D finite difference scheme is therefore a seven-point stencil, as illustrated in Figure 6.6, which can be used to solve for the potential inside the grid at all positions. In cylindrical coordinates, since the solution must be periodic in θ, we must introduce the boundary condition

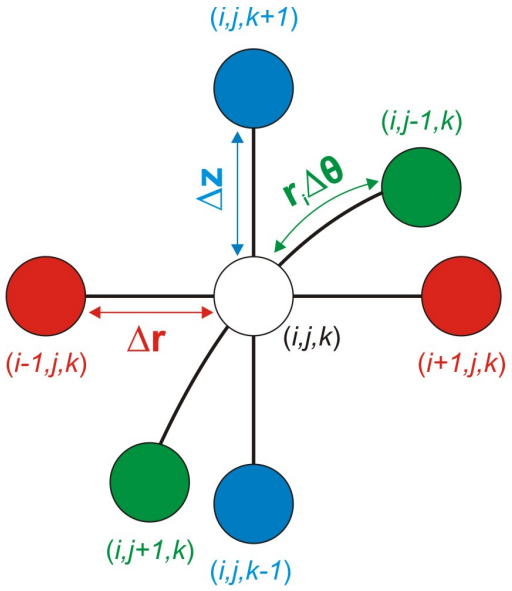

FIGURE 6.6: Computational molecule for the finite difference algorithm.

$\Phi(r, 0, z) = \Phi(r, 2\pi, z)$. Further, to overcome the pole at $r = 0$, we adopt a staggered grid approach, where r_i is defined as $r_i = (i + 1/2)\Delta r$. At $i = 0$, $r_{i-1/2} = 0$ such that the $i - 1$ term drops out from the equation.

The algorithm thus defined can successfully reproduce the field pattern from a dipole in vacuum: the next step is how to incorporate any area of different permeability. The situation is described by the Generalised Laplace Equation,

$$\nabla \cdot (\mu \nabla \Phi) = \nabla \mu \cdot \nabla \Phi + \mu \nabla^2 \Phi = \nabla \cdot \mathbf{M} = \mathbf{0}, \tag{6.15}$$

which we will again need to discretise to solve for numerically. The centred difference approach can be used for the first derivatives $\nabla\mu$ and $\nabla\Phi$ to arrive at an expression which consists of Equation 6.14 and an extra contribution due to the varying permeability. However, in practice, this does not deal well with the abrupt change in μ at the air-mu metal interface. To get around this problem, we note that there is no need for μ to be sampled at identical points as Φ, so using a finer grid for μ may help in tackling the discontinuity. We define a permeability grid with half the spacing of the potential grid. An alternative finite difference

expression can then be derived using the divergence theorem:

$$\int_V \nabla \cdot (\mu \nabla \Phi)\, \mathrm{d}V = \int_V \nabla \cdot \mathbf{M}\, \mathrm{d}V = 0$$
$$= \oint_S \mu \nabla \Phi \cdot \mathrm{d}\mathbf{A} = \oint_S \mu \left(\frac{\partial \Phi}{\partial r}\hat{\mathbf{r}} + \frac{1}{r}\frac{\partial \Phi}{\partial \theta}\hat{\boldsymbol{\theta}} + \frac{\partial \Phi}{\partial z}\hat{\mathbf{z}} \right) \cdot \hat{\mathbf{n}} \mathrm{d}A. \quad (6.16)$$

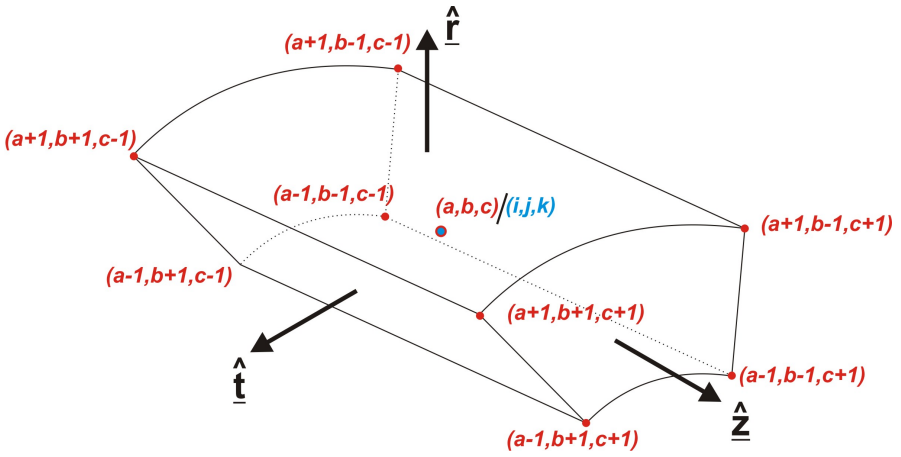

FIGURE 6.7: Volume element centred on point (i, j, k) of the potential grid, showing the finer permeability grid at half the spacing, with indices (a, b, c) at the point in question.

If the volume element encloses one sample point for the potential, the right-hand side is a closed surface integral around this point. We can split this into six contributions, for each of the faces that compose the closed surface, as shown in Figure 6.7. The integral for the top face becomes

$$\int_{S_1} = \int_{z-\frac{z}{2}}^{z+\frac{z}{2}} \int_{\theta-\frac{\theta}{2}}^{\theta+\frac{\theta}{2}} \mu \left(\frac{\partial \Phi}{\partial r}\hat{\mathbf{r}} + \frac{1}{r}\frac{\partial \Phi}{\partial \theta}\hat{\boldsymbol{\theta}} + \frac{\partial \Phi}{\partial z}\hat{\mathbf{z}} \right) \cdot \hat{\mathbf{r}}\, r \mathrm{d}\theta \mathrm{d}z = \int_{z-\frac{z}{2}}^{z+\frac{z}{2}} \int_{\theta-\frac{\theta}{2}}^{\theta+\frac{\theta}{2}} \mu \frac{\partial \Phi}{\partial r}\, r \mathrm{d}\theta \mathrm{d}z. \quad (6.17)$$

This surface forms the intersection between the volume element for $\Phi_{i,j,k}$ and $\Phi_{i+1,j,k}$. The partial differential can be approximated by a central difference formulation and the permeability is taken as the average of the values at the four corners of the given face, assuming it to be constant across it, which are grid points

for the staggered μ grid. This yields an expression for this first integral of

$$\int_{S_1} \approx \frac{(\mu_{a+1,b-1,c+1} + \mu_{a+1,b+1,c+1} + \mu_{a+1,b+1,c-1} + \mu_{a+1,b-1,c-1})}{4} \frac{\Phi_{i+1,j,k} - \Phi_{i,j,k}}{\Delta r} r_{i+\frac{1}{2}} \Delta\theta\Delta z. \tag{6.18}$$

Similar expressions can be derived for the other five faces so that the entire closed surface integral may be written as

$$\oint_S \approx p_1\Phi_{i+1,j,k} + p_2\Phi_{i,j+1,k} + p_3\Phi_{i,j,k+1} + p_4\Phi_{i,j,k-1} + p_5\Phi_{i,j-1,k} + p_6\Phi_{i-1,j,k} - p_0\Phi_{i,j,k}, \tag{6.19}$$

where

$$p_1 = \frac{(\mu_{a+1,b-1,c+1} + \mu_{a+1,b+1,c+1} + \mu_{a+1,b+1,c-1} + \mu_{a+1,b-1,c-1})}{4\Delta r} r_{i+\frac{1}{2}} \Delta\theta\Delta z,$$

$$p_2 = \frac{(\mu_{a+1,b+1,c-1} + \mu_{a+1,b+1,c+1} + \mu_{a-1,b+1,c-1} + \mu_{a-1,b+1,c+1})}{4r_i\Delta\theta} \Delta r\Delta z,$$

$$p_3 = \frac{(\mu_{a+1,b+1,c+1} + \mu_{a+1,b-1,c+1} + \mu_{a-1,b+1,c+1} + \mu_{a-1,b-1,c+1})}{4\Delta z} r_i\Delta\theta\Delta r,$$

$$p_4 = \frac{(\mu_{a+1,b+1,c-1} + \mu_{a+1,b-1,c-1} + \mu_{a-1,b+1,c-1} + \mu_{a-1,b-1,c-1})}{4\Delta z} r_i\Delta\theta\Delta r,$$

$$p_5 = \frac{(\mu_{a+1,b-1,c-1} + \mu_{a+1,b-1,c+1} + \mu_{a-1,b-1,c-1} + \mu_{a-1,b-1,c-1})}{4r_i\Delta\theta} \Delta r\Delta z,$$

$$p_6 = \frac{(\mu_{a-1,b+1,c-1} + \mu_{a-1,b-1,c-1} + \mu_{a-1,b-1,c+1} + \mu_{a-1,b+1,c+1})}{4\Delta r} r_{i-\frac{1}{2}} \Delta\theta\Delta z,$$

$$p_0 = p_1 + p_2 + p_3 + p_4 + p_5 + p_6.$$

This provides a new, seven-point stencil and one can verify that, with no mu metal and thus a constant permeability of unity across the volume, Equation 6.19 reduces to Equation 6.14. Exploiting the azimuthal symmetry of the problem, we need only specify four μ grid points per potential point.

6.2.2 The Relaxation Scheme

So far, we have developed a formulation using finite difference approximations that relates the potential at a given point to that of its neighbours. This is a discretisation of the Generalised Laplace Equation that will enable us to solve it numerically. In this form, we essentially have a linear set of simultaneous equations, with the grid points as variables. The system is expressible as a sparse matrix, so that iterative methods that involve splitting the matrix, rather than direct elimination methods, are preferable for finding a solution. Here we expand upon the idea of an iterative procedure, sticking with the former notation.

An intuitive choice is the Jacobi relaxation method, which reevaluates $\Phi_{i,j,k}$ using the FDM stencil on each iteration to generate successive approximations to the answer. In other words, using the variables defined in Equation 6.19,

$$\Phi^{n+1}_{i,j,k} = \frac{p_1\Phi^n_{i+1,j,k} + p_2\Phi^n_{i,j+1,k} + p_3\Phi^n_{i,j,k+1} + p_4\Phi^n_{i,j,k-1} + p_5\Phi^n_{i,j-1,k} + p_6\Phi^n_{i-1,j,k}}{p_1 + p_2 + p_3 + p_4 + p_5 + p_6}, \tag{6.20}$$

where n signifies the iteration number. The evaluations can be easily parallelised, since they solely depend on the old values of the potentials, but require that we store both sets of numbers at the same time.

An improvement can be made by realising that there is no need to compute the entire grid in one go, and that values can be used as soon as they become available: this is the basis of the Gauss-Seidel method [152]. Our nearest neighbour scheme facilitates a variant known as red-black ordering. The grid can be divided analogously to a chess board, with the red or odd points evaluated in one pass, and the black points in the second sweep, using the new red values. This removes any interdependence between the calculations such that they can once more be parallelised, since the order in which we determine the red/black points does not matter. The procedure is therefore very similar to the Jacobi method, with the

odd values updated as in Equation 6.20, and the even points with these just evaluated, $n + 1$ numbers. The revised values can be added straight into the same grid such that only one is needed, and the number of computations is effectively reduced by a half, meaning that it can converge up to twice as quickly as the Jacobi method.

It is here that we make the distinction between the iteration and convergence rates: the former is the inverse of the time taken to complete one full sweep of the grid, whereas the latter is related to the number of iterations required to reach the convergence criterion, which is nominally set to 10^{-8}. This represents the threshold that the residual must fall to. The residual is not equivalent to the error, but the two are intimately connected, such that a small residual corresponds to a more accurate answer: Appendix B.3 and Section 6.2.3 add depth to this discussion.

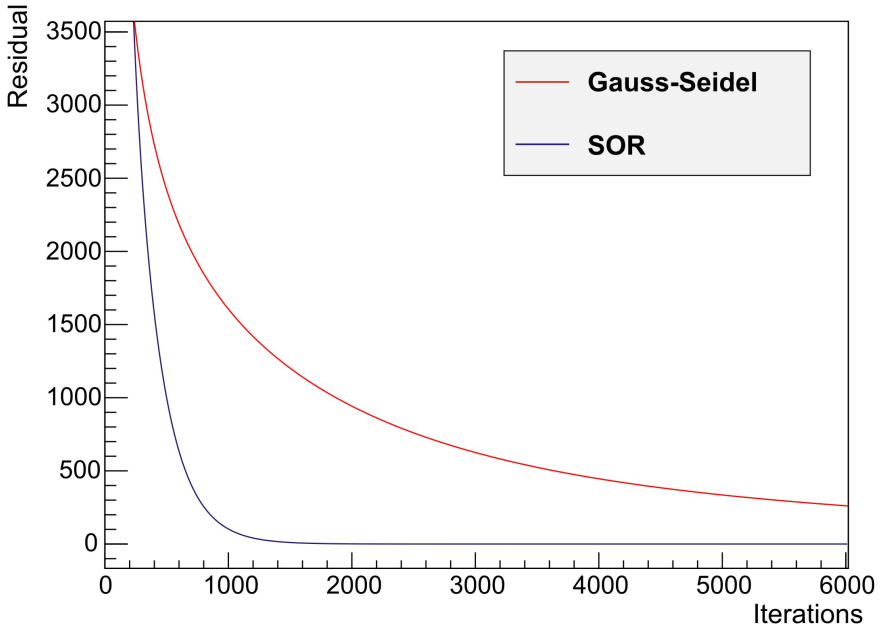

FIGURE 6.8: The Successive Overrelaxation (SOR) relaxation scheme shows improved convergence properties compared to the unadulterated Gauss-Seidel technique. For this particular instance, the latter reaches a residual of 10^{-8} after $\approx 1.2 \times 10^5$ iterations, 20 times more than required by the SOR method.

Implementing the Gauss-Seidel technique improves upon the convergence rate of

the Jacobi method but still leads to, relatively speaking, a slowly convergent process. The problem could benefit from a way to somehow 'damp' the solver to accelerate convergence (Figure 6.8). This is historically what is is referred to as 'relaxation,' although the term has now come to encompass any iterative procedure that converges upon a solution. The specific relaxation scheme we use here is that of Successive Overrelaxation (SOR) [153], wherein the value of $\Phi_{i,j,k}$ is overcorrected on each pass:

$$\Phi^{n+1}_{i,j,k} = \omega\Phi^{*}_{i,j,k} + (1-\omega)\Phi^{n}_{i,j,k}, \tag{6.21}$$

where $\Phi^{*}_{i,j,k}$ is the Jacobi or Gauss-Seidel description for the potential as in Equation 6.20, and ω is the overrelaxation factor, $1 < \omega < 2$. The effect of the SOR is analogous to critically damping the oscillator, which was formerly overdamped under the Gauss-Seidel prescription.

To improve the iteration rate, we capitalise on the opportunity for parallelisation. OpenMP is used to call a directive that generates multiple threads ahead of an iteration. The workload is distributed evenly amongst eight threads, allocated on a single machine/node – employing more threads or using dynamic scheduling carries diminishing returns, since more communication is required to apportion the workload and then collate the information in the master thread. A sweep of a $100 \times 100 \times 100$ grid takes $\sim$1 s when parallelised, compared to $\sim$5 s when not.

With these adjustments, a $100 \times 100 \times 100$ grid with no shields is solved in roughly 20 s. However, the convergence rate goes roughly with the inverse of the grid dimension and the inverse of the highest permeability (Figure 6.9). This means that the same problem including the shields, would take around 60 hours – this is infeasible if the process must be repeated many times to generate sufficient data. An alternative method to improve the convergence was therefore investigated.

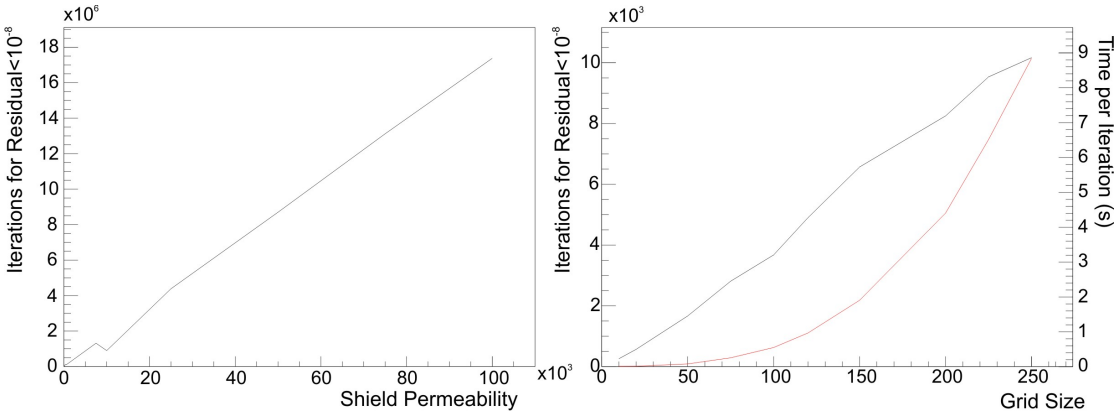

FIGURE 6.9: *Left:* The number of iterations required to reach the convergence criterion of 10^{-8} goes with the largest permeability in the problem. *Right:* The same parameter increases linearly with n, the number of grid points in each dimension (black curve), but the time taken for an iteration goes as the cube (red curve) – a fit to the curve gives $y = 5.7 \times 10^{-7} x^3$. This as expected, since there are that many more points to compute in an iteration. The total time taken to solve the problem therefore goes as n^4.

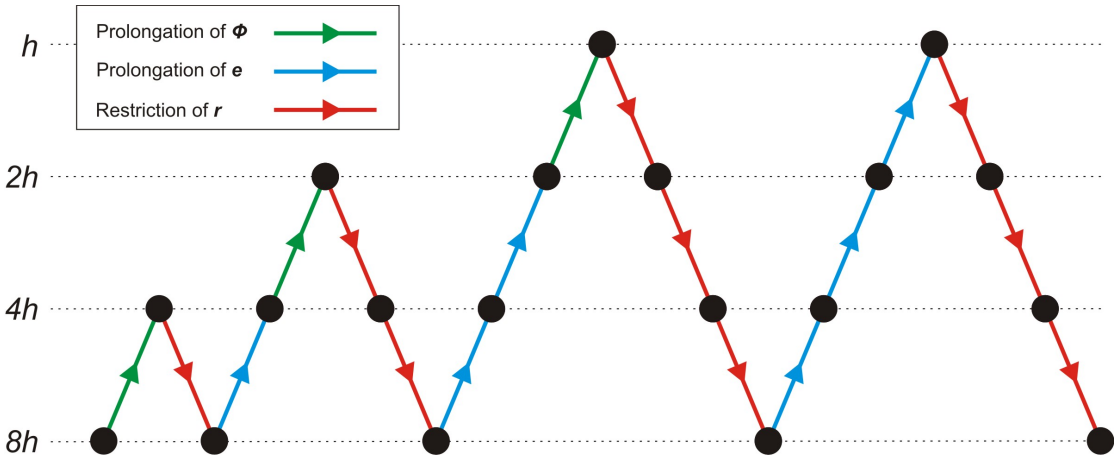

FIGURE 6.10: A diagrammatic representation of the movement of information in a Full Multigrid (FMG) method. On the coarsest grid, with spacing $8h$, the error is solved for exactly; on the finest grid, on which we desire the solution, the approximation is corrected after each V-Cycle using the propagated error. The V-Cycles can continue ad infinitum until the convergence criterion is reached, and the number of levels is arbitrary.

The first thing to note about the previous procedure is that it is a local process. By this we mean that since each point relies only on its neighbours, the correction at each step is applied on a small scale. This is effective at smoothing out the error, but slow at reducing its lower frequency modes, which may be imperceptible at that scale. An idea is to attack the problem at multiple orders, using the same

local process, such that at the highest of scales and the coarsest of grids, the low frequency error components become high frequency, merely roughness that can be efficiently corrected for. This is the premise behind the full multigrid (FMG) scheme (Figure 6.10), which is detailed in Appendix B.

It transpires that the errors can be related to the residuals, much in the same way as the potentials to the sources in Equation 6.15. The residuals from the finest grid are cast down to the coarsest grid, at each stage using these as the 'sources' to iterate on the errors. These errors are transferred back up to the finest grid to correct the potentials on a larger scale, cutting the time needed to reach convergence. This required specialised methods to transfer information between grids in cylindrical coordinates (Section B.2).

The FMG that was implemented did accelerate convergence, but was not as efficient in doing so for these problems as the SOR method. If SOR is used as the smoothing operator in the FMG, a small improvment in the computation time is seen over the SOR technique on its own. A novel idea was then to combine this new FMG with a gradual incrementation in the permeability, taking advantage of the quicker convergence rates at lower permeabilities (Appendix B.4). This protocol, a schematic of which is provided in Figure B.5, proved very effective, such that the aforementioned 60 hour problem can be completed in around 80 minutes.

