Magnetometry for cryoEDM

Michael Andrew McCann

Oriel College, Oxford

Thesis submitted in fulfilment of the requirements for the degree of Doctor of Philosophy at the University of Oxford

Trinity Term, 2012

Abstract

The existence of the matter in the universe is still an unsolved puzzle. After the Big Bang, both matter and antimatter should have been created in equal amounts, and subsequently annihilated. The leading theories to explain the existence of matter require an imbalance in the production of matter and antimatter in the early universe. This in turn requires CP violation, an asymmetry of the laws of physics between matter and antimatter. cryoEDM is designed to explore the total amount of CP violation and resolve this issue.

cryoEDM is a next-generation neutron electric dipole moment search in a commissioning phase of development at the Institut Laue-Langevin, Grenoble. A critical requirement of EDM searches is knowledge of the magnetic environment. This work is concerned with the development, implementation, and performance of the currently operating magnetometry system based on SQUID magnetometers.

An analysis scheme to provide magnetometry data over the volume occupied by the neutrons, from measurements using the available magnetometers, is developed. An updated method to calibrate the magnetometers using internal sources of magnetic fields is presented, and found to give good agreement with independent measurements. A new method of calibration using the neutrons as a reference is discussed, and tests on an example arrangement are shown to be promising. Algorithms for detecting and correcting for hardware induced artefacts in the data are produced, and demonstrated to reconstruct the field with good agreement in all but the noisiest environments. A software framework is developed to combine these into a real-time analysis that provides feedback and diagnostics to the experiment.

Using this new system the resolution of the magnetometers installed in cryoEDM is found to be limited by the environmental noise, and would give a false EDM signal that is greater than the statistical uncertainty in neutron counting. However, the resolution has been somewhat artificially limited to reduce the susceptibility to the RF interference present. This still allows the magnetometry to act as a useful diagnostic tool on any issues in the current magnetic environment, even if in a sub-optimal configuration. For example, investigation of the magnetic shielding of the experiment finds a reduction in the shielding relative to the design, a situation which is being addressed with the design of additional shielding. Once this shielding is installed the resolution of the magnetometers will improve as well as the slew rate of the SQUIDs, which is found to be lower than the $47\,\mu\mathrm{Ts}^{-1}$ required to measure AC fields applied during a measurement.

The current system can also determine sources of magnetic perturbations created within the experiment, which will require addressing before a full EDM run can be performed. For example, cryogenic effects are observed to occur approximately hourly causing large shifts in the magnetic field. Also operation of valves controlling the flow of neutrons around the experiment are found to produce both AC magnetic fields from the driving motors, and shifts in the field from their movement. Situations which can be resolved by reexamination of installation and operational procedures.

Acknowledgements

In the compilation of this thesis it would be remiss of me not to thank certain people for their guidance and support. I would like to thank the following people; My supervisors, Prof Hans Kraus, and Dr Maurits van der Grinten for giving me the opportunity to contribute to such an interesting experiment. Hans gave exemplary guidance and support during my time in Oxford, and Maurits gave valuable advice and support while I was in Grenoble. A thank you to Dr Sam Henry, whose advice has been useful during the completion of my thesis.

I'd like to thank the whole of the cryoEDM collaboration, whose existence has enabled me to work in this important field and fostered my desires for research.

A special thank you to the support staff in Oxford, especially for Kim and Sue whose abilities in the dark arts of admin are formidable, and without whom I'd have been lost, likely wandering the streets naked, far too long ago.

A thank you to my lab colleagues, whose ease of distraction has comforted me on many an occasion. I should thank my friends, family and dog for general reassurance and assistance.

I would also like to thank STFC for funding this work.

Contents

List of Figures

List of Tables

Glossary

C-symmetry Charge-conjugation symmetry, a system is C-symmetric if it appears the same when particles are replaced with their charge conjugates.

CKM matrix Cabibbo-Kobayashi-Maskawa matrix. A rotation matrix that transforms quark wavefunctions from mass eigenstates to weak eigenstates.

CP-symmetry A combination of C- and P-symmetries.

CPT-symmetry A combination of C-, P-, and T-symmetries. CPT symmetry is required by Lorentz invariant quantum field theories with Hermitian Hamiltonians.

DAQ Data acquisition.

EDM Electric dipole moment.

FIR filter Finite impulse response filter. A digital filter using only data samples to produce its output.

IIR filter Infinite impulse response filter. A digital filter using feedback of previous values in addition to data samples.

P-symmetry Parity Symmetry, a system is P-symmetric if it appears the same under the co-ordinate transformation $\boldsymbol{r} \mapsto -\boldsymbol{r}$.

PMNS matrix Pontecorvo-Maki-Nakagawa-Sakata matrix. A rotation matrix that transforms neutrino wavefunctions from mass eigenstates to weak eigenstates.

PMT Photomultiplier tube.

Rabi resonance A measure of the population of two non-degenerate states as a function of the frequency of radiation inducing a transition between the states.

Ramsey resonance Similar to a Rabi resonance with a phase accumulation period carried out in an intermediate state. The outer envelope is that of a Rabi resonance.