6.2.3 Conversion to Fields

The FMG algorithm yields a map of potentials, which we must convert into magnetic field to be able to find the flux which crosses each sensing area. The centred difference approximation is again used to find the fields at each grid point with the relation $\vec{B} = -\nabla\Phi$, which are subsequently converted into Cartesian fields, (B_x, B_y, B_z). For each measurement, the dipoles need to be scaled by their present magnitude before they are incorporated. Figure 6.11 gives an sample output.

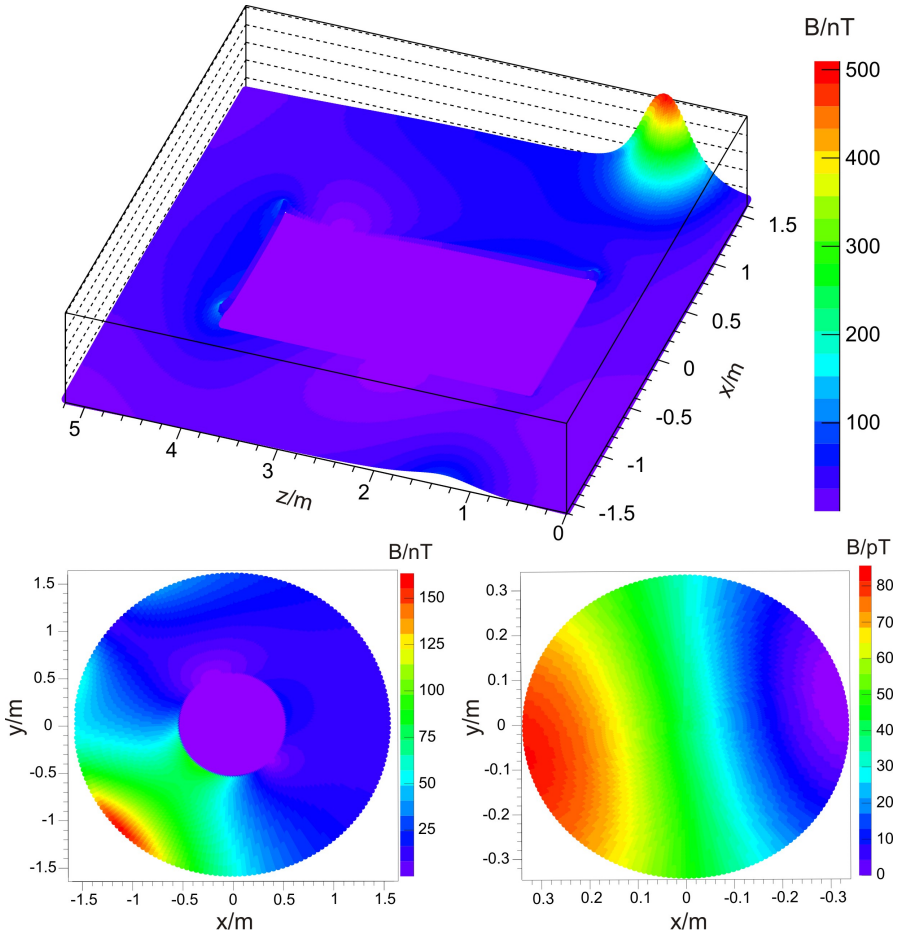

FIGURE 6.11: 2D cross-sections of the field magnitudes for a given simulated magnetic environment consisting of 30 dipoles. *Above*: a longitudinal, x-z cross-section, where there is a visible shadow for the 0.36 m diameter, 3.2 m long mu metal shields, and where the contribution of a near-field dipole source is noticeable from the positive x-direction. *Bottom left*: a transverse, x-y cross-section through the centre of the cylindrical shields, again illustrating the field suppression. *Bottom right*: a close-up of the former figure, showing the field variation internal to the shielding.

The fields at sample points on the sensing area, divided as a grid as illustrated in Figure 6.3, are then found via trilinear interpolation of the fields at the vertices of the cylindrical element encasing the sample point, as shown in Figure 6.12. For a pickup loop, the component of the field in the direction of the normal to the loop is determined and multiplied by the area of the grid element to give an estimate of the flux through it: these are summed to give the total flux through the loop, which can be converted into an average field, if calibration is required, simply by dividing by the area of the pickup loop. In the case of a neutron cell, the field

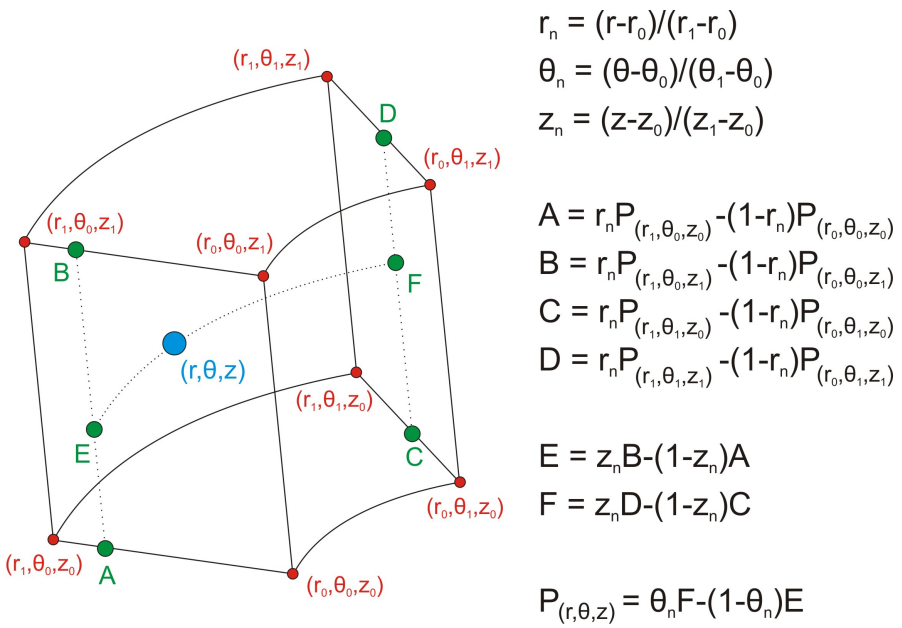

FIGURE 6.12: Trilinear interpolation to find the fields at the blue point, which represents a sample point on the sensing area, from those of the nearest grid points, in red. Distance weighted averages are used first in the radial direction, then along the height of the cylinder, yielding intermediate results, and finally in the azimuthally. Each component of the field undergoes the same treatment.

is combined with the static field generated by the solenoid, whose contribution at each point is calculated using the equation of an off-axis field from a loop (see Appendix A) for each of its 10860 coils, normalised such that the central field is $5\,\mu$T. This is averaged amongst all the points to give the volume-averaged field measurement.

Figure 6.13 shows how the relative size of the average field across a pickup loop as generated by these simulations compares to that when there are no shields. Whilst these are mainly tightly distributed about a given value, some field magnitudes, in particular those of the near field dipoles for an axial loop, cannot easily be predicted - in this instance, the separation of the dipole is of similar order to the length of the shields, so the amount of flux that can permeate the shields will be very much contingent on its position relative to the open ends of the superconducting layer. The results show that there is merit to simulating the shields, since the spread in the level of field suppression suggests that there is no one-size-fits-all shielding factor that can be applied to any dipole field.

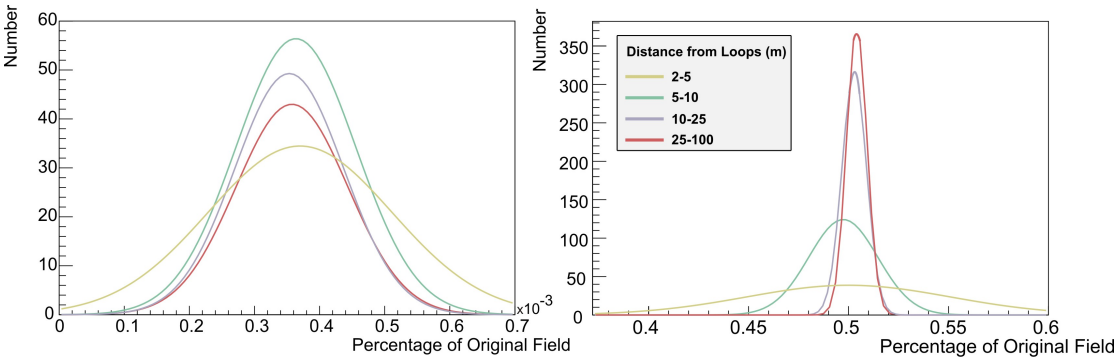

FIGURE 6.13: The values of the average fields sensed by a pickup loop as reduced by the shielding, as a percentage of the field they would have seen without them, for a transverse (left) and an axial pickup loop (right). The transverse fields have been attenuated around 1000 times greater than the axial fields, and do not demonstrate as great a variability in their suppression.

In the process of deriving these fields, the discretisation error resulting from the definition of the grid is compounded with that due to the finite sampling of the loop or cell. When there are no shields, we remarked that the flux across a pickup loop can be calculated analytically, and the error from the latter was seen to be negligible (Section 6.1.2). Further, with the convergence criterion set at 10^{-8} to minimise the algebraic error, the main source of error is likely the discretisation of the volume into a cylindrical grid and the trilinear interpolation. To make sure that this is insignificant, the Laplace Equation is first solved on the grid with no shields in place and the fluxes evaluated are compared to those of the discretised, analytical calculation. Only simulations in which the difference is less than 2% for every pickup loop are allowed to proceed to the next stage, where the shields are introduced (the first steps of the procedure in Figure B.5). This is almost always the case in the far field, but in the near field, where there are fine scale variations, only around a third of the simulations met this criterion. Having said this, there is no apparent bias in the positions of the successfully modelled dipoles, which are seen to be distributed evenly about the shields.

If this 2% error does indeed carry forward to the shielded case, if the contributions of around 30 sources are superposed, the combined error is $\approx 11\%$. With the simulated SQUID data typically corresponding to variations of millivolts, this is still

less significant than the intrinsic noise that is added to the signals (Section 6.1.2).

6.3 Conclusions

This chapter detailed the development of a model for the magnetic environment, which allows us to generate simulated magnetometry data. The main advantage over commercial software was that this could be especially tailored for this problem to provide the most efficient method for reaching a solution. The finite difference Laplace equation solver that has been developed can find the contributions of a source to the signals in a few hours, which makes it viable to execute for a number of sources. The outputs can be combined to yield a magnetic background that provides a realistic interpretation of the fields at ILL.

Future work could further improve this model. We have noted that the sources could be made to exhibit random motion, and that the superconducting shielding could benefit from a vector formulation, which could monitor its effects with time. A possible way to add a temporal dimension to these simulations would be to rework them under a Finite Difference Time Domain (FDTD) framework, instead involving the Maxwell and London Equations [154].

Whilst these additions would make the simulation more authentic, it is difficult to see what they would add in the context of providing feasible magnetic data with which to assess the performance of the scheme described in Chapter 7. As has been remarked, there is good reason to believe that the fields are dominated by low frequency components, and that changes in the field over the course of data-taking should be not greater than a μT. If the shielding properties change little with time and the environment is truly slowly varying, a realisation under a FDTD framework adds little to its interpretation.

Chapter 7

Determining the Fields in the Neutron Cells

CryoEDM required that the volume-averaged magnetic field, $\langle B \rangle$, in the neutron cells be known to a precision of 0.1 pT per measurement cycle for it to establish a limit of $d_n < 10^{-27}\, e\,\mathrm{cm}$ within the foreseen timetable of operations, as derived in Chapter 3. This chapter aims to answer the question of whether the SQUID system, with its current configuration of pickup loops detailed in Chapter 3, is capable of tracking $\langle B \rangle$ in both Ramsey cells with such a resolution, and if not, how well it can perform.

To monitor the field to an adequate resolution, the chosen magnetometer(s) must be either directly sensitive to $\langle B \rangle$, as in the case of a comagnetometer, or else, if positioned outside the cells, be used to infer the value of $\langle B \rangle$ through a combination of measurements. This is because it cannot be assumed that the field magnitude measured at a position in the vicinity of the neutron cell (as with a SQUID) will be identical to that inside it to within the required precision (see Figure 7.1). The approach taken here is to determine a mapping of the SQUID signals onto the neutron cell measurements in order to infer the value of $\langle B \rangle$ for

a specified time period. Integral to this idea are the notions of a calibration period, over which the mapping algorithm is trained, and a successive test period, to determine how accurately the algorithm can recreate $\langle B \rangle$ over time. The aim of the work in this chapter is not only to outline a suitable interpolation scheme, but to investigate possible lengths of calibration and test periods, examining for how long the residuals between the output of the algorithm and the true value of $\langle B \rangle$ remain within the 0.1 pT requirement.

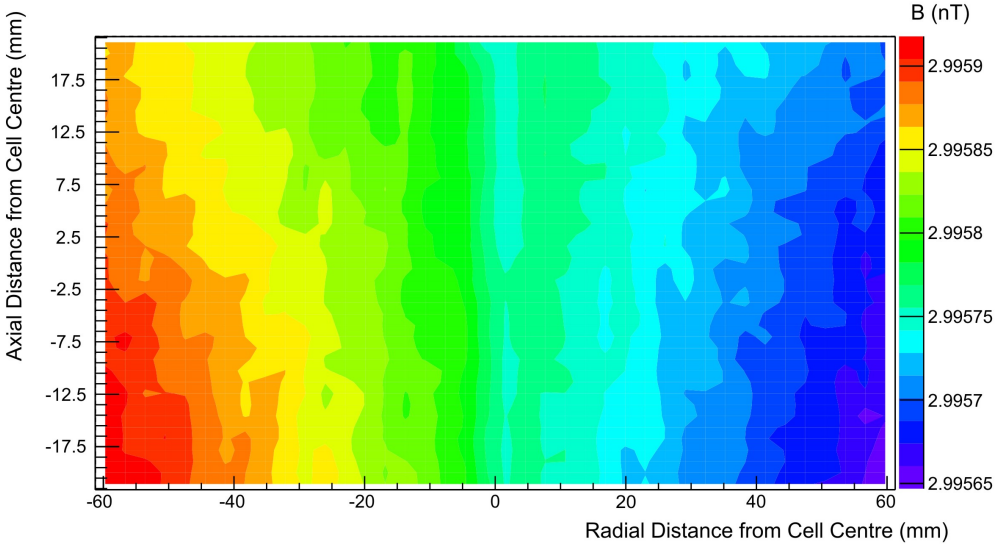

FIGURE 7.1: An example cross-section through one of the neutron cells, showing the field magnitude at each point due to external sources. The variation from one edge to its opposite is as much as 0.3 pT, which means at a position a further 10 cm from the first of the pickup loops, the field will not be within 0.1 pT of $\langle B \rangle$ in this instance. The inhomogeneity in the cell could potentially make fitting to the volume-averaged field difficult.

7.1 Introduction to Proposed Interpolation

There are two essential approaches to the problem of interpolating the $\langle B \rangle$, neutron cell signals using external magnetometers: one is to relate them spatially; the other, temporally. For the former, one can envisage, in a similar vein to the solution of fields in Chapter 6, formulating Laplace's equation for a volume encircling the neutron cells to find $\langle B \rangle$ within them. However, this is not viable with the present

arrangement of pickup loops as they cannot provide boundary conditions on the outer surface of this volume. Further, as the SQUIDs track changes in field over time, they would need to be calibrated, as described in Section 3.6.3, in order to yield these required, absolute field values.

With the loops placed to one side of the Ramsey cells, the situation presents itself as one of an extrapolation rather than an interpolation. Indeed, if the predominant variations are parallel to the central shield axis, the coaxial loops could yield a useful approximation for $\langle B \rangle$. This concept can be developed further by including contributions from the off-axis loops and those oriented perpendicular to the cell axis to make adjustments to this estimate. Thus we arrive at the basis of a framework that uses a combination of the fields registered by these sensing areas to recreate the neutron cell signals. As the experiment operates with successive Ramsey resonance cycles, the goal is to deduce such a combination to transform or map the SQUID to the resultant neutron cell measurements that is applicable throughout the duration of data-taking. This points towards exploiting any temporal relation between these signals.

First and foremost, this requires time-varying data to develop and test an interpolation algorithm on. The field simulations described in Chapter 6 are used to generate 200-hour long 'runs' of SQUID and neutron cell data, corresponding to 2400 measurement points. For each run, these are derived from a unique magnetic environment composed of 20-40 dipole sources, whose contributions to the integrated fluxes across each corresponding pickup loop and the volume-averaged fields in each neutron cell are summed and averaged for every 300 s measurement, as discussed in futher detail in Sections 6.1.1 and 6.1.2. This complexity of environment is seen to generate variations of field in these mock SQUID signals that mimic those seen in real magnetometry data in the laboratory.

A run is segmented into calibration and test data sets, with the purpose of usng the former to establish the mapping protocol to be used in the latter (see Figure 7.2).

This division is arbitrary, and a number of partition ratios are compared throughout the course of this chapter. The calibration section represents concurrent data samples from both the SQUIDs and neutron cells with no electric field present, when $\langle B \rangle$ can be determined from the neutron signals, albeit with a statistical error (see Section 6.1.3); the test period signifies data obtained with an electric field, where $\langle B \rangle$ should be reconstructed in the real experiment using the SQUID signals. Here, the simulated, 'true' $\langle B \rangle$ data allocated to the test set (that is, with no deviations due to neutron counting) are compared to the output of the interpolation to judge its performance, typically documenting the time for which the residuals or RMS differences (RMS errors) between the actual and recreated measurement points remain within 0.1 pT of each other.

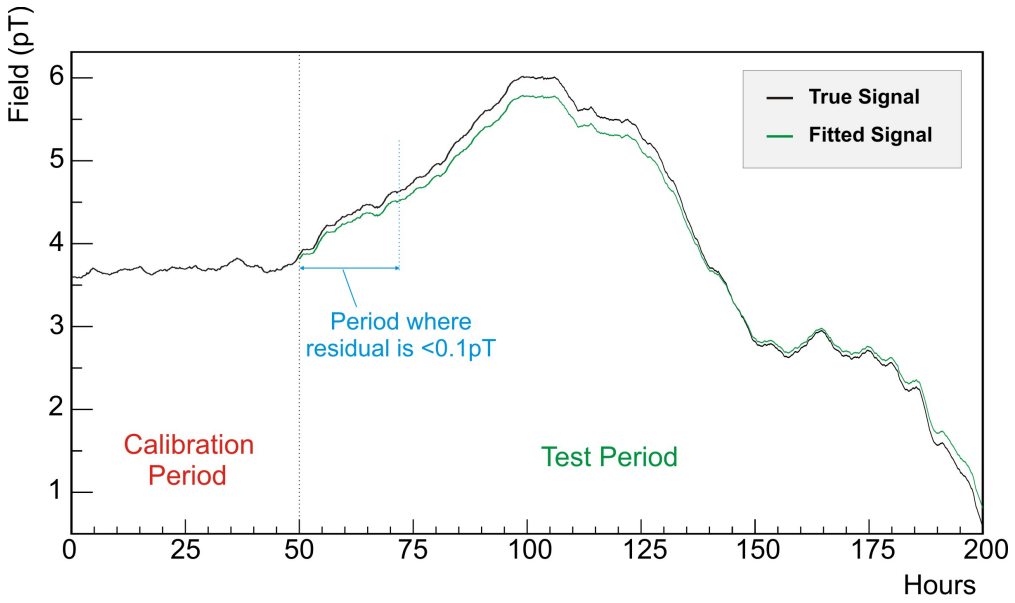

FIGURE 7.2: A 200-hour run, split into a 50-hour calibration period and 150 test period, illustrating the fitted signal for the latter. The fit is trained on the calibration period, and it is the difference betweeen it and the original signal that is examined. The performance of the fit is quantified not only through an RMS difference between it and the true signal, evaluated over a select time length, but also the duration directly after the calibration period over which the residuals remain smaller than 0.1 pT.

This $\langle B \rangle$ data represent the ideal output to which the transformation converges by tweaking its internal parameters. In this respect, with the fitting algorithm essentially calibrated using the neutron measurements, an absolute calibration

for each of the SQUID channels (that is, an output voltage to magnetic field conversion) is not strictly necessary, provided that there is no systematic drift that alters the relationship between the measured variable and the field over time. This in turn eliminates a potential source of error that would be present if exact field values from each of the SQUIDs were a requirement.

The feasibility of this technique relies on the premise that an accurate mapping can be established during calibration, and that it remains valid for at least some part of the test section. This is reasonable to expect as, whilst the magnitudes of the fields at the positions of the SQUIDs may not be the same as $\langle B \rangle$ to within fractions of a picotesla, they should be temporally correlated. Moreover, if the magnetic sources are far away and their primary contributions are slowly varying, one might anticipate that the derived relationship will remain valid for a long timespan so long as the interpolation algorithm has been properly calibrated.

In general, a SQUID signal could be expressed as $S_i = \sum_{i=1}^{n} A_i x_i(t)$, where n is the number of dipoles, $x_i(t)$ is the signal of a given dipole and A_i the corresponding transfer function between the dipole and the pickup loop (see Section 6.1.2). If the neutron cell signal, S_{NC} can be similarly written, but with different coefficients from A_i, then we might then say $S_i = \sum_{i=1}^{n} A_i' x_i(t) + S_{NC}$, where the A_i' represent new coefficients. If, for every instant in time, $\sum_{i=1}^{n} A_i' x_i(t) < 0.1\,\mathrm{pT}$, then this SQUID signal can be solely used to determine $\langle B \rangle$ under the requirements of the experiment. As has been detailed, there is enough spatial variation in the field in most simulations to suggest this would not be the case, but it is not far-fetched to consider a scenario in which a given SQUID signal would be sufficiently similar to a neutron cell signal so as to only require fine-tuning to yield the latter, given the proximity of the pickup loops to the cells. Ultimately, we require a fitted signal,

$$S_{\mathrm{Fit}} = \sum_{i=1}^{12} C_i(t) S_i + C_0 = \sum_{i=1}^{n} D_i(t) x_i(t) + S_{NC}, \tag{7.1}$$

taking a general combination with an offset, C_0, summing over all twelve SQUID signals, S_i, that are weighted by coefficients, $C_i(t)$, and where the criterion is that $\sum_{i=1}^{n} D_i(t)x_i(t) < 0.1\,\mathrm{pT}$ for the duration of interest.

If it is possible to find a mapping such that this condition holds during calibration, then it is conceivable that it will continue to be realised for a subsequent time frame under certain circumstances. This is aided by our interpretation of the ambient fields, where there are no dramatic changes due to the rapid movement of existing magnetic objects or the introduction of new sources, as was reasoned to be the case in a well-controlled environment suitable for an nEDM experiment. With some degree of constancy in the magnetic background, despite the sources being time-varying, it could be possible to determine an underlying relationship between the SQUID and neutron signals. A given transformation is therefore highly contingent on the magnetic environment, such that a mapping computed in one set of conditions will not, in general, be suitable for another. It should also be noted that it is solely the neutron signals that are of interest and not the sources themselves: whilst it would be useful, characterising the background is unnecessary, and there is simply not enough information to pinpoint the sources. Likewise, there are not enough degrees of freedom available to deduce S_{NC} exactly for each measurement, but the point was made that the fields may well be attenuated sufficiently by the shielding and that they could be similar enough at the positions of the sensing areas that any extra contribution from a fit parametrised as in Equation 7.1, $\sum_{i=1}^{n} D_i(t)x_i(t)$, can be suppressed such that it remains within $0.1\,\mathrm{pT}$ of the true value of $\langle B \rangle$ for a given set of measurements.

As we have discussed, each different field environment presents a new problem, and thus the chosen interpolation methods are tested on 200 sets of data (pairs of calibration and test periods) to see under what conditions they perform effectively and to judge the reproducibility of their outcomes. This chapter will examine two different schemes, which essentially find the coefficients $C_i(t)$ in Equation 7.1: employing a neural network approach, and in the special case where the C_i can be

modelled as constants, a multiple linear regression. The following investigations will look into not only optimising the structure of the fitting frameworks, but also the lengths of the calibration and testing periods for the fits to achieve and maintain a sufficient level of precision.

7.2 Linear Regression Method

This method is an extension of standard linear regression to encompass more independent variables, i.e. our twelve SQUID channels. The aim is to find the best fit to the neutron cell data expressed as a linear combination of these voltage signals:

$$S_{\mathrm{Fit}} = \sum_{i=1}^{12} C_i S_i + C_0 \,, \tag{7.2}$$

where C_i are the linear coefficients specifying the contribution of each of the twelve SQUID signals, S_i, and C_0 is an arbitrary offset. This is achieved by a least squares estimation, whereby the sum of squares of the residuals, R^2, is minimised with respect to each of the linear coefficients,

$$R^2 = \sum_{\text{data points}} (S_{\mathrm{NC}} - S_{\mathrm{Fit}})^2 \text{ with } \frac{\partial R^2}{\partial C_i} = 0 \,, \tag{7.3}$$

where S_{NC} is the value of the chosen neutron cell signal for a specific data point. This yields a set of simultaneous equations which can be written in matrix form as

$$\begin{pmatrix} \sum S_1 S_1 & \sum S_1 S_2 & \cdots & \sum S_1 S_{12} & \sum S_1 \\ \sum S_2 S_1 & \sum S_2 S_2 & \cdots & \sum S_2 S_{12} & \sum S_2 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ \sum S_{12} S_1 & \sum S_{12} S_2 & \cdots & \sum S_{12} S_{12} & \sum S_{12} \\ \sum S_1 & \sum S_2 & \cdots & \sum S_{12} & n \end{pmatrix} \begin{pmatrix} C_1 \\ C_2 \\ \vdots \\ C_{12} \\ C_0 \end{pmatrix} = \begin{pmatrix} \sum S_1 S_{\mathrm{NC}} \\ \sum S_2 S_{\mathrm{NC}} \\ \vdots \\ \sum S_{12} S_{\mathrm{NC}} \\ \sum S_{\mathrm{NC}} \end{pmatrix} ,$$

where n is the number of data points in the calibration data set, and the summations are found over all of them. These are solved by finding the lower upper (LU) decomposition of the matrix [155] on the left hand side, which simplifies the system of linear equations such that they can be tackled through substitution. This is preferable to finding the psuedo-inverse of the matrix, a process which is not only more computationally expensive, but more prone to inaccuracies due to the working precision of the solver.

As each SQUID measurement has magnitude O(pT), the data should be scaled accordingly to avoid such machine errors. They should also be standardised such that larger contributions from a select few channels do not dominate the calculations. This is particularly important if we wish to treat each channel equally, since there are discrepancies in the amplitudes of the signals in the different orthogonal directions. For both the SQUID and neutron data, the average baselines are subtracted and the resultant signals normalised. These operations lead to the following transformations for the coefficients, C_n:

$$S_i \to \frac{S_i - O_i}{m_i} \quad \begin{cases} \text{For } n \neq 0,\ C_n \to \frac{m_{\mathrm{NC}} C_n}{m_n} \\ C_0 \to m_{\mathrm{NC}} C_0 + O_{\mathrm{NC}} + \sum_1^n \frac{m_{\mathrm{NC}} O_n}{m_n} \end{cases} \tag{7.4}$$

with the m_i representing the scale factors associated with normalisation, and the O_i the applied offsets. These must be cast back into the original coefficients via the reverse transformations, so as the fit is compatible with the rest of the data.

Initial observations for successful computations are:

1. The addition of noise does not detract significantly from the power of the regression process. This is as expected as, unless the length of calibration were particularly short, the statistical errors on each of the signals should average out over the fitting sample.

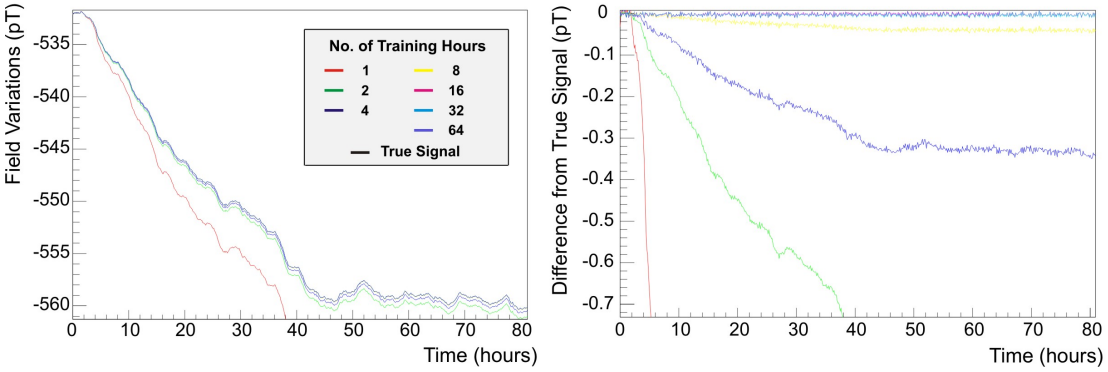

FIGURE 7.3: The outcome of linear fits on a given set of data for varying calibration periods, with the fits and the true signal illustrated on the left, and the differences between them on the right. This shows the typical pattern of a close fit during calibration that diverges afterwards. In this case, the quality of the fit gets better with a longer calibration, such that those with 16 hours or more remain within the 0.1 pT range.

2. A longer calibration period does not guarantee a better fit. This may be because there is significant drift during calibration that the relationship no longer remains linear. This means that the calibration data themselves are not as well reconstructed as if the fit were executed for smaller segments of that period: in other words, we obtain a fit that tries to generalise to a wide range of data without performing particularly well on any part of it.

3. Following from this, the majority of the fits are seen to diverge from the real signal after a given time, which is consistent with the hypothesis that the relationship between the SQUIDs and the neutron cells is a dynamic one. Despite this, some of these will remain within the 0.1 pT range for the whole 200 hours; some will drastically deviate after calibration; yet others will reconverge with the true signal due to periodic elements in the magnetic environment (Figure 7.3). Such differing behaviours are not necessarily problematic so long as they can be predicted (see Section 7.5).

Around 1 in 5 times, despite conditioning the matrix, the algorithm fails altogether due to the matrix being singular (Figure 7.4). This could be the result of a pair of

channels having variations which are sufficiently similar that there is linear dependence within the matrix. There are a number of ways that this could potentially be dealt with. The first is simply to reduce the tolerance of the solver until the matrix is no longer registered as singular: but, this could in turn exacerbate the issue, since instead of flagging the run as problematic, the algorithm will produce an outcome that will grossly exaggerate the contribution of certain channels, amplifying their accompanying noise. This happens since the near-singular entries in the matrix will in effect be the smallest, such that they will be the largest in the inverse formulation to solve for the coefficients, dominating the calculations and leading to wild misestimates.

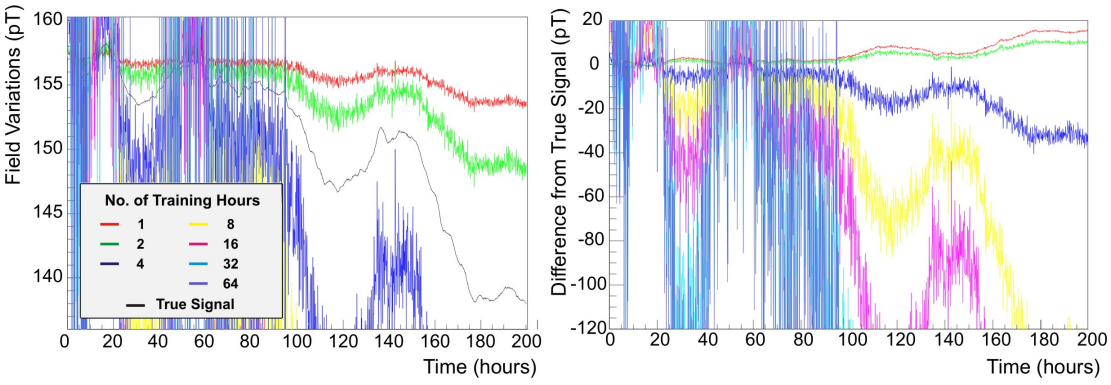

FIGURE 7.4: Linear fits for another set of data for which the matrix is noted to be singular to a given level of tolerance. Lowering this threshold such that the mapping is computed does not do the fit a great service, as some channels are over-represented and the quality of fit remains poor.

A second tactic would be to identify the problem channels and eliminate one of them to rid the matrix of any linear dependence. The likely candidates are the three double loops since, as we have argued, if the sources are far away such that there is in effect a linear gradient parallel to their common normal, they should then, being of identical sensing area and coaxial, sense very similar changes in field. We find, without exception, that it is possible to construct a non-singular matrix by removing a channel, and that the equations can then be resolved (see Figure 7.5). The quality of the fit, however, is normally noticeably poorer than in the cases where the matrix does not have to be amended.

The main disadvantage of using this techinque is the presumption that the SQUID and neutron signals are linearly related and must remain identically so for the entire duration of data-taking, since the coefficients remain rigidly defined after calibration. With evidence that this relationship does not continue to hold in several instances, we look to an alternative method to see whether it could potentially offer better performance.

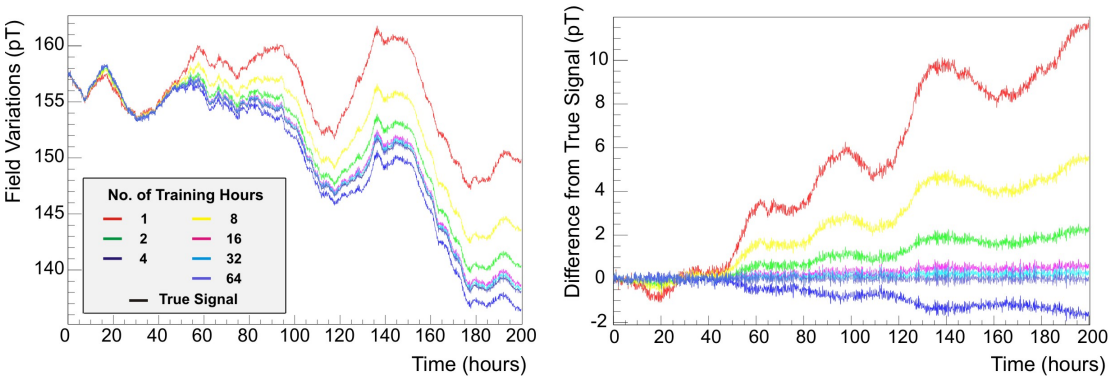

FIGURE 7.5: A linear fit using the same data as in Figure 7.4, but excluding two of the double loops. These fits, in general, show more rapid divergence from the original signal. If the field in the axial direction can be easily recreated, as we might assume if the extra double loops are not necessary, then the limitation could well be the reconstruction of fields in the perpendicular directions.

7.3 Neural Network Approach

A neural network offers a somewhat different approach from that of a least squares regression. The network can learn and generalise relationships between the input and output variables with no *a priori* assumptions about their nature: this means that, in our scenario, it is not restricted to determining a linear mapping of the SQUID signals, unlike the regression technique. This could be useful since, unless the fields are slowly varying and/or arise from far away sources, we would expect that a linear approximation would not be valid for long periods of time.

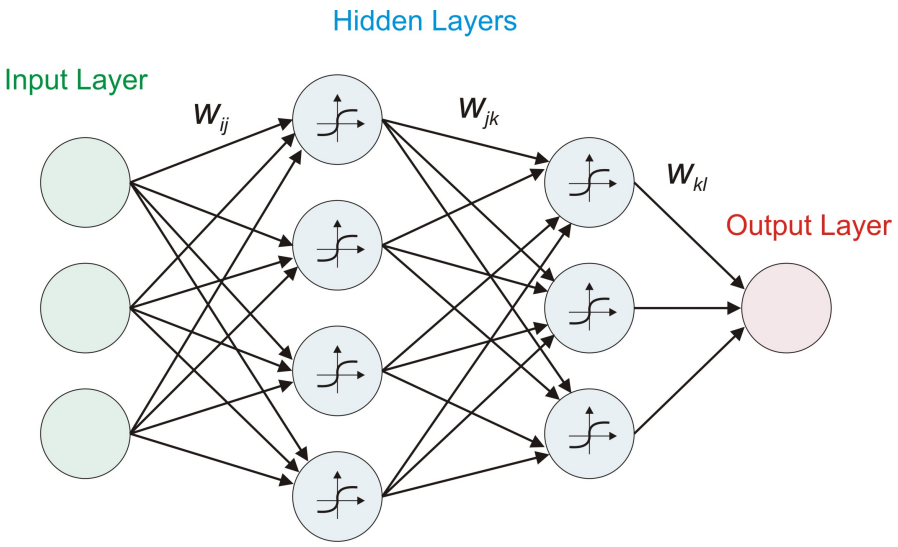

FIGURE 7.6: An example feedforward neural network architecture for a mapping of three input variables to an output. The hidden layers consist of neurons with sigmoidal activation functions, which translate the sum of the inputs to an output. The weights are adjusted when training the network to produce an output that best matches the data it is trained on.

The structure of an artificial neural network (ANN) mimics that of its namesake in a biological system, with neurons connected together by synapses that transmit signals between them (Figure 7.6) [156]. Each node computes a sum of its inputs and translates the total into an output. Both the incoming and outgoing signals are weighted, and it is the weightings that control the flow of information and ultimately the network's response to a given stimulus. The strengths of the connections are adjusted during training as the network tries to reproduce the desired outcome from the data it is presented with. An ANN is therefore an adaptive system capable of finding associations between parameters where there is no straightforward relation, a powerful tool for use in pattern recognition, classification and regression analysis.

The type of network that we invoke here is a multilayer perceptron (MLP) using ROOT's in-built TMultiLayerPerceptron class [157]. An MLP is a feedforward network in which information is transferred unidirectionally through a series of neurons organised in different stages. This comprises the input layer, with nodes corresponding to the independent variables, an output layer for the dependent

variable(s) and a given number of hidden layers which handle the data in between: without them, connections run directly from the inputs to the output, and the resultant transformation is tantamount to linear or logistic regression. It is the hidden layers which enable the network to represent non-linear relationships, if they suitably process the signals. Each neuron in the layer is given a sigmoidal activation function which configures the response for a given sum of inputs.

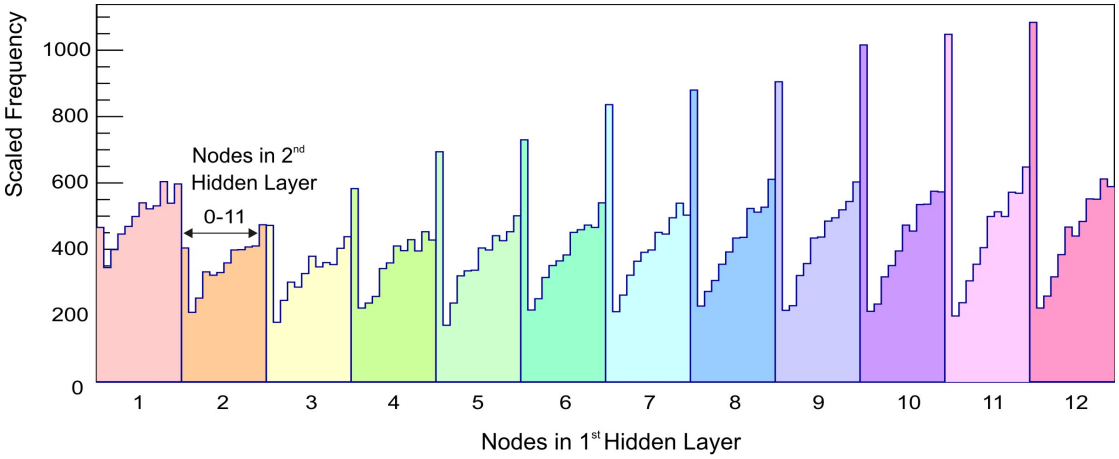

FIGURE 7.7: A comparison of different neural network architectures, based on the inverses of their RMS errors for data after a 16-hour calibration period - those with smaller RMS errors are therefore greater represented by the scaled frequencies. This suggests that networks with a single hidden layer generally perform better than ones with two, and that the more neurons in this layer, the better the network can generalise.

This is motivated by the observation that, in theory, any continuous function can be expressed as a superposition of sigmoid functions [158]. Under the Universal Approximation Theorem, an MLP should not need more than one hidden layer to reproduce such a function to any degree of accuracy [159]. However, whilst this posits that such a network exists, it does not offer insight into how it is constructed. Complex problems are sometimes better served with a network with further hidden layers that can more easily represent a sophisticated mapping between the variables. On the other hand, too many parameters and the network will overlearn the data it is trained on, with the superposed noise, and be unable to spot an overarching pattern that will allow it to generalise to later data.

It follows that a primary concern in the design of an ANN is its structure. To begin to narrow down the possibilities, we investigate the use of up to two hidden layers with 1-12 nodes in the first hidden layer and up to 11 in the second. This is in accordance with a rule of thumb that, in general, the number of neurons in the hidden layers does not need to exceed the total number of neurons in the output and input layers [160]. Figure 7.7 shows the results for 16 hours of calibration data, comparing these networks by their RMS errors for ten successive 12 hour periods after 10^4 iterations of training. The best network architecture will be the one which consistently produces networks with the lowest RMS errors over a broad spectrum of time frames. For each set of data, the networks are weighted inversely with the given RMS, these values summed together for all possible network configurations, and each expressed as a percentage of this total. The fraction for that run is scaled by an arbitrary factor, which is the same for all runs, and the value totalled over hundreds of data sets, giving a scaled frequency for that configuration. Figure 7.7 illustrates clearly that a second hidden layer adds unnecessary abstraction that does not generally improve the network performance. It should be noted that, whilst Figure 7.7 combines the data from all ten 12-hour sections, the same upwards trend is seen with neuron number regardless of which of these periods we focus on. This suggests that more neurons in a single layer appear beneficial, although it cannot be told whether more than 12 neurons will prove useful from this figure.