SCV Superfluid containment vessel, the container within the magnetic shields where a measurement is performed holding the electrodes and magnetometry and where the neutrons spend most time.

SQUID Superconducting quantum interference device, a sensitive magnetic flux to voltage converter.

SUSY Supersymmetry. An additional symmetry that transforms between fermionic and bosonic states.

T-symmetry Time-reversal symmetry, a system is T-symmetric if it appears the same under the transformation $t \mapsto -t$

UCN Ultracold Neutrons, neutrons with a kinetic energy lower than the Fermi potential of certain materials, typically less than 1 μeV. They can be confined due to total reflection from these materials at any incidence angle.

Introduction

cryoEDM is a next generation cryogenic neutron electric dipole moment (EDM) search, currently being constructed by a collaboration with over 15 years of experience with neutron EDM searches, hoping to shed light on beyond standard model sources of CP violation. The aim is to obtain a two orders of magnitude improvement in the world's best limit on the neutron EDM, currently held by this group, reaching a sensitivity of 1×10^{-28} e cm. To accomplish this goal it requires a magnetometry system to operate in the sub-Kelvin temperatures within the experiment, and able reach sensitivities beyond the capabilities of current systems used in such experiments. In this thesis the work done to develop the magnetometry system for cryoEDM is detailed. This work is the culmination of three years of development at the University of Oxford and one year on-site at the Institut Laue-Langevin in Grenoble, France where the experiment is housed.

The physics goals of cryoEDM, and electric dipole moment searches in general, will be explored in a description of CP violation and EDMs in the first chapter. In the second chapter the measurement principles of previous neutron EDM searches, and the improvements over them that constitute cryoEDM will be described to provide context for the magnetometry requirements needed to obtain the expected experimental limit.

The remainder of the thesis can be broken down into three broad categories:

1. Development and testing of the magnetometry system's software algorithms.

2. Ground work for developing the magnetometry system and a magnetically clean environment.

3. Performance of the magnetometry system as a whole, and characterisation of the magnetic environment as seen by the magnetometry.

These categories will be expanded upon in chapters three to seven. An analysis scheme capable of providing the magnetometry requirements from the available magnetometers is discussed and methods of performing a calibration of the magnetometers is investigated in the third chapter. The appearance of hardware artefacts in the magnetometry signals that require removal and compensation, and details and performances of algorithms to accomplish this correction, will be presented in the fourth chapter.

Details of investigations into the magnetic properties of materials for components to be used in cryoEDM will be given in the fifth chapter.

In the sixth chapter an analysis of the magnetic environment in cryoEDM and the performance of the magnetometers will be given, along with the applicable constraints on the neutron EDM limit derived from these. The seventh chapter will present a discussion of the work preformed and the conclusions of the thesis.

Five appendices are provided, the first deals with the method of determining the value of the neutron EDM from measurements. The second deals with the methods of extracting calibration data for the magnetometers. The third deals with the calibration of the equipment used to investigate the magnetic properties of materials. The last two appendices expand upon elements of the software framework developed.

In short, this work represents the development of the magnetometry system from a disjointed collection of hardware, software requests, and specifications into a functional system serving as a vital operational and analysis component, and useful diagnostic tool within the experiment cryoEDM.

Chapter 1

CP violation and EDMs

1.1 Matter-antimatter asymmetry

One of the longstanding mysteries in physics is why the universe is made of matter. Given a baryon symmetric Big Bang [1], where initials conditions have a zero baryon number, the matter and antimatter contents of the universe should be equal. The equal distributions should have totally annihilated with each other to produce a universe which is empty, apart from radiation. Clearly this is not the case as stars, planets, and indeed ourselves are made from matter.

Putting anthropic arguments aside, the Wilkinson Microwave Anisotropy Probe (WMAP) [2] can place a value on the baryon number in the universe expressed as a ratio to the number of photons [3], this value is

$$\eta = \frac{\eta_B - \eta_{\overline{B}}}{\eta_\gamma} = (6.116^{+0.197}_{-0.249}) \times 10^{-10} \ , \tag{1.1}$$

which is not compatible with zero and leaves an interesting problem.

Among the theories proposed to account for the observation of the existence of matter is baryogenesis [4]. Baryogenesis requires three conditions to be met in order to allow a matter-antimatter asymmetry in the early universe. These conditions were proposed by Sakharov

in 1967 [5]. These Sakharov conditions are

1. Baryon number non-conservation
2. Out-of-thermal-equilibrium interactions
3. C and CP violation

Baryon number non-conservation is possible within the Standard Model [6] through non-perturbative effects [7], and out-of-thermal-equilibrium interactions are to be expected during a first-order electroweak phase transition [8] of the early universe.

It is the third of these conditions that is of interest here. The work of Wu in 1957 [9] showed both the charge conjugation symmetry and the parity symmetry to be maximally violated in the weak interaction. It remains that CP violation is the limiting condition on the matter-antimatter asymmetry, and so this is of particular interest.

1.2 CP violation

Charge-parity symmetry (CP) is an operation on a wave function that replaces particles with their charge conjugates and inverts the co-ordinate system. Until 1964 it was assumed to be an unbroken symmetry of nature, however it is now known to be violated at a small level.