To test this, the same technique was extended to examine networks with 1-14 neurons in a single hidden layer, but now looking at a range of calibration lengths and training the networks for 10^5 iterations. This confirms that the networks with more neurons tend to perform better in general (Figure 7.8). However, there is an observed increase in the performance of networks with fewer neurons as the calibration period is lengthened. Since the sample size is two hundred entries, it is difficult to say whether this is statistically significant, and indeed, a comparison between networks with 4 and 13 neurons in their hidden layers respectively suggests

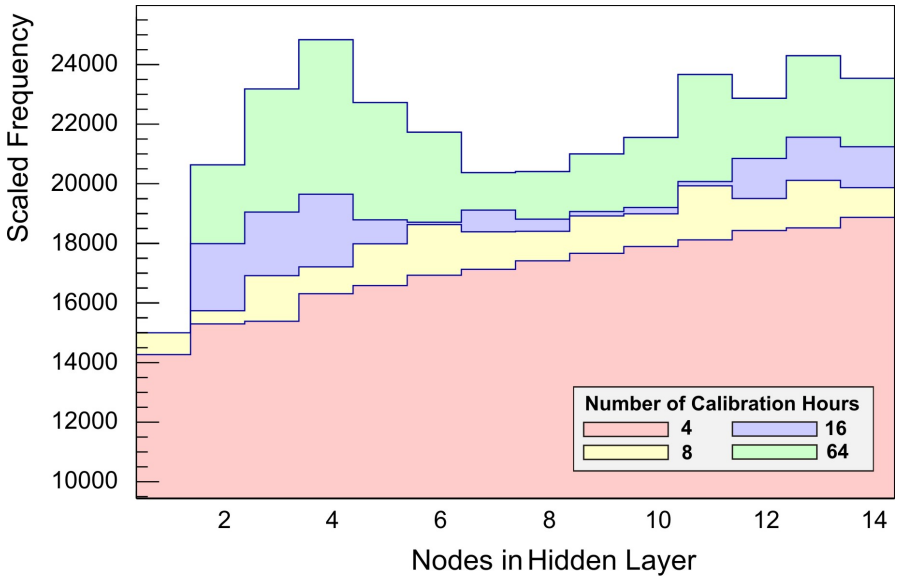

FIGURE 7.8: A continuation of the line of investigation in Figure 7.7, comparing the performances of networks with a single hidden layer with up to 14 neurons for a range of calibration periods. It should be noted that the graphs have been offset as they were overlaid so that all useful information is visible, thus there is no meaning in the differences between the heights of entries for the different calibration lengths. The network with 13 neurons is one of the best performers in every instance, and is the one which will be chosen to be further developed.

that the latter fares better across a range of calibration periods (Table 7.1). It is apparent, however, that none of the networks consistently demonstrate the level of performance required of the interpolation scheme. Whilst we will proceed with the network with 13 hidden layer neurons for the rest of this analysis, we must now look to see how its performance can be improved.

Calibration Time (Hrs)	Hidden Layer Neurons	Av. RMS	Max. RMS	Min. RMS	Networks with RMS <1 pT
4	4	6.2×10^{-10}	2.4×10^{-8}	1.2×10^{-14}	61
	13	5.5×10^{-10}	2.0×10^{-8}	1.5×10^{-14}	64
64	4	8.7×10^{-10}	2.7×10^{-9}	4.8×10^{-15}	118
	13	6.1×10^{-11}	1.8×10^{-9}	5.2×10^{-15}	127

TABLE 7.1: Comparison of figures of merit for the RMS errors for networks with 4 and 13 neurons in their hidden layer respectively, for calibration periods of 4 and 64 hours in length. The values are drawn from a sample of two hundred networks for each configuration, and the better figure is highlighted for each parameter.

7.3.1 Refining the Neural Network

To assess the different configurations, we aim to compare networks that are trained close to their full potential. However, there are a number of subtleties involved in the training process that can make judging this complicated. Training involves adjustment of the weights in the network to reduce the value of the performance function for the training data set, and thus ultimately improve the quality of the fit that it can provide. The performance function can be defined in several ways, but for a regression problem, typically takes the form

$$F = \frac{1}{n}\sum_{i=1}^{n}(t_i - s_i)^2, \tag{7.5}$$

i.e. a mean square sum of the residuals, where n is the total number of training data points, t_i is the desired output for one of these points and s_i the actual output. The most common form of algorithm used to optimise the network with respect to the performance function is that of backpropagation [161]. The value of the performance function is transferred back through the hidden layers with reverse computations and the weights are tweaked in direct proportion to the error in the associated outgoing neuron, thus decreasing its size. The TMultiLayerPerceptron class offers six different learning methods and of these, the Polak-Ribiere formula [162] was found to be the most effective in our case, in that it reduces the error the quickest. This is a variant of the technique of steepest gradient descent, in which the connection strengths are updated in the direction of greatest negative change on the error surface defined as a function of the weights (see Figure 7.9). With the weights initially randomised and no prior information about the geometry of this surface, it is more than likely that the method converges upon a local minimum, rather than the optimal point. This makes it imperative that we train multiple networks of the same topology and choose the one which performs the best. This approach is also driven by the fact that a lower error on the training set does not necessarily correspond to a lower error on the rest of the data, as now discussed.

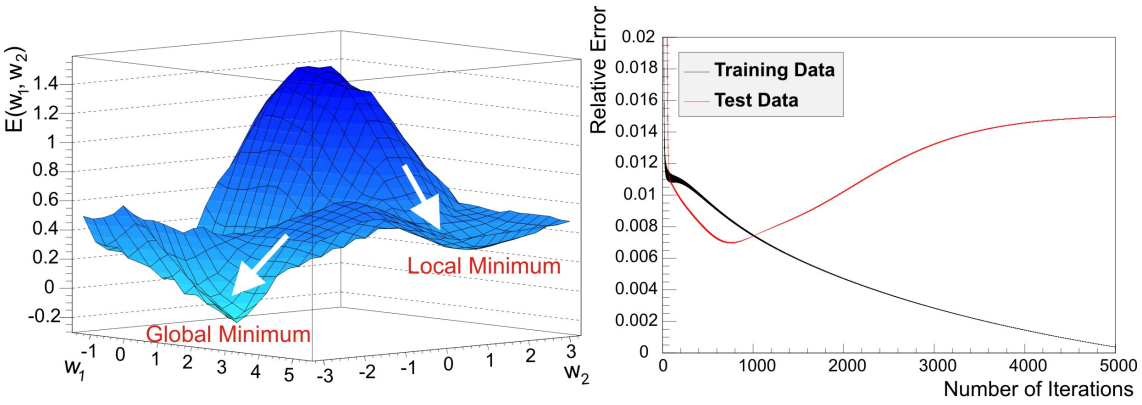

FIGURE 7.9: *Left:* an example error surface for a neural network with two adjustable weights. With no knowledge of its geometry and with the weights initally randomly assigned, there is no telling whether training is converging on a local minimum on the surface. This illustrates that by training many networks and thus choosing many starting points on the surface, one has a better chance of finding a network with a lower error. This is even more so when moving to 12 dimensions, with a much greater parameter space to explore. *Right:* unless both sets of data are truly representative, there is no guarantee that the error surfaces for the training and test data will coincide, and thus the error on the test set could start to diverge whilst the training error continues to improve.

The test data is used as a separate gauge to assess how well the network will generalise and fit to any subsequent data. In an ideal world, as the network is trained, the error on the test data would decrease, and this would correspond to smaller residuals in later data. There are two major issues that can mean that this is not the case in reality: the first is the problem of overfitting; the second is the likelihood that the test data may not be representative of later measurements - these are both consequences of a potentially small calibration data set of at most a few hundred data points, which is localised at the beginning of a run of data. The scenario is therefore unlike a standard regression, and there is always the danger that the network will be too specialised to earlier data and fail to capture the underlying temporal variations.

Overfitting occurs when the network begins to memorise the outputs of the training data and does not generalise, so that the test error eventually starts to climb [163]. In tackling this problem, we might ask the question of how many iterations of training is optimal. To attempt to answer this, several hundred networks were

trained on different simulated data with 4, 16 and 64 hours of calibration, for two million iterations each. Evidence of overfitting is sometimes seen, but on the whole, the test error is seen to keep decreasing for the entire period. A suspicion is that the test error is too similar to the training data, and what we are then seeing is overfitting to the calibration data as a whole. However, what we find when we plot the RMS error of 12-hour segments of data after calibration is that there is genuinely still room for improvement in the vast majority of cases, as these errors are also seen to decrease with further iterations. In fact, the general trend of these errors follows that of the test error: if the test error continues to smoothly decrease, so do the errors on the rest of the data; if the test error shows significant features, they also manifest, albeit slightly differently, in the other errors. This is borne out by a direct comparison of the fitted data to the true neutron signal, where we see in the former case, the fit edging closer to the signal whereas in the latter, any increases in the error correspond to a worse fit and herald overfitting (see Figure 7.10).

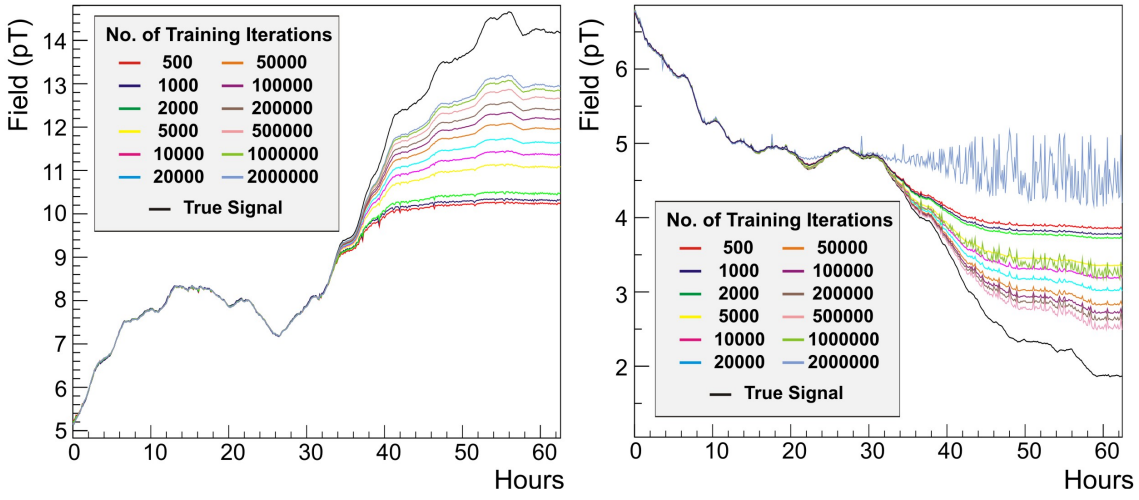

FIGURE 7.10: *Left:* The network demonstrates improved performace with further iterations of training, as it converges towards the true signal. *Right:* Here, the network provides a better fit with more iterations, but only up to a point. It begins to overfit at greater than 5×10^5 iterations, where the fit not only becomes poorer but is dominated by noise on the input channels.

As the test error is the only diagnostic we have available with which to predict the

network's eventual performance, it is a good sign that there is an observed connection between it and errors on subsequent data. However, there are a number of caveats in its interpretation. In none of the instances examined did the error in any period of data increase without being accompanied by a similar departure from a steady decrease in the test error. These features, however, do not necessarily coincide exactly, making combatting overfitting difficult. An obvious solution, terminating the training at the point where the test error begins to increase, a technique sometimes referred to as early stopping [164], may therefore overshoot and not provide the lowest error on the rest of the data amongst the sequence of iterations. Further, any blips in the test error will cause the training to be stopped in advance of the minimum point. To overcome both of these, early stopping is implemented comparing the error in larger intervals of a thousand iterations or more, with an increase being a clear indicator of worsening performance. Networks which are not well-behaved, aside from being more difficult to interpret, are unlikely to recover their performance.

This brings us to the question of how to find the best network for the given data. Clearly, whilst the observation that the test and RMS errors can continue to decrease for hundreds of thousands of iterations is encouraging, given that this can be equivalent to over a day's worth of training, and that the gains become increasingly smaller as the curve stagnates, it is not the most effective way of achieving the smallest error. It is better to train multiple networks on the same data and to pick the network which shows the most promise, rather than focussing efforts on one network. If we return to the picture of the error surface, but this time defined for a set period after calibration, this will be the network whose starting weights are most favourable for reaching the global minimum or at least, with a finite number of networks, the lowest local minimum that is accessed. If we are to judge this using solely the test error, we must establish whether there is a connection between the sizes of the test error and the RMS error from the desired period for multiple networks trained on the same data.

This calls into question how best to divide up the calibration period into test and training data sets. In the simulations thus far, every third point has been allocated to the test set, and the remainder to the training: the idea was to have enough data in the test set that is spread out across the calibration time from which to draw conclusions, so that the lower frequency drifts are well represented. Figure 7.11 and Table 7.2 give relevant results from the comparison of other possible partitions of the data - splitting in a 3:1 ratio rather than 2:1, looking not only at spreading the test points evenly within the calibration data, but also distributing them more heavily towards the end of the calibration period, since these may bear more resemblance to later data. What we see is that these weighted data distributions are better correlated with the RMS errors, particularly that of the 12-hour period immediately following calibration, but at the expense of worse values for these errors. These are not necessarily conflicting observations: the weighted partitions will have a similar test error to the RMS error over the next 12 hours, since the test data occur just before this duration, but as the training data are at the beginning of the calibration, the network does not generalise as well and these errors are both higher than their counterparts using the even distribution.

	Average RMS Error for Period (After Calibration)		
Partition Configuration	0-12 Hours	12-24 Hours	0-120 Hours
1	7.10×10^{-12}	2.71×10^{-11}	2.72×10^{-10}
2	4.37×10^{-11}	8.57×10^{-11}	4.04×10^{-10}
3	1.11×10^{-11}	3.40×10^{-11}	2.88×10^{-10}
4	7.04×10^{-12}	2.69×10^{-11}	2.70×10^{-10}
5	2.93×10^{-11}	6.23×10^{-11}	3.60×10^{-10}
6	1.21×10^{-11}	3.53×10^{-11}	2.91×10^{-10}
7	9.69×10^{-12}	3.11×10^{-11}	2.80×10^{-10}
8	1.02×10^{-11}	3.19×10^{-11}	2.83×10^{-10}

TABLE 7.2: Comparison of the average RMS errors for networks whose test error correlations were illustrated in Figure 7.11. The highlighted rows are the best and second best performing configurations respectively, which correspond to the 2:1 and 3:1 partitions where the test data are evenly distributed amongst the calibration data.

Merging the two to capitalise on both their advantages, Figure 7.11 illustrates a third possible avenue splitting the calibration data into three sections - a training,

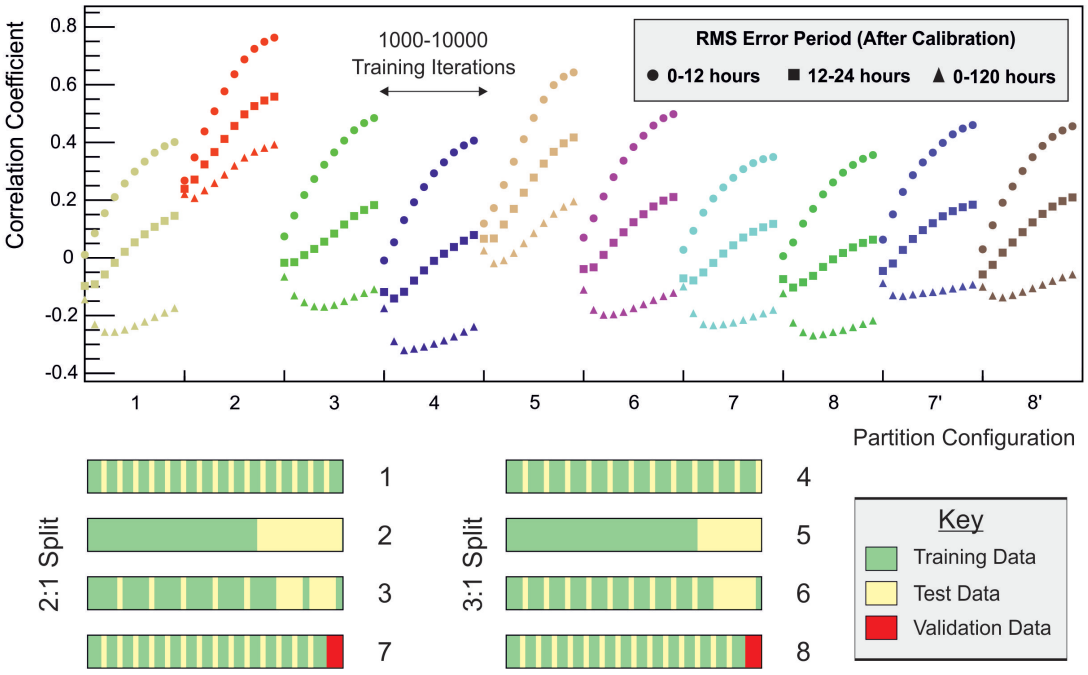

FIGURE 7.11: Figure showing the average correlation (for 150 runs with a calibration period of 16 hours) between the test error and the RMS error for three different periods of data after calibration and for eight possible partitions of the calibration data, illustrated below (plus the validation period for configurations 7 and 8, signified by the dash). The strongest correlations are seen for the fully weighted calibration partitions, where all the test data are distributed towards the end of the calibration. In some cases, a negative correlation is seen for later data, or over a long stretch of time. This could indicate that a pattern which describes the calibration data and any which closely follow does not extend as well to later data, pointing again to a lack of generalisation of the networks.

test and validation set. The test data are allocated evenly within the calibration period save the last hour, which is reserved for validation. The error on the test data is monitored as before, with early stopping implemented, but it is the performance on the validation set that is used to select the prime network. In this way, the network should be able to generalise well whilst we should see a good correlation between the error on the validation set. However, Figure 7.11 and Table 7.2 show that this method tries to accomplish too much and we in fact make sacrifices on both capabilities: on the one hand, as would be expected, the RMS error values fare worse since the network is trained on fewer data, but also, the correlation between the validation and the RMS errors is barely better than

with no validation at all (comparing 1 and 7′ or 2 and 8′ in Figure 7.11). This is simply due to it being too short a period, but lengthening it will constrain the data available for training and testing and impact on the network's performance.

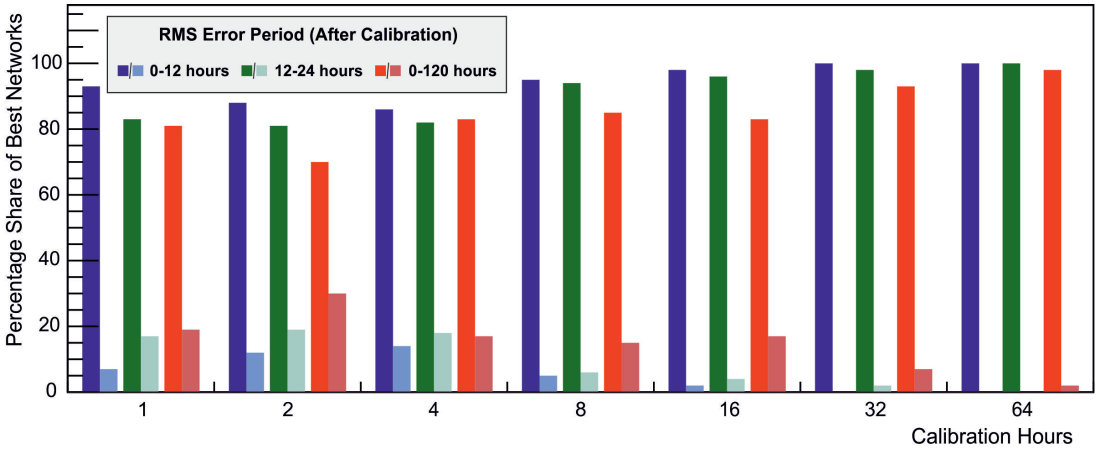

FIGURE 7.12: A comparison of the even and fully weighted distribution (1 and 2 respectively on Figure 7.11), with the lighter colours corresponding to the latter. For every run, we assess which of the two configurations is better by comparing the RMS error of their respective networks over a given period after calibration: the better network is added to the histogram. The figure shows conclusively that, despite the better correlations between the test error and the RMS errors that should lead to a more meaningful choice of best network for the weighted paritioning, the even distribution simply wins out by providing a more representative set of test data with which to assess the network's performance.

So we are left with the dilemma of which is the better option: one where we can better predict the outcome of the network, or one which is guaranteed to give a reasonable standard of performance. Using the 2:1 split, we train 200 networks for a given run and pick the one with the lowest test error after a certain number of iterations (10000 for the weighted and 3000 for the even partitions). This time, the best networks continue to be trained up to 10^5 iterations, or until they are terminated by early stopping. The even and fully weighted distributions, corresponding to partitions 1 and 2 in Figure 7.11 respectively, are compared by their RMS error for the same three periods as in Figure 7.11. The number of times one configuration wins out over the other is tallied and converted to a percentage out of the total of 200 runs. Figure 7.12 illustrates these fractions, and shows that networks with an even distribution better those with a fully weighted splitting

by a considerable margin for a broad range of training periods and durations of interest after calibration, even having been initially selected with fewer training iterations.

In summary, we therefore adopt an even distribution of test data within the calibration period, as illustrated by 1 in Figure 7.11, and elect to use a neural network with 13 neurons in its hidden layer to compare and contrast its performance with the linear regression method detailed in Section 7.2.

7.4 Comparison of Neural Network and Linear Regression Methods

Two techniques have been examined which fit a combination of the SQUID signals to the neutron cell measurements, and they have been refined for this purpose. What remains is to compare them and consider under what circumstances they may succesfully be used.

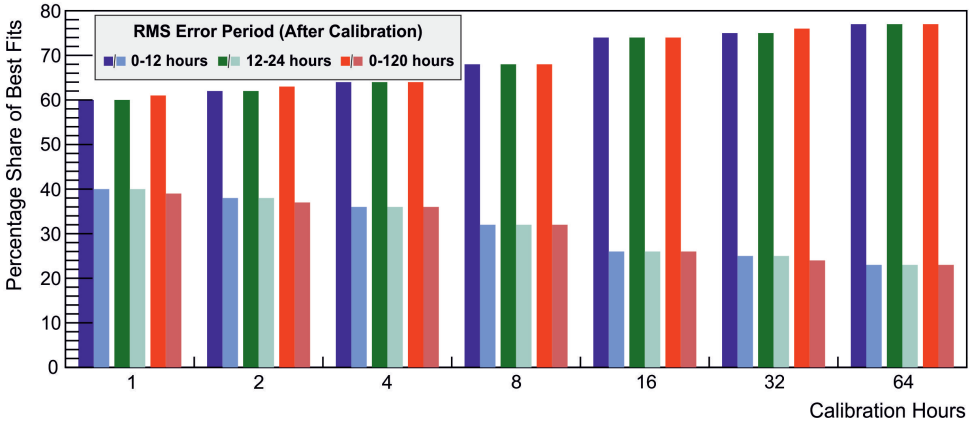

FIGURE 7.13: Much as in Figure 7.12, we compare the neural network and linear fits by their RMS errors and tally the best method in each instance. The darker shades represent the linear fit and the lighter, the neural network. The linear fit performs better in most cases and much more so the greater the number of calibration hours used.

In much the same vein as with the comparison of different data partitions in Figure 7.12, we begin by straight off comparing which of the two methods performs better for varying calibration lengths at different stages of data-taking. In Figure 7.13 it is observed that the linear fit outdoes the neural network fit in the majority of cases, with this majority significantly rising as the number of calibration hours is increased. As the quality of both fitting techniques was previously observed to generally get better with more calibration hours, this suggests that the rate of improvement of the linear transformation with further input data is greater than that of the neural network. This is somewhat surprising since, as we have remarked, the coefficients of the linear regression are static whereas the neural network has the ability to adapt. What this appears to tell us is that the linear fit has the greater ability to generalise when trained on more data: whilst the input data are no longer as localised, which dilutes the quality of the fit for these earlier sections, the exposure to more data enables the fit to better reconstruct following sections of the neutron signal. This is of course helped by the fact that the relationship between the SQUIDs and neutron signals in some of these simulations looks to be quasi-linear such that a simple linear fit may suffice, but this hinges on any departures from linearity being small enough that the fit still conforms to the specified 0.1 pT resolution. Further, Figure 7.13 should not be the be-all and end-all with regards the comparison between the linear and network mappings, and it must not be overlooked that, even in the least favourable of circumstances, the neural network is the better choice in almost 1 in 4 instances.

The next step in breaking down this analysis is to look at the percentage of runs where the fits meet the 0.1 pT requirement, since this is the ultimate goal of this exercise. Figure 7.14 illustrates this for a variety of calibration periods and junctures during data-taking. On the one hand, we see that for the linear fit, this percentage climbs for all points examined during data-taking with further calibration hours. After 16 hours, the mapping in any successful iteration is well-established such that there is no degradation in performance with time, in accordance with our

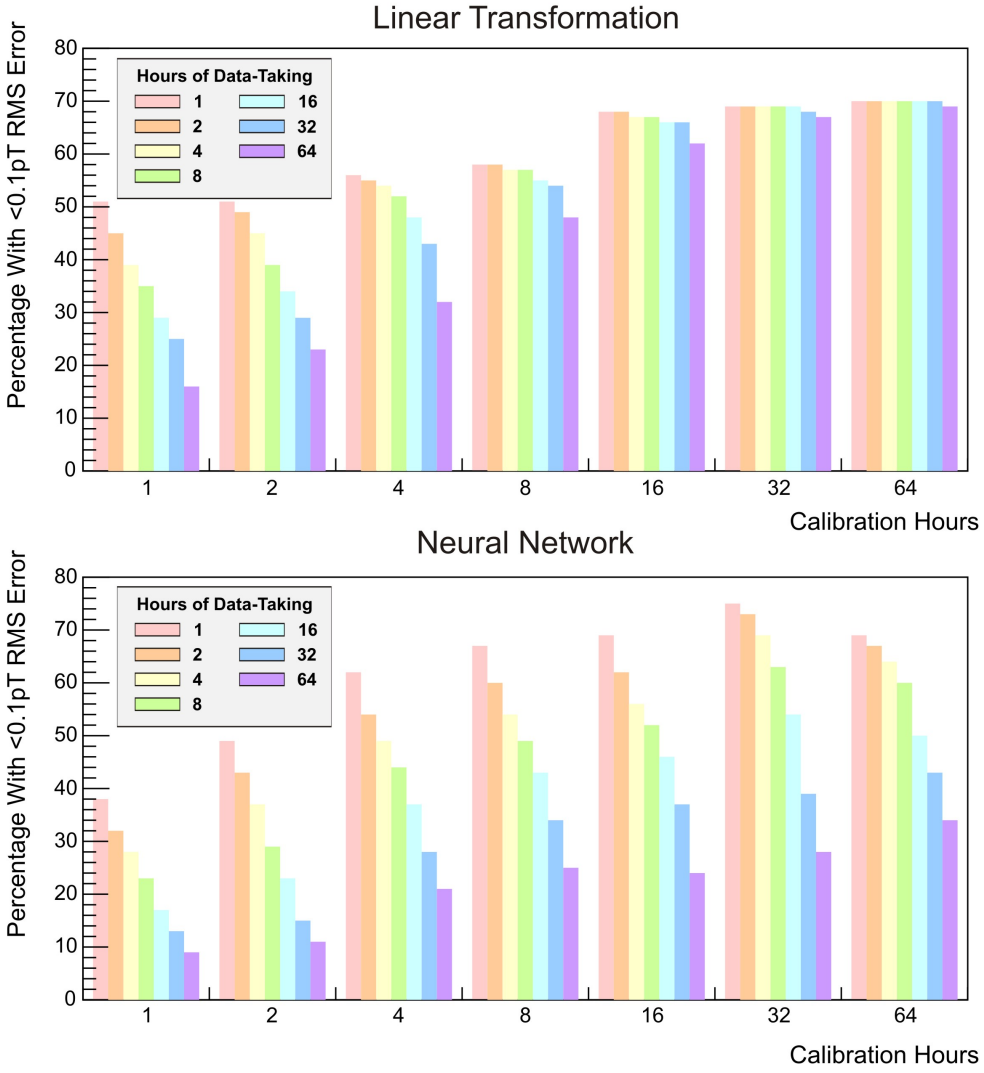

FIGURE 7.14: A graph showing the percentage of runs which produce fits with an RMS error lower than 0.1 pT for both methods, for different calibration lengths and at different stages after calibration.

previous statement that the relationship in these scenarios is indeed close to linear. This is promising, but equally, suggests a length of calibration period which may be unrealistic, if we are to be sure that the fit is performing well enough. Moreover, almost a third of these simulations do not fall under this category and therefore must be identified.

On the other hand, whilst we cannot make such strong claims for the neural network, we see that it outperforms the linear transformation in the immediate aftermath of the calibration period on a number of occasions. Although the neural

network is inherently less predictable than the linear technique, this is a remark that may prove useful should the focus be on a much shorter timeframe for both calibration and observation.

		No. of Days of Calibration			
		1	2	3	4
Average RMS Difference of Fit for Day (pT)	1	8.69	20.45 (+135%)	20.36 (+134%)	25.71 (+195%)
	2	10.86	17.62 (+62%)	15.21 (+40%)	13.94 (+28%)
	3	7.30	42.40 (+480%)	45.21 (+519%)	56.85 (+678%)
	4	20.43	30.93 (+51%)	28.20 (+38%)	21.15 (+4%)

TABLE 7.3: The RMS differences of the neural network fits from the true signals for each of four consecutive days of simulated data, after further calibration for a further 8 hours on each of those days, averaged over 100 runs. This is accompanied by the average percentage increase in the RMS compared to the initial calibration on day 1.

With this observation in mind, another protocol was proposed involving further calibration of the neural network at set intervals. Keeping the weights from the previous calibration, they can be adjusted through training on later data such that the network adapts dynamically to the signals with time. This is in contrast to the linear regression method, where this concept holds no meaning since this is tantamount to reevaluating the coefficients from scratch. The data-taking period is divided into days, of which the first third is used for calibration.

In many cases, the fits become worse as more data are added (Table 7.3). This suggests that the error surfaces are sufficiently dissimilar between the different calibration data sets that carrying forward the weights from one to the next does not help in reaching a minimum in the latter, the results being no different than if the weights had been randomised. This implies that we are better off calibrating afresh, again picking the best network from its test error, and corroborates the idea that the chosen neural network architecture inherently cannot identify the underlying pattern connecting the SQUID signals to the neutron cells.

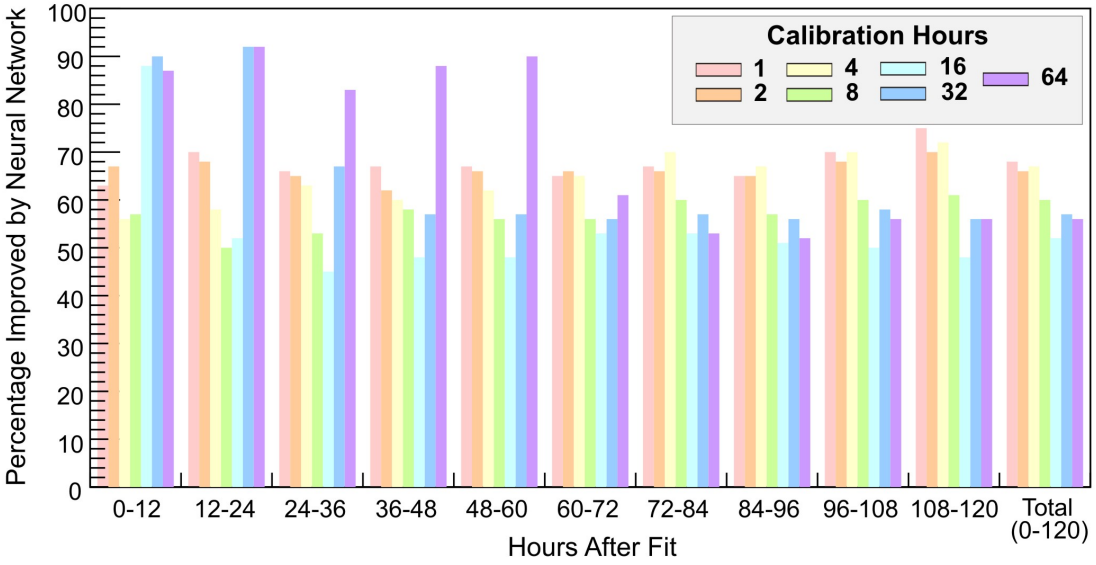

FIGURE 7.15: A graph showing the percentage of runs which produce fits with an RMS error lower than 0.1 pT for both methods, for different calibration lengths and at different stages after calibration.

As a last investigation, the neural network was trained on the differences between the linear fit and the true neutron signal to try and improve its performance. This removes a substantial part of the relation between the signals, allowing the network to focus on the non-linearities. Figure 7.15 compares the original fit with that augmented by the neural network. The results mirror those in Figure 7.13, with a network which is trained for longer demonstrating a greater likelihood of improving the corresponding linear fit in the short term, compared to a fit with a shorter calibration period. But there is little difference in performance in the long term, greater than 60 hours after calibration. Without a guarantee that the network will ameliorate the fit, it becomes a liability to include as standard as an addition to the linear regression.

What these analyses have shown us is that the linear regression technique is the more consistent performer, compared to the neural network, and offers a simpler interpretation. This is a definite advantage since ultimately we need to be able to predict whether the fit will meet the 0.1 pT criterion under a given set of conditions.

7.5 Characterising the Fit

In the previous section, linear regression was shown to produce a fit with a precision within the requirement around two-thirds of the time. This success rate is not on par with the statistics required of a high precision experiment, and therefore we must determine the circumstances under which the fit will perform sufficiently. Although there will likely be correlations with a number of variables, for instance, the field gradients across the neutron cells, these are not accessible in a true experiment, and therefore of little use in the context of devising a protocol for use with real data from CryoEDM. The discussions here have therefore been limited to parameters associated with the SQUID signals since these are, in principle, the only information we have to work with.

To start, we look for correlations between the size (by which we mean the RMS average) or range of any of the signals and the RMS errors up to 120 hours after the calibration, the former over the entire data-taking period. Clearly, for a given run, scaling all of the signals by the same amount would similarly scale the RMS error. For a given run, therefore, there will be a threshold size of a given signal below which the fitting technique will always achieve a 0.1 pT precision. One might consider that this principle could extend to all runs, such that there would be a connection between the signal size and the fit outcome. The correlations between these suggest otherwise – defining R as the Pearson's correlation coefficient, ± 1 signifying a perfect postive/negative correlation, most are at best weakly correlated ($R < 0.1$), if at all, and the strongest is seen for the ranges of the neutron signals ($R \sim 0.4$). What this implies is that there is no one SQUID contribution that is always dominant, and that the variability of the neutron signals is the most important factor.

Since we are fitting to the neutron cell signals, the range is only a viable parameter over the calibration period: using the ranges over 64 hours of calibration, the correlation coefficients fall ($R \sim 0.3$). If we restrict the analysis to just near-field

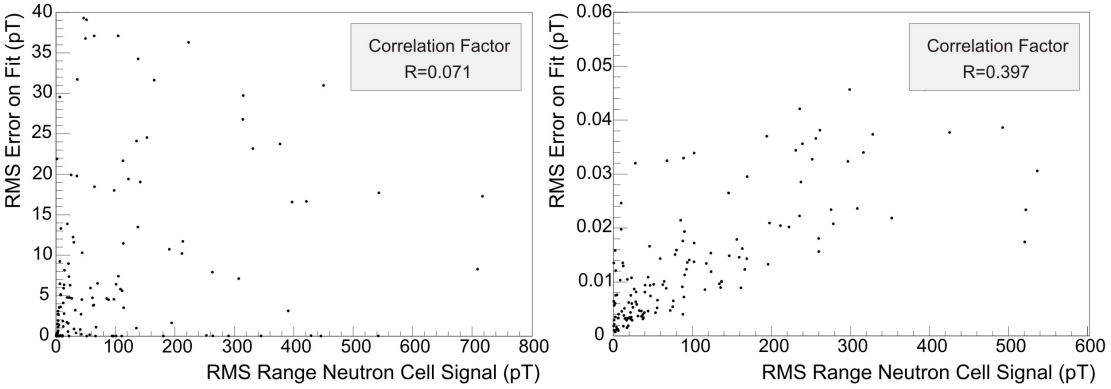

FIGURE 7.16: RMS difference of the fit over 120 hours of fitting after calibration for a 64-hour calibration period, plotted against the range of corresponding neutron cell signal over that period, for the far-field (left) and near-field cases: in the latter, we illustrate a subset of data, which account for around 60% of all the data. The far-field data are spread out over a bigger scale, and are more variable. However, there is a caveat since the signals in both cases were scaled to a similar level, giving a field magnitude of around 0.1 μT at the position of a fluxgate near the mu metal shields. In reality, we would expect that the far-field sources, being at distance, would result in a smaller signal at the positions of the pickup loops, and therefore this difference in scales between the two may be reconciled, even if that in the spread of the data cannot.

or far-field sources (between $2-5$ m and greater than 100 m respectively), the correlation coefficient is apparently seen to increase. However, on closer inspection, this is seen to be highly affected by a handful of outliers which skew the distribution, and excluding such points lowers the coefficient back to the former level. In the near-field there is a nucleation of points around much smaller values of both parameters, that are consistent with the 0.1 pT limit, whereas in the far-field the scales are much greater (see Figure 7.16).

To go further, if we were to reproduce the last column of Figure 7.14 for both of these regimes, i.e. look at the percentage of fits which have a RMS error less than 0.1 pT 64 hours after a 64-hour calibration period, one finds that 62% of those with dipoles solely in the far-field meet the criterion compared to 77% in the near-field and 69% with a mixture. This is in accordance with the spread of points seen in Figure 7.16, where it is apparent that few are below the 0.1 pT mark. The lower percentage for the far-field cases could be as a result of them generating

a comparatively spatially homogenous field pattern in the vicinity of the pickup loops. As we saw in Section 7.2, the signals could then be similar enough so as to cause singularities in the matrix associated with the linear transformation, that were seen to negatively affect the quality of the fit. This is useful information, but it must be reiterated that isolating runs in the two regimes would not be possible unless the environment were characterised.

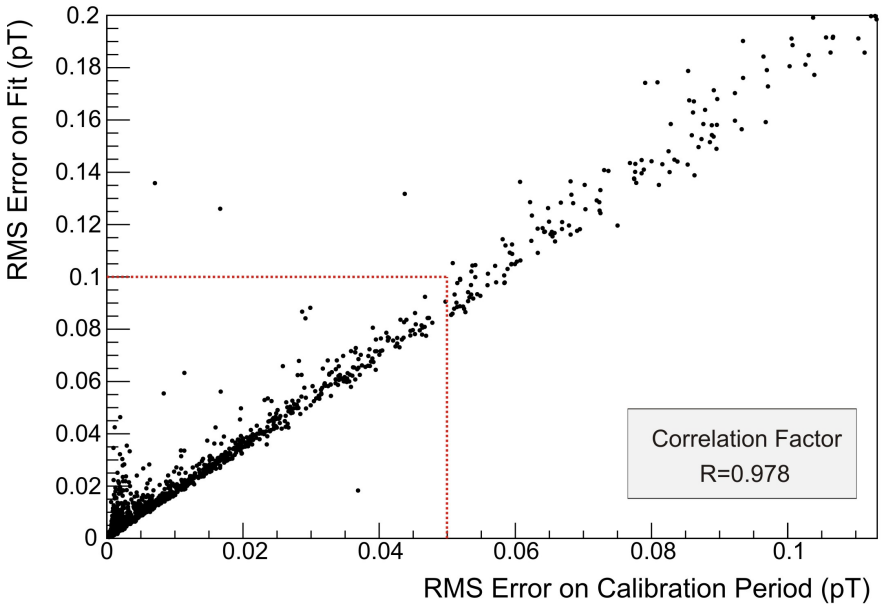

FIGURE 7.17: RMS difference of the fit over 120 hours of fitting after calibration for a 64-hour calibration period, plotted against the RMS difference of the fit for that period. The strong correlation continues up to higher values for both parameters than shown. The red lines illustrate a cut that could be made on the data, such that $> 99\%$ of the remaining fits have RMS errors within the 0.1 pT requirement.

Lastly, we examine whether there is a correlation between the performance of the fit during the calibration period and that for data taken afterwards. In Figure 7.15 we saw that, for fits trained with 16+ hours of data, if a fit demonstates an RMS error below the 0.1 pT threshold within the first few hours, it maintains a sufficient level of performance for many hours after calibration. By extension, it is not unreasonable to expect that how well it fits on the calibration data may have bearing on how well it does on any later data. What we see is indeed a strong correlation between the two, which grows with the size of the calibration period.

As an example, we look at a sample of 2000 runs with 64 hours of calibration. If all points with a demonstrated RMS error over this period of greater than 0.05 pT are excluded, we are left with 1428 runs, of which 8 have an RMS error over the subsequent 120 hours of data of more than 0.1 pT (see Figure 7.17). By making a simple cut on the data, discounting around a quarter of the runs, we can achieve a success rate of over 99% in terms of fits which meet the 0.1 pT requirement. Of those which do not make the cut, there is a higher incidence of near-field scenarios, but there is no other indication as to why they fail. For these near-field runs, these will be the points that are not illustrated in Figure 7.16, where the results do not conform with the majority.

7.6 Conclusions

This chapter has investigated two techniques to determine the fields within the neutron cells: a linear regression method and one based on a neural network. This type of procedure, whereby the SQUID signals are combined to fit the neutron cell signals, was seen as the only viable choice for the current setup of the SQUID system – the pickup loops are situated on one side of the neutron cells, which makes physically interpolating the signal an impossibility. These fitting processes required that we split the data into a calibration period on which they are trained, where there is no electric field and thus the outcomes are known, and a data-taking period thereafter, on which we test how well it performs.

The two methods were refined to optimise their performance, the linear network by reworking the matrix in the event of singularities resulting from similarities between the SQUID signals; the neural network by examining different possible architectures and calibration procedures. A neural network with one hidden layer with 13 nodes was chosen, but was not seen to produce competitive results over a long range of time. Provided that the linear fit had been trained with enough

data (16 hours or more), it was seen to yield a fit that was within 0.1 pT of the true neutron signal, tens of hours after calibration, around 70% of the time.

To identify which runs these correspond to, the correlations between the RMS errors and the parameters of the SQUID and neutron signals were investigated to see if there is any positive indication, but no significant connection could be found. However, a strong correlation could be seen between the performance of the fit over the calibration period and the rest of the data. With a cut made on the data for linear fits trained for 64 hours, 99% of the remaining runs are within the 0.1 pT requirement for the next 120 hours.

This bodes well, but is of course heavily contingent on the model that we have developed over the course of Chapter 6. This was reasoned not to be lacking for our investigation, but further characterisation of the magnetic environment would be beneficial. Such information could inform changes to the model, for example, leading to a better way to simulate the temporal variation of the field.

A future design could, if the infrastructure allows, consist of pickup loops packed on the surface of a sphere around the neutron cells, providing more channels to aid in a better fit that should be more sensitive to fields in all three orthogonal directions, but also allowing for the interpolation of the field using the Laplace equation. This provides a second means of determining the fields, which could serve as a useful cross-check of the results. Many nEDM experiments incorporate two different magnetometers in order to achieve this, and if this additional function could be realised, it would give a boost towards a future cryogenic experiment.

Conclusions and Further Outlook

The quest for a neutron electric dipole moment is well-motivated, as was discussed in Chapter 1, since it is sensitive to new sources of CP violation that can potentially explain the baryon asymmetry that exists in the universe today. Current and next generation nEDM experiments incorporate a vast array of different technologies in an effort to improve their sensitivity to d_n. Almost all of these are bottled UCN experiments that seek to maximise the neutron density, the storage time and the electric fields that the neutrons are exposed to during the Ramsey measurement, as well as maintain a high degree of spin polarisation leading up to the measurement and a high neutron detector efficiency. Chapter 2 emphasised that a stable and well-known magnetic field is also a necessity, and therefore a way of monitoring the field is imperative. As experiments move towards cryogenic conditions to capitalise upon improved UCN production and magnetic field generation, so it becomes necessary to find an alternative to traditional magnetometers that can operate in sub-Kelvin temperatures.

SQUIDs are a prominent option, but they have been little explored in the context of an nEDM experiment. This thesis has served as an exposition of the design considerations, construction and expected performance of a SQUID magnetometry system developed for CryoEDM, a low-temperature nEDM experiment, that demonstrates that SQUIDs are a definite possibility for an nEDM magnetometer in future cryogenic projects.

Chapter 3 presented a brief introduction to the CryoEDM experiment and the operation of DC SQUIDs. Crucially, it established that the SQUIDs in CryoEDM must track the magnetic field in a neutron cell with a precision of 0.1 pT. It further outlined the infrastructure of the CryoEDM SQUID system and discussed improvements that were made, including the construction of new readout cabling, which reduces the heat load, and new pickup loops to act as sensing areas for the

SQUIDs. Again, these are all considerations and features that are applicable to prospective nEDM experiments.

A characterisation of the shielding efficiency of the superconducting capillary shields installed on the input lines was performed and reported in Chapter 4. It was estimated that a shielding factor of $\sim 10^4$ was needed, such that the magnetic pickup on the input lines of the SQUID due to fluctuations in the ILL experimental hall would be suppressed below 0.1 pT. A three-axis Helmholtz coil setup was devised to apply fields to both shielded and exposed lengths of twisted wire pairs attached to SQUIDs inputs. Through comparing their responses, the conservative limits of $\sim 10^7$ and $\sim 10^6$ for the shielding factors in the transverse and longitudinal directions are inferred, which are both orders of magnitudes above the requirement. These experiments also showed that the shielding factor of shields damaged through wear and tear can be restored to some extent. They represent the first systematic study of this shielding and, moreover, constitute one of few documents concerning the performance of solder as a magnetic shield that have entered into the literature.

The effect of Cryoperm on the intrinsic noise of the CryoEDM SQUIDs was investigated for the first time in Chapter 5. The work demonstrated that there was an increase in the low-frequency, $1/f$ part of the noise spectrum for SQUIDs operated without Cryoperm, particularly when cooled in a substantial magnetic field. However, for the conditions in CryoEDM, it was estimated that this translates into a difference in the noise voltage of less than a factor of 1.5.

This chapter also documented magnetostatic simulations undertaken to determine whether the Cryoperm, due to it being a high permeability material in close proximity to the neutron guide, could cause depolarisation of neutrons in transport. They implied that the Cryoperm was detrimental, and proposed a compensation coil to counter its effect. However, this adds another degree of complexity to the system. Combining the results of both lines of investigation, it was reasoned that

if the SQUIDs remained in the same position, that it would be better to run them without Cryoperm for CryoEDM, since the impact on the spin polarisation appears more drastic than that on the SQUID noise. However, it was suggested that precautions should be taken, for example, in mediating the field the SQUIDs are cooled in. In short, this work revealed a definite tension between the two consequences of using Cryoperm, and that the SQUIDs should ideally be located away from the neutrons in a prospective design to avoid the need to compromise.

A novel simulation of SQUID and neutron magnetic signals was presented in Chapter 6, with the aim of producing data that can be used to investigate potential methods for tracking the fields in the neutron cells. The magnetic environment was modelled as temporally-variant, stationary dipole sources, and the resultant flux which traverses the pickup loops and the volume-average field in the neutron cells were found numerically. It was noted that the discretisation error in the calculation of the fluxes and fields is negligible, and reasoned that this description of the fields was adequate for determining a prescription to reconstruct a neutron cell measurement, especially as the background is expected to be dominated by low-frequency fields. The biggest challenge was to include the mu metal and superconducting shields: this was achieved using a generalised Laplace equation solver for magnetic scalar potential formulated in cylindrical coordinates. Techniques to accelerate the rate of convergence were investigated, and it was found that using a full multigrid method with successive overrelaxation and a gradual incrementation of the relative permeability of the mu metal from $\mu_r = 1$ to $\mu_r = 10000$ was most effective. The field map was determined from the resulting potentials through trilinear interpolation, to then find the SQUID and neutron measurements.

With a statistically large data set, the work detailed in Chapter 7 looked to find a prescription for determining the neutron cell fields. It was suggested to use a combination of the SQUID signals to fit to the neutron signals, and two distinct techniques were examined for that purpose: a linear, least squares regression, and a method based on a neural network. These were optimised, and in the case of

the latter, it was found that a network with 13 neurons in a single hidden layer was the most suitable architecture. However, this did not perform as well as the linear fit on a consistent basis, with the regression technique being successful in finding a fit within 0.1 pT of the true signal for tens of hours around two-thirds of the time. It was found that, as an example, with a calibration period of 64 hours and a data-taking period of 120 hours, if the RMS error on the data that this fit was trained on was less than 0.05 pT, then the fit would meet the 0.1 pT criterion for the rest of the data $> 99\%$ of the time (see Section 7.5).

All of this work together paints the picture of a functional SQUID system that can reach the target 0.1 pT resolution, which bodes well for the future. But in some respects, the design of the SQUID system for CryoEDM was driven out of necessity to conform to the pre-existing layout of the apparatus. This resulted in, for example, long input lines that required shielding, the location of the pickup loops to one side of the neutron cells, which complicates finding the fields in these volumes, and the positioning of the SQUIDs around the neutron guide, whose Cryoperm shielding could cause spin depolarisation. This thesis has shown that, to some extent, each of these issues can be resolved, but they are not ideal situations to have to contend with and that they are best avoided. A future experiment should be designed with the magnetometry in mind from the very start. This is especially given the foreseen importance of magnetometry, as systematic errors associated with the magnetic field begin to catch up with, and eventually overtake, the neutron statistics as the limitation on the experimental sensitivity.

If the SQUIDs are to be the sole magnetometer, then it is important that they are utilised as best as they can be. Chapter 7 proposed a design whereby the pickup loops are packed on the surface of a sphere around the neutron cells, so that their signals could be used not only to fit to that of the neutrons, but also to interpolate the fields via a multipole expansion. This provides a second method of establishing the measurement as a cross-check. It remains to be seen whether this would be viable, and whether it would be possible to acheive full solid-angle

coverage – there is the issue, for instance, of proximity to the high voltage end, and the possibility of cross-talk between the loops.

At any rate, the future design could aim for shorter connections between the SQUIDs and the loops, which will still require shielding but are, nevertheless, less of a liability, with the SQUIDs mounted within their own vessel away from the neutrons, such that they can be run with Cryoperm with no foreseen issue. If this vessel could be attached and detached, and could be mounted in a different cryostat, it provides a degree of modularity in the system such that the SQUIDs and their connections could be tested prior to a full cooldown.

There are numerous potential modifications and improvements that could be made for a proposed SQUID system in a new experiment. However, this thesis has shown, via examinations of its components and predicted performance, that the present design for CryoEDM is adequate to meet its sensitivity specification. It has laid the groundwork, which can be built upon, that clearly demonstrates SQUIDs to be a competitive option for an nEDM magnetometer.

Appendix A

Formulae for Coil-Generated Magnetic Fields

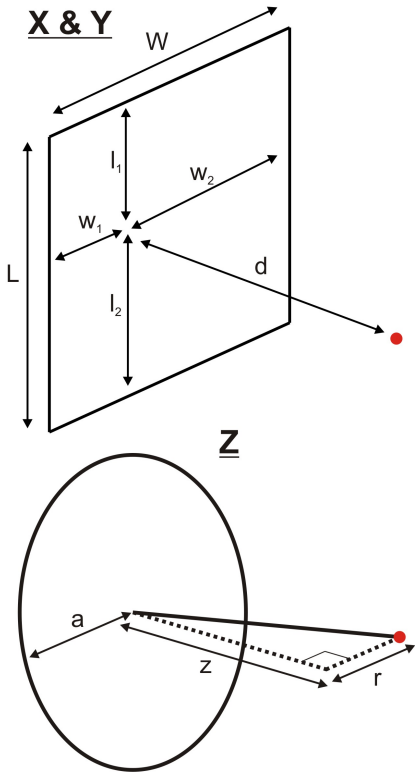

Z Coils

Field parallel to coil axis

$$\frac{B_0}{\pi\sqrt{Q}}\left[E(k)\,\frac{1-\alpha^2-\beta^2}{Q-4\alpha}+K(k)\right]$$

Radial field

$$\frac{B_0\gamma}{\pi\sqrt{Q}}\left[E(k)\,\frac{1+\alpha^2+\beta^2}{Q-4\alpha}-K(k)\right]$$

$$\alpha=\frac{r}{a},\ \beta=\frac{z}{a},\ \gamma=\frac{z}{r}$$
$$Q=\left[(1+\alpha)^2+\beta^2\right],\ k=\frac{4\alpha}{Q}$$

where B_0 is the field at the centre of the coil; $K(k)$ is the complete elliptic integral function of the first kind and $E(k)$ that of the second kind

X & Y Coils

Integrated field parallel to coil axis (x/y)

$$\frac{\mu_0 nI}{4\pi}\left(\frac{l_1(\sqrt{l_1^2+w_2^2+d^2}-\sqrt{l_1^2+w_1^2+d^2})}{l_1^2+d^2}+\frac{l_2(\sqrt{l_2^2+w_2^2+d^2}-\sqrt{l_2^2+w_1^2+d^2})}{l_2^2+d^2}\right)$$
$$+\frac{\mu_0 nI}{4\pi}l\left[\operatorname{arctanh}\left(\frac{\sqrt{l_1^2+w_1^2+d^2}}{l_1}\right)+\operatorname{arctanh}\left(\frac{\sqrt{l_2^2+w_1^2+d^2}}{l_2}\right)\right.$$
$$\left.+\operatorname{arctanh}\left(\frac{\sqrt{l_1^2+w_2^2+d^2}}{l_1}\right)+\operatorname{arctanh}\left(\frac{\sqrt{l_2^2+w_2^2+d^2}}{l_2}\right)\right]$$

Integrated field pointing along TWP (z)

$$\frac{\mu_0 nI}{4\pi}\left[\arctan\left(\frac{l_1 w_1}{d\sqrt{l_1^2+w_1^2+d^2}}\right)+\arctan\left(\frac{l_1 w_2}{d\sqrt{l_1^2+w_2^2+d^2}}\right)\right.$$
$$\left.+\arctan\left(\frac{l_2 w_2}{d\sqrt{l_2^2+w_2^2+d^2}}\right)+\arctan\left(\frac{l_2 w_1}{d\sqrt{l_2^2+w_1^2+d^2}}\right)\right]$$

Integrated field perpendicular to both coil axis and TWP (y/x)

$$\frac{\mu_0 nI}{4\pi}\left(-\frac{d\sqrt{w_1^2+l_1^2+d^2}}{l_1^2+d^2}+\frac{d\sqrt{w_1^2+l_2^2+d^2}}{l_2^2+d^2}-\frac{d\sqrt{w_2^2+l_1^2+d^2}}{l_1^2+d^2}+\frac{d\sqrt{w_2^2+l_2^2+d^2}}{l_2^2+d^2}\right)$$

Appendix B

The Finite Multigrid Method (FMG)

B.1 The Algorithm

A discretised, general linear system of equations on a grid spacing of h can be expressed as

$$A_h u_h = b_h. \tag{B.1}$$

A_h in our case is the finite difference approximation to the Laplace operator; u_h denotes the exact solution of our potentials; b_h is the source term, which here is zero, but we retain for clarity. Representing the approximation we obtain to u_h at a given iteration as $\tilde{u}_h$, the *error* or correction may be written as

$$e_h = u_h - \tilde{u}_h. \tag{B.2}$$

We further define the *residual* as the difference between the two sides of Equation B.1 for an approximate solution, that is

$$r_h = b_h - A_h \tilde{u}_h. \tag{B.3}$$

If we were to apply the linear operator A_h to both sides of Equation B.2, we are left with the expression

$$A_h e_h = r_h. \tag{B.4}$$

How do these equations fit into the FMG? It is not the solutions at each stage that will be translated between the grids, but rather the errors and residuals. This is because an approximation on a fine grid cannot be accurately represented on a coarser grid due to aliasing, so the transfer of this information and subsequent relaxation could begin to undo the work on the fine grid. The idea is then to iterate on the error on the coarse grid, using Equation B.4, and then use this to amend the error on the fine grid on a larger scale, without losing the local refinements.

The backbone of the FMG is the V-Cycle, which represents a single application of this principle. A summary of the scheme for two grids of spacing h (fine) and $2h$ (coarse) is as follows:

1. **Pre-smoothing:** iterate $A_h u_h = b_h$ to yield the initial approximation, $\tilde{u}_h$.
2. Find and transfer (**restrict**) the residual $r_h = b_h - A_h \tilde{u}_h$ to the coarse grid.
3. Set e_{2h} to be zero and iterate $A_{2h} e_{2h} = r_{2h}$ to find the low frequency error.
4. Transfer (**prolongate**) this error to the fine grid and **correct** the approximation: $\tilde{u}_h \rightarrow \tilde{u}_h + e_h$.
5. **Post-smoothing:** iterate $A_h u_h = b_h$, with the updated approximation as an initial guess.

The process can be extended to multiple grids (Figure B.1) by imposing the condition that steps 2-4 will only execute if not on the coarsest of grids, and by adding a step 3.5, to begin executing another V-Cycle for all grids of lower dimension. This recursive method will keep repeating for each level until it is called upon the lowest, where it solves exactly for the error, and this correction is interpolated back up to the highest level. Many of these V-Cycles are strung together in a FMG, repeatedly impoving upon the finest grid approximation. A further advancement

is made by first iterating on the coarser grids, each time prolongating $\Phi_{i,j,k}$ to the next finer grid. This provides better initial conditions than a null grid at the finest level, from which we can make adjustments. This yields the classic full multigrid scheme, as is seen in Chapter 6 in Figure 6.10.

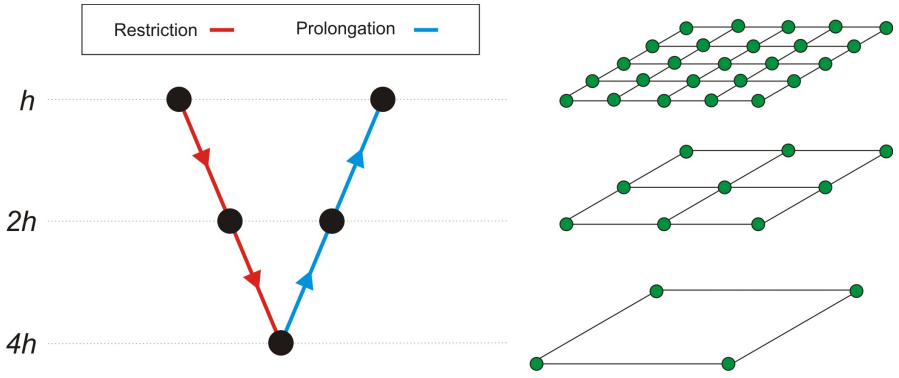

FIGURE B.1: A three stage V-Cycle.

The FMG is generally incompatible with Succesive Overrelaxation (Section 6.2.2), and thus standard Gauss-Seidel is used. SOR overcorrects the approximation, which is useful for tackling the larger frequency modes in the error. In doing so, however, it destructs any smoothing at high frequencies. As the FMG targets the roughness at every stage, the SOR is not a good choice for a smoothing operator.

B.2 The Transfer of Information

Methods of transferring information were especially developed for the FMG scheme, since there are no common prescriptions for a cylindrical, staggered grid. The first problem we encounter is dealing with the innermost points, since their potential values cannot be as easily related to those of nearby points. Labelling these the 0th elements, in accordance with their radial numbering, we define a temporary 'ghost' point for each $r - \theta$ cross-section of the cylinder. The ghost lies at the centre point of the circle and is taken as the average of the 0th elements in the initial grid. When a transfer process calls for an element of one lower index in

the radial direction, the ghost is called for each of the 0th points. The procedures devised here are therefore only slightly different for these points compared to the rest of the grid, as shall now be detailed.



FIGURE B.2: Points along a given radius from a coarse grid (white) with separation $h \equiv \Delta r$ and fine grid (black) with half the separation, $h/2 \equiv 0.5\Delta r$, along with the centre of the circle (red). Due to the staggered grid configuration, the coarse grid points are not a subset of the fine grid elements. The prolongation process, which weights points according to distance, will, for instance, estimate the point at $3/4h$ as $3/4$ that at $1/2h$ plus $1/4$ that at $3/2h$.

Prolongation takes data from the coarse grid and interpolates it to each of the positions in the fine grid. The staggered grid approach means that the coarse and fine grids do not overlap to any degree, and therefore that every one of the final grid points must be found through some combination of values from the initial grid, as illustrated in Figure B.2. The process takes a weighted average of the points either side of the fine grid point, taking account of their relative separation. After performing a radial sweep of the grid, interpolation azimuthally and vertically is more straightforward, although we must make sure to take into account the continuation in going from 2π to 0. The end result is a procedure similar to that of trilinear interpolation, as described in Section 6.2.3.

Restriction is the reverse process of transferring data from a finer to a coarser resolution. There are two general types of restriction: injection and full weighting. In injection, the coarse grid data point is directly read in from the corresponding point on the fine grid, whilst with full weighting, the surrounding grid points are averaged. The second is preferable as it better approximates the fine grid function in the frequency components that can be accurately represented by the coarse grid. Aliasing becomes a problem with injection as it causes high frequency contributions to manifest in the low frequency modes of the coarse grid. Further,

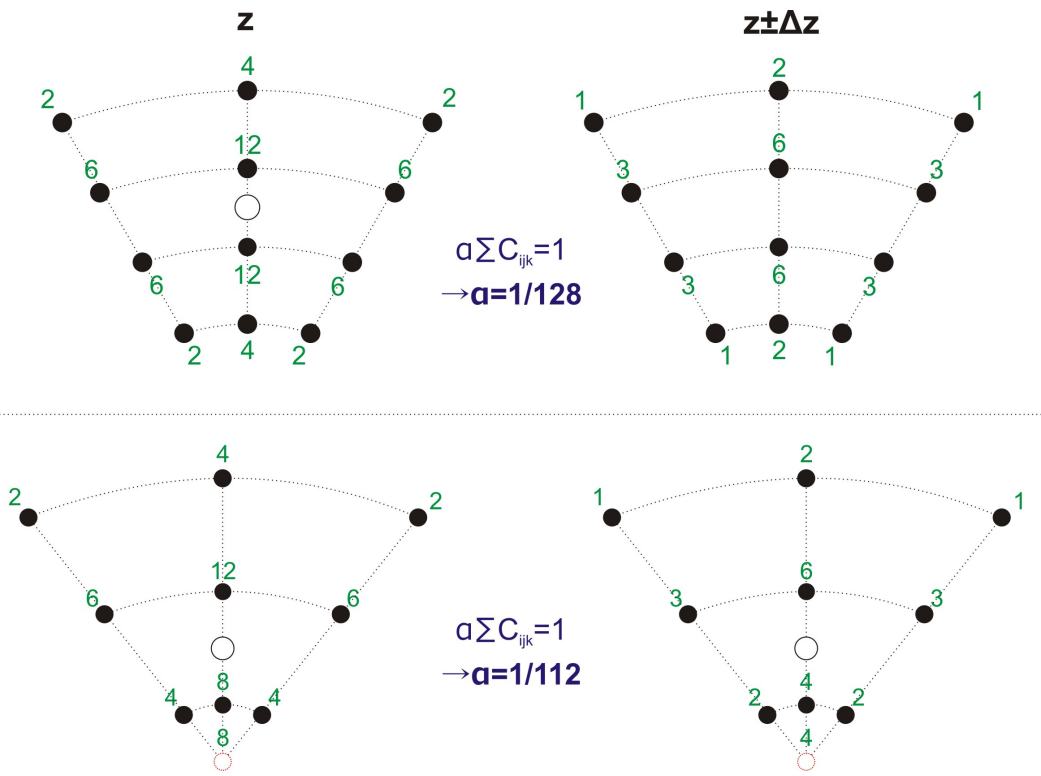

FIGURE B.3: The stencil for restriction, showing all the fine grid points that contribute to a given coarse grid point for a point at $i > 1$ above, and a point at $i = 0$ below. The weightings of these points, C_{ijk}, in green are based on their indices with repect to the coarse grid point, and thus roughly based on their relative distance. These are normalised by an α factor, which is the reciprocal of the total of all the weightings.

as previously discussed, as there is no overlap between the two grids in our case, there must at least be some radial interpolation applied.

To implement a full weighting for our grid, we use the nearest neighbours (and next nearest neighbours in the radial case) and loosely weight them according to index and thus approximate distance. The prescription is strictly the inverse of the prolongation operator that we have defined, when represented in matrix form, that is then normalised, as illustrated in Figure B.3.

B.3 Computational Cost and Performance

Chapter 6 describes how the time to reach convergence increases with the size of the grid, and how the FMG counteracts the increase by reducing the number of iterations that are needed to reach convergence. The term 'convergence criterion' has been lightly used in Chapter 6 to describe the end point of relaxation. The definition with respect to the residual is a consequence of the fact that we of course do not know the true answer on the grid, and thus the error at each stage. As the Laplace Equation has a unique solution, having used the same discretised prescription, so must Equation B.4. This means that a null residual vector must correspond to a null error vector. Although we can never reduce the residual to zero, given that the boundary values for Equation B.4 are zero, one would expect that the entries of the error vector will be small if those of the residual are.

Another thing to clarify is that the error quoted so far is the difference between the current approximation and the solution as discretised on the grid, that is, not the true analytical solution. The error therefore consists of this difference, termed the algebraic error, and that due to the finite sampling, known as the discretisation error. With the true solution labelled as u, our discretised solution as u_h and our current approximation as $\tilde{u}_h$, the total error can be related to these errors via

$$||u - \tilde{u}_h|| = ||u - u_h + u_h - \tilde{u}_h|| \leq ||u - u_h|| + ||u_h - \tilde{u}_h||, \tag{B.5}$$

using the Cauchy-Schwarz inequality. It is the discretisation error that is the limitation as nothing can be done to modify it for a specified grid size in a given problem, whilst we can keep reducing the algebraic error with each iteration.

To begin to judge the efficiency of the FMG, we look at the computational cost of the V-Cycle at its heart. If there are N computations in an iteration of the fine grid, a V-Cycle consisting of an iteration at each level will have at most $2N(1+{}^1\!/_8+{}^1\!/_{64}+...) = {}^{16}\!/_7 N$. Thus it is not necessarily computationally expensive.

As we have discussed, there is little point in decreasing the algebraic error beyond the discretisation error. The discretisation error for a 3D grid of equal spacing $ah = {}^{a}/_{n}$ and number of points ${}^{n}/_{a}$ in each dimension (n for the finest grid), will go as $O(h^a) = O(n^{-a})$. Assuming a static convergence factor, γ, for the V-Cycle, we can formulate the requirement that

$$\gamma^{\nu} = O(n^{-a}), \tag{B.6}$$

and therefore that the number of iterations of the V-Cycle required, $\nu = O(\log n)$. As the computational cost of a single V-Cycle is $O(N) = O(n^d)$, where $d = 3$ is the number of dimensions, the cost of converging to the level of the discretisation error is $O(n^d \log n)$. Moreover, as the FMG method provides an initial guess at each stage of the grid, the cost can fall to $O(n^d)$ - see [165] for further discussion.

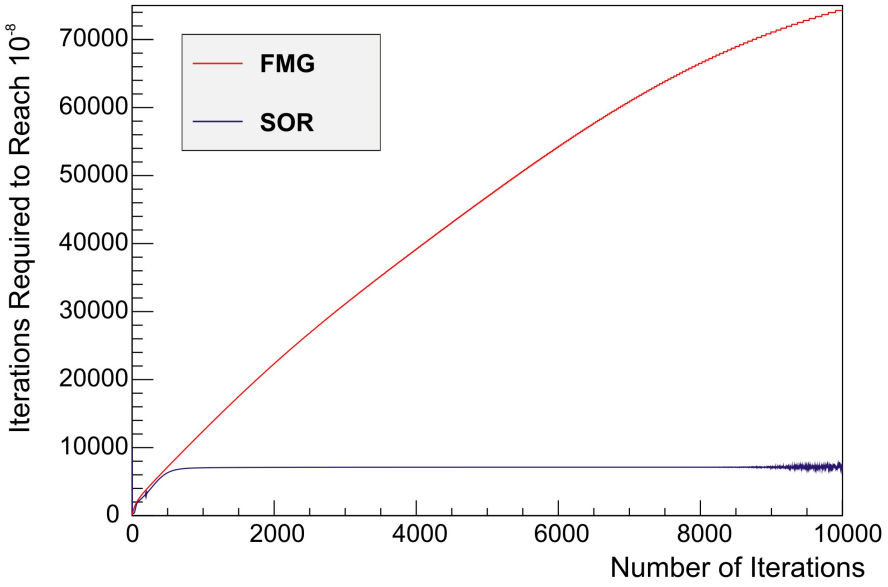

FIGURE B.4: The number of iterations it would take to achieve a 10^8 reduction in the residual using the instantaneous convergence rate on a $160 \times 160 \times 160$ grid. The FMG becomes slower and slower as the solver proceeds, whilst the SOR always has the better convergence rate and peters out. With larger grids, the trend in the FMG rate does not change considerably, whereas the rates for the SOR are noticeably slower. At some grid size, the FMG may be preferable, but this will be at a point when the grid size is much bigger and the time for solution much, much longer, more than would be viable for running the simulation multiple times.

Now it remains to be seen how this translates to the performance of the FMG in our scenario. We take as an example a problem without shields on a $160 \times 160 \times 160$ grid, and define the FMG with four levels of grid, and use two iterations on each level. The initial accelerated convergence is encouraging, but with more and more iterations, the rate of convergence becomes slower and slower. This is because the available correction from the coarser grids becomes iteratively smaller and smaller, as the roughness is smoothed more quickly than that on the finest. The convergence factor tends asymptotically towards that of the Gauss-Seidel regime which, as we saw in Figure 6.8, requires many times more iterations than the SOR method to solve a problem. Thus, if a problem reaches the convergence criterion within a few iterations of the FMG convergence factor surpassing that of the SOR method, then the FMG will find the solution quicker. However, in this instance, this point is reached in 59 V-Cycles, well in advance of the iterations needed to achieve the threshold for convergence: the problem completes in 961 s, as opposed to the 211 s with SOR technique. For computing the solutions to these grids, therefore, the FMG as it is is not a viable alternative to the SOR operation.

B.4 Improved Relaxation Method

It is difficult to suggest whether changes can be made to the FMG framework to better its performance, or whether it is inherently unsuitable for these problems. For instance, the circumferential distance as used in the azimuthal case gets bigger with distance from the centre of the cylinder, such that the concept of fine scale changes across the grid. This could lead to differing rates of improvement in the error that could eventually hinder the progress made for the inner points. Much of the work that is detailed here is as a result of trial and error to see if an improved prescription could be found.

Despite the reasoning against combining the FMG with SOR, we look to see how this would perform in comparison. The residual is, as expected, seen to climb

straight after a V-Cycle, before resettling and starting to decrease once more. Having modified the number of iterations at each stage to 200, the rebound results in a residual less than the SOR equivalent at the same stage, suggesting a little of the fine scale information must be retained. In another example problem without shields on a $160 \times 160 \times 160$ grid, the solution was found using this new FMG in 203 s, as opposed to 223 s with the SOR method alone. The decrease in time is typically around 10%, which may be small, but is an improvement nonetheless.

The FMG is designed to overcome the computational penalty in solving on a finer grid, but does not address the impact of a larger permeability. In Chapter 6, it was noted that the convergence proceeds with the inverse of the highest permeability present in the problem. By gradually incrementing the permeability we can not only take advtange of the quicker convergence rates at lower permeabilities, but also present the problem at higher permeability with a better initial guess for the grid. The FMG offers a natural place to integrate this at the beginning of each V-Cycle, where the permeability is increased if the residual is within a certain threshold and is stable. The best size of the increment is difficult to assess, because the impact of an increase is greater for lower permeabilities than at higher. This inevitably leads to some divergent behaviour as the grid adapts to the new permeability, but not noticeably more than with the standalone FMG with SOR smoothing. Figure B.5 illustrates a simplified step-by-step procedure, which was reached in a conscious effort to minimise this behaviour. A convergence analysis would be needed to determine the mathematically optimal sequence, but this is beyond the scope of this thesis. The method in Figure B.5 proves very effective: taking the example of the supposed 60-hour problem discussed in Section 6.2.2, this is seen to complete in around 80 minutes. This makes it feasible to run this program for several hundred dipoles, as is statistically required.

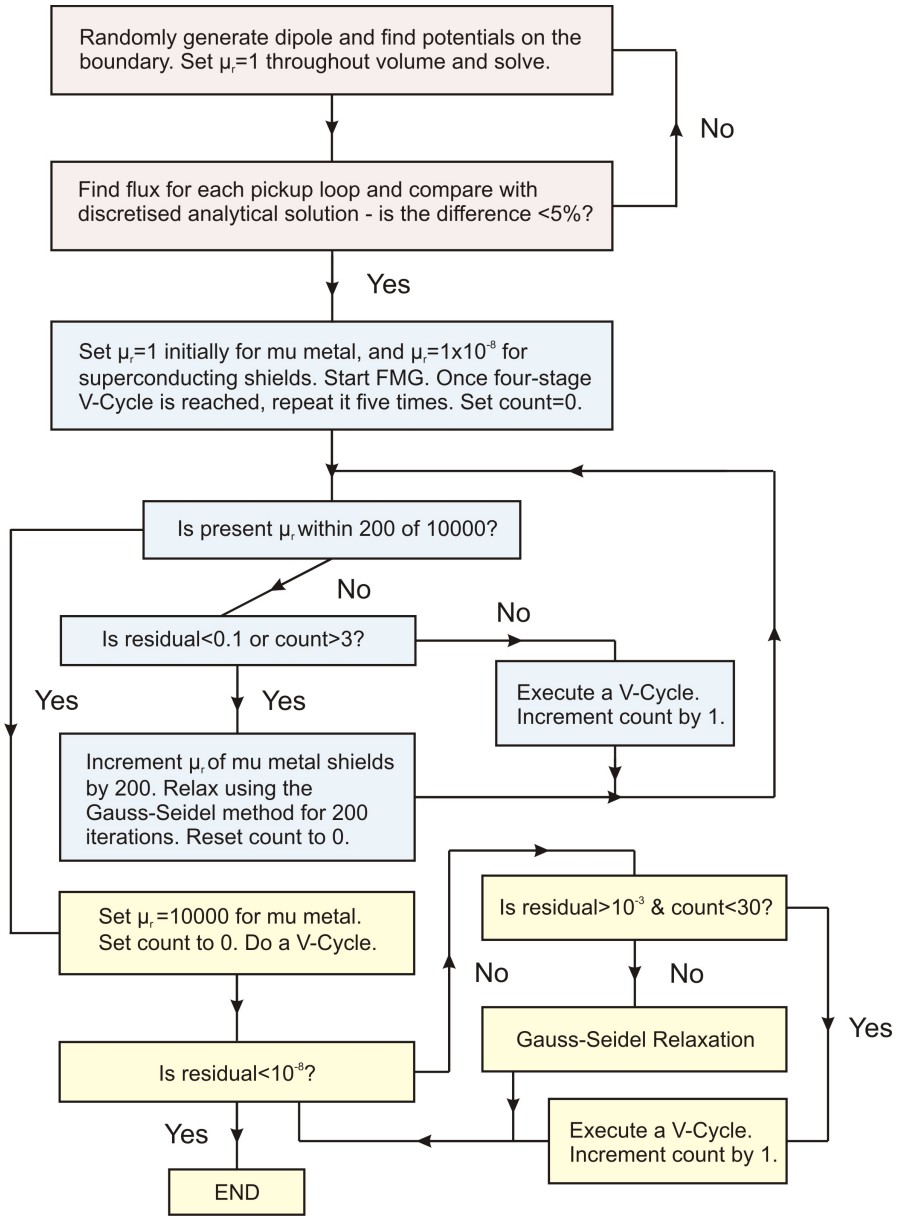

FIGURE B.5: Flow chart depicting the procedure used to find the magnetic scalar potentials within the shields. This can be divided into three main stages: the first, to establish that the problem can be solved to a given tolerance (as further explained in Section 6.2.3); the second the ramping up from $\mu_r = 1$ to the full shielding of $\mu_r = 10000$; and third, a refinement of this solution until it meets the convergence criterion of 10^{-8}.

Bibliography

[1] H. Abele. The neutron. Its properties and basic interactions. *Progress in Particle and Nuclear Physics*, 60(1):1–81, 2008.

[2] D. Dubbers and M. G. Schmidt. The neutron and its role in cosmology and particle physics. *Reviews of Modern Physics*, 83(4):1111, 2011.

[3] C. Baker, D. Doyle *et al.* Improved experimental limit on the electric dipole moment of the neutron. *Physical review letters*, 97(13):131801, 2006.

[4] K. Olive *et al.* The Particle Data Group Collboration. *Chin. Phys. C*, 38: 090001, 2014.

[5] W. Gibson and B. Pollard. *Symmetry Principles Particle Physics*. Cambridge Monographs on Physics. Cambridge University Press, 1980. ISBN 9780521299640.

[6] T. Soldner, L. Beck *et al.* New limit on T violation in neutron decay. *Physics Letters B*, 581(1):49–55, 2004.

[7] G. Lüders. Proof of the tcp theorem. *Annals of Physics*, 2(1):1–15, 1957.

[8] W. Pauli, L. Rosenfeld, and V. Weisskopf. Niels Bohr and the development of physics. *Niels Bohr and the Development of Physics*, 1, 1955.

[9] V. A. Kosteleckỳ and N. Russell. Data tables for Lorentz and CPT violation. *Reviews of Modern Physics*, 83(1):11, 2011.

[10] C. S. Wu, E. Ambler *et al.* Experimental test of parity conservation in beta decay. *Phys. Rev*, 104:1315, 1956.

[11] T. D. Lee and C. N. Yang. Question of parity conservation in weak interactions. *Physical Review*, 104(1), 1956.

[12] L. Landau. On the conservation laws for weak interactions. *Nuclear Physics*, 3(1):127–131, 1957.

[13] J. Christenson, J. Cronin *et al.* Evidence for the 2π decay of the K0. *Phys. Rev. Lett*, 13:138, 1964.

[14] V. Fanti, A. Lai *et al.* A new measurement of direct CP violation in two pion decays of the neutral kaon. *Physics Letters B*, 465(1):335–348, 1999.

[15] B. Aubert, D. Boutigny *et al.* Observation of CP violation in the B0 meson system. *Physical Review Letters*, 87(9):091801, 2001.

[16] K. Abe, R. Abe *et al.* Observation of large CP violation in the neutral B meson system. *Physical Review Letters*, 87(9):091802, 2001.

[17] R. Aaij, C. A. Beteta *et al.* Evidence for CP violation in time-integrated D0 h-h+ decay rates. *Physical review letters*, 108(11):111602, 2012.

[18] A. D. Sakharov. Violation of CP invariance, C asymmetry, and baryon asymmetry of the universe. *JETP Lett.*, 5:24–27, 1967.

[19] P. Ade, N. Aghanim *et al.* Planck 2013 results. xvi. cosmological parameters. *A&A*, A16(571), 2014.

[20] P. Huet and E. Sather. Electroweak baryogenesis and standard model CP violation. *Physical Review D*, 51(2):379, 1995.

[21] F. An, J. Bai *et al.* Observation of electron-antineutrino disappearance at Daya Bay. *Physical Review Letters*, 108(17):171803, 2012.

[22] R. Barbieri, P. Creminelli *et al.* Baryogenesis through leptogenesis. *Nuclear Physics B*, 575(1):61–77, 2000.

[23] M. Aguilar, J. Alcaraz *et al.* The alpha magnetic spectrometer (AMS) on the international space station: Part I–results from the test flight on the space shuttle. *Physics Reports*, 366(6):331–405, 2002.

[24] A. Czarnecki and B. Krause. Neutron electric dipole moment in the standard model: complete three-loop calculation of the valence quark contributions. *Physical review letters*, 78(23):4339, 1997.

[25] S. Dar. The neutron EDM in the SM: A review. *arXiv preprint hep-ph/0008248*, 2000.

[26] T. P. Cheng and L.-F. Li. Why the naive quark model can yield a good account of the baryon magnetic moments. *Phys. Rev. Lett.*, 80:2789–2792, 1998.

[27] M. Pospelov and A. Ritz. Electric dipole moments as probes of new physics. *Annals of Physics*, 318(1):119–169, 2005.

[28] S. Narison. A fresh look into the neutron EDM and magnetic susceptibility. *Physics Letters B*, 666(5):455–461, 2008.

[29] R. D. Peccei. The strong CP problem and axions. In *Axions*, pages 3–17. Springer, 2008.

[30] K. Nakamura, P. D. Group *et al.* Axions and other similar particles. *JPG*, 37:075021, 2010.

[31] R. N. Mohapatra and G. Senjanović. Natural suppression of strong P and T non-invariance. *Physics Letters B*, 79(3):283–286, 1978.

[32] S. M. Barr. Natural class of non-Peccei-Quinn models. *Physical Review D*, 30(8):1805, 1984.

[33] M. J. Ramsey-Musolf and S. Su. Low-energy precision tests of supersymmetry. *Physics Reports*, 456(1):1–88, 2008.

[34] C.-P. Liu, M. Ramsey-Musolf *et al.* Atomic electric dipole moments: The Schiff theorem and its corrections. *Physical Review C*, 76(3):035503, 2007.

[35] J. Hudson, D. Kara *et al.* Improved measurement of the shape of the electron. *Nature*, 473(7348):493–496, 2011.

[36] A. C. Vutha, W. C. Campbell *et al.* Search for the electric dipole moment of the electron with thorium monoxide. *Journal of Physics B: Atomic, Molecular and Optical Physics*, 43(7):074007, 2010.

[37] J. Baron, W. C. Campbell *et al.* Order of magnitude smaller limit on the electric dipole moment of the electron. *Science*, 343(6168):269–272, 2014.

[38] M. E. Pospelov and I. Khriplovich. Electric dipole moment of the W boson and the electron in the Kobayashi-Maskawa model. *Soviet Journal of Nuclear Physics*, (53):638–640, 1991.

[39] G. Bennett, B. Bousquet *et al.* Improved limit on the muon electric dipole moment. *Physical Review D*, 80(5):052008, 2009.

[40] F. Farley, K. Jungmann *et al.* New method of measuring electric dipole moments in storage rings. *Physical review letters*, 93(5):052001, 2004.

[41] W. Griffith, M. Swallows *et al.* Improved limit on the permanent electric dipole moment of hg 199. *Physical review letters*, 102(10):101601, 2009.

[42] M. Rosenberry and T. Chupp. Atomic electric dipole moment measurement using spin exchange pumped masers of 129 xe and 3 he. *Physical review letters*, 86(1):22, 2001.

[43] V. Fedorov, M. Jentschel *et al.* Perspectives for nEDM search by crystal diffraction. test experiment and results. *Nuclear Physics A*, 827(1):538c–540c, 2009.

[44] F. M. Piegsa. New concept for a neutron electric dipole moment search using a pulsed beam. *Physical Review C*, 88(4):045502, 2013.

[45] J. Smith, E. Purcell, and N. Ramsey. Experimental limit to the electric dipole moment of the neutron. *Physical Review*, 108(1):120, 1957.

[46] C. Baker, Y. Chibane *et al.* Apparatus for measurement of the electric dipole moment of the neutron using a cohabiting atomic-mercury magnetometer. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 736(0):184 – 203, 2014.

[47] N. F. Ramsey. A molecular beam resonance method with separated oscillating fields. *Phys. Rev.*, 78:695–699, 1950.

[48] S. Balashov, K. Green *et al.* A proposal for a cryogenic experiment to measure the neutron electric dipole moment (nEDM). *arXiv preprint arXiv:0709.2428*, 2007.

[49] R. Golub and J. Pendlebury. The interaction of ultra-cold neutrons (UCN) with liquid helium and a superthermal ucn source. *Physics Letters A*, 62(5): 337–339, 1977.

[50] C. Baker, S. Balashov *et al.* Experimental measurement of ultracold neutron production in superfluid He-4. *Physics Letters A*, 308(1):67–74, 2003.

[51] H. Yoshiki, H. Nakai, and E. Gutsmiedl. A new superleak to remove He-3 for UCN experiments. *Cryogenics*, 45(6):399 – 403, 2005.

[52] M. Mukhopadhyay. *Fundamentals of Cryogenic Engineering*. PHI Learning, 2010. ISBN 9788120330573.

[53] A. Saunders, T. Bowles *et al. AIP Conference Proceedings*, 529(1).

[54] A. Saunders, J. Anaya *et al.* Demonstration of a solid deuterium source of ultra-cold neutrons. *Physics Letters B*, 593(1):55–60, 2004.

[55] S. Arzumanov, L. Bondarenko *et al.* Storage of ultracold neutrons in vessels whose walls are made from graphite, fluorine polymer oil, or heavy-water ice. *Physics of Atomic Nuclei*, 66(10):1820–1830, 2003.

[56] T. Brys, M. Daum *et al.* Diamond-like carbon coatings for ultracold neutron applications. *Diamond and related materials*, 15(4):928–931, 2006.

[57] F. Atchison, B. Blau *et al.* Measurement of the Fermi potential of diamond-like carbon and other materials. *Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms*, 260(2): 647 – 656, 2007.

[58] F. Atchison, B. Blau *et al.* Diamondlike carbon can replace beryllium in physics with ultracold neutrons. *Physics Letters B*, 642(1):24–27, 2006.

[59] F. Atchison, B. Blau *et al.* Storage of ultracold neutrons in a volume coated with diamondlike carbon. *Physical Review C*, 74(5):055501, 2006.

[60] A. Serebrov, V. Varlamov *et al.* Measurement of the neutron lifetime using a gravitational trap and a low-temperature Fomblin coating. *Physics Letters B*, 605(1):72–78, 2005.

[61] J. Stahn and D. Clemens. A remanent Fe/Si supermirror transmission polarizer. *Applied Physics A*, 74(1):s1532–s1534, 2002.

[62] A. Serebrov, M. Lasakov *et al.* Superconducting UCN polarizer for a new EDM spectrometer. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 545(1–2):490 – 492, 2005.

[63] T. Brys, M. Daum *et al.* Measurement of the loss and depolarization probability of ucn on beryllium and diamond like carbon films. *Journal of Research of the National Institute of Standards and Technology*, 110(3):279–281, 2005.

[64] R. Golub and P. Huffman. Search for a neutron electric dipole moment. *Journal of Research of the National Institute of Standards and Technology*, 110(3):169–172, 2005.

[65] A. Lynch. *A Cryogenic Scintillation UCN Detector for a Neutron EDM Experiment*. PhD thesis, University of Oxford, 2014.

[66] C. Baker, S. Balashov *et al.* Development of low temperature solid state detectors for ultra-cold neutrons within superfluid ^{4}He. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 501(2):517–523, 2003.

[67] M. Lasakov, A. Serebrov *et al.* A large area silicon UCN detector with the analysis of UCN polarization. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 545(1–2):301 – 308, 2005.

[68] G. Ban, X. Flechard *et al.* ^{6}Li doped glass scintillators for ultra cold neutrons detection. In *Nuclear Science Symposium Conference Record, 2004 IEEE*, volume 2, pages 829–832. IEEE, 2004.

[69] M. Klein and C. J. Schmidt. CASCADE, neutron detectors for highest count rates in combination with ASIC/FPGA based readout electronics. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 628(1):9–18, 2011.

[70] R. Bouclier, M. Capeans *et al.* The gas electron multiplier (GEM). *Nuclear Science, IEEE Transactions on*, 44(3):646–650, 1997.

[71] V. Hélaine. The neutron electric dipole moment experiment at the Paul Scherrer Institute. In *EPJ Web of Conferences*, volume 73, page 07003. EDP Sciences, 2014.

[72] J. Karamath. *He-II under High Electric Field Conditions for the Cryogenic Neutron Electric Dipole Moment Experiment.* PhD thesis, University of Sussex, 2007.

[73] W. Rogowski. Die elektrische festigkeit am rande des plattenkondensators. *Archiv für Elektrotechnik*, 12:1–15, 1923.

[74] J. Zenner. *The search for the neutron electric dipole moment.* PhD thesis, Johannes Gutenberg-Universität, Mainz, 2013.

[75] B. Rebel, C. Hall *et al.* High voltage in noble liquids for high energy physics. *Journal of Instrumentation*, 9(08):T08004, 2014.

[76] A. J. Davidson. *High Voltage Breakdown in the Ramsey Cell of the CryoEDM experiment. - An experimental study of some relevant parameters.* PhD thesis, University of Sussex, 2011.

[77] F. Paschen. Ueber die zum funkenübergang in luft, wasserstoff und kohlensäure bei verschiedenen drucken erforderliche potentialdifferenz. *Annalen der Physik*, 273(5):69–96, 1889.

[78] Y. Masuda, K. Asahi *et al.* Neutron electric dipole moment measurement with a buffer gas comagnetometer. *Physics Letters A*, 376(16):1347 – 1351, 2012.

[79] T. M. Ito, D. H. Beck *et al.* High electric field development for the SNS nEDM experiment. *eprint arXiv:1401.5435*, 2014.

[80] M. Parker. The kerr magneto-optic effect (1876–1976). *Physica B+ C*, 86: 1171–1176, 1977.

[81] N. Nouri and B. Plaster. Comparison of magnetic field uniformities for discretized and finite-sized standard , solenoidal, and spherical coils. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 723(0):30 – 35, 2013.

[82] D. U. Gubser, S. A. Wolf, and J. E. Cox. Shielding of longitudinal magnetic fields with thin, closely spaced, concentric cylinders of high permeability material. *Review of Scientific Instruments*, 50(751-756), 1979.

[83] T. J. Sumner, J. M. Pendlebury, and K. F. Smith. Convectional magnetic shielding. *Journal of Physics D: Applied Physics*, 20(9):1095, 1987.

[84] C. P. Bidinosti and J. W. Martin. Passive magnetic shielding in static gradient fields. *AIP Advances*, 4(4):047135, 2014.

[85] S. Afach, G. Bison *et al.* Dynamic stabilization of the magnetic field surrounding the neutron electric dipole moment spectrometer at the Paul Scherrer Institute. *Journal of Applied Physics*, 116(8):084510, 2014.

[86] T. Bryś, S. Czekaj *et al.* Magnetic field stabilization for magnetically shielded volumes by external field coils. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 554(1–3):527–539, 2005.

[87] J. M. Pendlebury, W. Heil *et al.* Geometric-phase-induced false electric dipole moment signals for particles in traps. *Phys. Rev. A*, 70:032102, 2004.

[88] P. G. Harris, J. M. Pendlebury, and N. E. Devenish. Gravitationally enhanced depolarization of ultracold neutrons in magnetic-field gradients. *Phys. Rev. D*, 89:016011, 2014.

[89] J. Pendlebury, K. Smith *et al.* Search for a neutron electric dipole moment. *Physics Letters B*, 136(5–6):327 – 330, 1984.

[90] K. Smith, N. Crampin *et al.* A search for the electric dipole moment of the neutron. *Physics Letters B*, 234(1–2):191 – 196, 1990.

[91] K. Green, P. Harris *et al.* Performance of an atomic mercury magnetometer in the neutron EDM experiment. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 404(2–3):381 – 393, 1998.

[92] Y. Kim and S. Clayton. Tests of a SQUID-based He-3 co-magnetometer readout for a neutron EDM experiment. *Applied Superconductivity, IEEE Transactions on*, PP(99):1, 2014.

[93] S. Gröger. *Laser-pumped cesium magnetometers for the PSI-nEDM experiment*. PhD thesis, University of Fribourg, 2005.

[94] A. J. Fairweather and M. J. Usher. A vector rubidium magnetometer. *Journal of Physics E: Scientific Instruments*, 5(10):986, 1972.

[95] B. Patton, E. Zhivun *et al.* All-optical vector atomic magnetometer. *Phys. Rev. Lett.*, 113:013001, 2014.

[96] A. Vershovskii. Project of laser-pumped quantum MX magnetometer. *Technical Physics Letters*, 37(2), 2011.

[97] S. J. Seltzer and M. V. Romalis. Unshielded three-axis vector operation of a spin-exchange-relaxation-free atomic magnetometer. *Applied Physics Letters*, 85:4804–4806, 2004.

[98] H. B. Dang, A. C. Maloof, and M. V. Romalis. Ultrahigh sensitivity magnetic field and magnetization measurements with an atomic magnetometer. *Applied Physics Letters*, 97(15):151110, 2010.

[99] D. Drung, C. Assmann *et al.* Highly sensitive and easy-to-use SQUID sensors. *Applied Superconductivity, IEEE Transactions on*, 17(2):699–704, 2007.

[100] C. Baker, G. Ban *et al.* The search for the neutron electric dipole moment at the Paul Scherrer Institute. *Physics Procedia*, 17:159–167, 2011.

[101] I. Altarev, D. Beck *et al.* A next generation measurement of the electric dipole moment of the neutron at the FRM-II. *Nuovo Cimento-C*, 35(4):122, 2012.

[102] I. Altarev, G. Ban *et al.* Towards a new measurement of the neutron electric dipole moment. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 611(2):133–136, 2009.

[103] K. Matsuta, Y. Masuda *et al.* An EDM measurement with a new comagnetometer and a high density UCN source. In *11th Conferenc on the Intersections of Particle and Nuclear Physics (CIPANP 2012)*, volume 1560, pages 152–154. AIP Publishing, 2013.

[104] R. Golub and P. Huffman. Search for a neutron electric dipole moment. *Journal of Research of the National Institute of Standards and Technology*, 110(3):169–172, 2005.

[105] T. M. Ito, n. Collaboration *et al.* Plans for a neutron EDM experiment at SNS. In *Journal of Physics: Conference Series*, volume 69, page 012037. IOP Publishing, 2007.

[106] B. Filippone, M. Cooper, and P. Huffman. Search for a neutron electric dipole moment at the SNS.

[107] R. Golub and S. K. Lamoreaux. Neutron electric-dipole moment, ultracold neutrons and polarized ^{3}He. *Physics Reports*, 237(1):1–62, 1994.

[108] S. Balashov, K. Green *et al.* Internal note - cryoedm parameters. 2007.

[109] J. Bardeen, L. N. Cooper, and J. R. Schrieffer. Microscopic theory of superconductivity. *Physical Review*, 106(1):162, 1957.

[110] B. S. Deaver and W. M. Fairbank. Experimental evidence for quantized flux in superconducting cyclinders. *Physical Review Letters*, 7(2):43–46, 1961.

[111] R. Doll and M. Näbauer. Experimental proof of magnetic flux quantization in a superconducting ring. *Physical Review Letters*, 7(2):51, 1961.

[112] H. Weinstock. *SQUID Sensors: Fundamentals, Fabrication and Applications*. Nato Science Series E:. Springer Netherlands, 2012. ISBN 9789401156745.

[113] D. E. McCumber. Effect of ac impedance on dc voltage-current characteristics of superconductor weak-link junctions. *Journal of Applied Physics*, 39 (7):3113–3118, 1968.

[114] W. C. Stewart. Current-voltage characteristics of josephson junctions. *Applied Physics Letters*, 12(8):277–280, 1968.

[115] Supracon. SQUID current sensor model CSBlue: Technical parameters, 2007. Supracon AG, Jena, Germany.

[116] Supracon. SQUID current sensor model CSBlue: Description, 2007. Supracon AG, Jena, Germany.

[117] S. Ingleby. *Cryodetector Readout for Direct Dark Matter Searches*. PhD thesis, University of Oxford, 2012.

[118] X. Zhang. *EDELWEISS*. PhD thesis, University of Oxford, 2015.

[119] D. Benford, T. Powers, and S. Moseley. Thermal conductivity of Kapton tape. *Cryogenics*, 39(1):93–95, 1999.

[120] G. T. Dyos and T. Farrell. *Electrical resistivity handbook*. Number 10. IET, 1992.

[121] R. L. Powell, H. M. Roder, and W. M. Rogers. Low-temperature thermal conductivity of some commercial coppers. *Journal of Applied Physics*, 28 (11):1282–1288, 1957.

[122] S. Henry, M. McCann, and C. Clarke. Internal communications, 2007.

[123] D. Shoenberg. *Superconductivity*. Cambridge monographs on physics. Cambridge University Press, 1962. ISBN 9780608106755.

[124] M. McCann. *Magnetometry for CryoEDM*. PhD thesis, University of Oxford, 2012.

[125] F. Grover. *Inductance calculations, working formulas and tables*. Dover Publications, 2004. ISBN 9780486495774.

[126] S. Henry, M. Pipe *et al.* Characterisation of superconducting capillaries for magnetic shielding of twisted-wire pairs in a neutron electric dipole moment experiment. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 763: 155–162, 2014.

[127] B. Seeber. *Handbook of Applied Superconductivity*. Number v. 2. Taylor & Francis, 1998. ISBN 9781420050271.

[128] A. Mager. Magnetic shields. *Magnetics, IEEE Transactions on*, 6(1):67–75, 1970.

[129] W. H. Warren and W. G. Bader. Superconductivity measurements in solders commonly used for low temperature research. *Review of Scientific Instruments*, 40(1):1, 1969.

[130] A. Shepelev. chapter The Discovery of Type II Superconductors (Shubnikov Phase). INTECH Open Access Publisher, 2010. ISBN 9789533071077.

[131] J. L. Kirschvink. Uniform magnetic fields and double-wrapped coil systems: improved techniques for the design of bioelectromagnetic experiments. *Bioelectromagnetics*, 13(5):401–411, 1992.

[132] S. Denis, M. Dirickx *et al.* Field penetration into hard type-II superconducting tubes: effects of a cap, a non-superconducting joint, and non-uniform superconducting properties. *Superconductor Science and Technology*, 20(5): 418, 2007.

[133] D. Karmakar. Studies of magnetization and related properties for hard type-II superconductors. *Indian J.Phys*, 10(79):1107–1147, 2005.

[134] R. H. Koch, J. Clarke *et al.* Flicker (1/f) noise in tunnel junction DC SQUIDs. *Journal of low temperature physics*, 51(1-2):207–224, 1983.

[135] R. H. Koch, D. P. DiVincenzo, and J. Clarke. Model for 1/f flux noise in SQUIDs and qubits. *Physical review letters*, 98(26):267003, 2007.

[136] L. Faoro and L. B. Ioffe. Microscopic origin of low-frequency flux noise in Josephson circuits. *Physical review letters*, 100(22):227005, 2008.

[137] S. Anton, J. Birenbaum *et al.* Magnetic flux noise in DC SQUIDs: temperature and geometry dependence. *Physical review letters*, 110(14):147002, 2013.

[138] S. Machlup. Noise in semiconductors: Spectrum of a two-parameter random signal. *Journal of Applied Physics*, 25(3):341–343, 1954.

[139] A. Khintchin. Korrelationstheorie der stationren stochastischen prozesse. *Mathematische Annalen*, 109(1):604–615, 1934.

[140] K. Taconis, J. Beenakker *et al.* Measurements concerning the vapour-liquid equilibrum of solutions of ^{3}He in ^{4}He below 2.19 ° K. *Physica*, 15(8):733–739, 1949.

[141] P. J. Rousseeuw and K. Van Driessen. Computing lts regression for large data sets. *Data mining and knowledge discovery*, 12(1):29–45, 2006.

[142] J. Gail, M. Mück, and C. Heiden. Pinning and depinning of single vortices in niobium thin film DC superconducting quantum interference devices. *Applied physics letters*, 73(18):2663–2665, 1998.

[143] W. Pogosov and V. Misko. Vortex quantum tunneling versus thermal activation in ultrathin superconducting nanoislands. *Physical Review B*, 85(22): 224508, 2012.

[144] A. Wallraff, A. Kemp, and A. Ustinov. chapter Quantum Dynamics of Vortices and Vortex Qubits, pages 162–185. Wiley-VCH Verlag GmbH & Co. KGaA, 2006. ISBN 9783527606085.

[145] A. Katsika. *Magnetic Aspects of the Cryo-nEDM Experiment.* PhD thesis, University of Sussex, 2011.

[146] *CryoEDM Collaboration Meeting - Effect of Cooling Procedures on Field Homogeneity*, 2005.

[147] J. R. Grozier. *The Cryogenic Neutron Electric Dipole Moment Experiment: Magnetic Challenges and Solutions.* PhD thesis, University of Sussex, 2007.

[148] C. Crawford. Precision mangetic fields for fundamental neutron symmetries. Technical report, University of Kentucky, 2012.

[149] W. H. Press, S. A. Teukolsky *et al.* *Numerical recipes in C*, volume 2. Citeseer, 1996. ISBN 9780521431088.

[150] J. R. Nagel. Solving the generalized poisson equation using the finite-difference method (FDM). *Department of Electrical and Computer Engineering University of Utah, Salt Lake City, Utah*, 2011.

[151] I. Yavneh. Why multigrid methods are so efficient. *Computing in science & engineering*, 8(6):12–22, 2006.

[152] N. Black and S. Moore. *MathWorld–A Wolfram Web Resource.*

[153] R. Barrett, M. W. Berry *et al.* *Templates for the solution of linear systems: building blocks for iterative methods*, volume 43. Siam, 1994. ISBN 9780898713282.

[154] K. S. Kunz and R. J. Luebbers. *The finite difference time domain method for electromagnetics*. CRC press, 1993. ISBN 9780849386572.

[155] K. Riley, P. Hobson, and S. Bence. *Mathematical Methods for Physics and Engineering: A Comprehensive Guide*. Cambridge University Press, 2006. ISBN 9781139450997.

[156] K. Gurney. *An Introduction to Neural Networks*. Taylor & Francis, 2003. ISBN 9780203451519.

[157] C. Delaere. TMultiLayerPerceptron. *CERN ROOT. November*, 1, 2012.

[158] G. Cybenko. Approximation by superpositions of a sigmoidal function. *Mathematics of control, signals and systems*, 2(4):303–314, 1989.

[159] K. Hornik, M. Stinchcombe, and H. White. Multilayer feedforward networks are universal approximators. *Neural networks*, 2(5):359–366, 1989.

[160] A. Blum. *Neural Networks in C++: An Object-Oriented Framework for Building Connectionist Systems*. Number v. 1 in Wiley professional computing. John Wiley & Sons, 1992. ISBN 9780471552017.

[161] G. H. D. E. Rumelhart and R. J. Williams. Learning representations by back-propagating errors. *Nature*, 323:533–536, 1986.

[162] S. J. Wright and J. Nocedal. *Numerical optimization*, volume 2. Springer New York, 1999. ISBN 9780387400655.

[163] R. C. S. L. L. Giles. Overfitting in neural nets: Backpropagation, conjugate gradient, and early stopping. In *Advances in Neural Information Processing Systems 13: Proceedings of the 2000 Conference*, volume 13, page 402. MIT Press, 2001.

[164] L. Prechelt. Early stopping-but when? In *Neural Networks: Tricks of the trade*, pages 55–69. Springer, 1998.

[165] W. L. Briggs, V. E. Henson *et al. A multigrid tutorial*. Siam, 2000. ISBN 8580000923568.