Electromagnetism respects the CP state of the reactants and products during interactions [10]. Despite a generic term in the QCD Lagrangian allowing potentially large CP violating effects,

$$\mathcal{L} = -\frac{n_f g^2 \theta}{32\pi^2} F_{\mu\nu}\tilde{F}^{\mu\nu} \ , \tag{1.2}$$

where n_f is the number of fermions, g is the strong coupling constant, $F_{\mu\nu}$ is electromagnetic tensor with tilde representing its dual, and θ is an arbitrary constant, no experimental evidence of CP violation in the strong sector has been discovered so far [11]. Indeed it is limits

on the electric dipole moment of the neutron that constrain the CP violating parameter such that $|\theta| < 0.7 \times 10^{-11}$ [12], a value so much lower than the order unity one might naturally expect that it prompted Peccei and Quinn to suggest a new U(1) symmetry that promotes θ to become a field [13]. This symmetry, when spontaneously broken, gives the CP violating term its small value. Excitations in this field lead to the so far unobserved particles, axions [14][15].

The CP implication of gravity is hotly debated [16][17] and not fully resolved, however experimental investigation is underway [18][19] using cooled anti-hydrogen. Currently it is the weak interaction that is the only known force to violate CP.

CP violation was first observed indirectly in the neutral kaon sector in 1964 [20]. The combination of $\mathrm{K}_1 = \frac{1}{\sqrt{2}}(\mathrm{K}_0 + \overline{\mathrm{K}}_0)$ and $\mathrm{K}_2 = \frac{1}{\sqrt{2}}(\mathrm{K}_0 - \overline{\mathrm{K}}_0)$, which are CP $+1$ and CP -1 eigenstates, would be equivalent to the weak eigenstates with definite lifetimes K_S and K_L if the CP symmetry was not violated. In this experiment a beam of neutral kaons was prepared such that the short lived K_S component had decayed, leaving a beam of the long lived K_L. The K_L being an expected CP -1 eigenstate should decay predominantly to the CP -1 eigenstate 3π. A small number of the decays produced the CP $+1$ eigenstates 2π with a branching ratio of 2×10^{-3}, hence violating CP. The implication was that the weak eigenstates were in fact a combination $\mathrm{K}_S = \mathrm{K}_1 + \varepsilon \mathrm{K}_2$ and $\mathrm{K}_L = \mathrm{K}_2 + \varepsilon \mathrm{K}_1$ with $\varepsilon = 2.3 \times 10^{-2}$. This is known as indirect CP violation as it manifests in the mixing of the states rather than directly in a decay.

Measurements of direct CP violation have also been made in the kaon sector at CERN [21] and Fermilab [22] through measurements of the four decays:

$$K_{L/S} \rightarrow \pi^0\pi^0 \tag{1.3}$$

$$K_{L/S} \rightarrow \pi^+\pi^- \tag{1.4}$$

the ratio of which is dependent on the direct CP violating parameter $\Re\left(\frac{\varepsilon'}{\varepsilon}\right)$ such that

$$\Re\left(\frac{\varepsilon'}{\varepsilon}\right) \approx \frac{1}{6}\left[\frac{\Gamma(K_L \to \pi^+\pi^-)/\Gamma(K_L \to \pi^+\pi^-)}{\Gamma(K_L \to \pi^0\pi^0)/\Gamma(K_L \to \pi^0\pi^0)} - 1\right]. \tag{1.5}$$

The experiments gave values of this parameter of $(18.5 \pm 4.5(\text{stat}) \pm 5.8(\text{syst})) \times 10^{-4}$ and $(28.0 \pm 3.0(\text{stat}) \pm 2.8(\text{syst})) \times 10^{-4}$.

Measurements have also been performed in the B meson sector at the BaBar [23] and Belle [24] experiments looking at the decays of the B meson. Differences in the decay rates between antimatter partners have been observed for many of these mesons extending the known CP violation beyond just the kaon system. Recently LHCb has measured CP violation in the charmed meson sector at 3.5σ [25] through differences in the decay rates between the decays $D^0 \to K^+K^-$ and $D^0 \to \pi^+\pi^-$, and their CP conjugates. The CP violation measured so far is located solely in the quark sector, and currently all CP measurements are consistent with a CP violating phase in the CKM matrix [26].

However, the amount of CP violation currently measured in the CKM matrix is too small to account for the full extent of the matter-antimatter asymmetry [27][28], providing a baryon asymmetry of only

$$|\eta| < 10^{-26}\ , \tag{1.6}$$

which is several orders of magnitude lower than the measured value given in equation 1.1. It is therefore necessary to look for sources of CP violation beyond the Standard Model.

1.3 Beyond Standard Model sources

Generally all extensions to the Standard Model contain possible sources of CP violation. The neutrino sector can contain a CP violating phase in the PMNS matrix [29] in the same manor as the CKM matrix in the quark sector. Additionally, if neutrinos are Majorana particles, a fermion that is its own antiparticle, two extra phases are introduced [30] into the matrix: with neutrino and antineutrino fields equal fewer phases can be absorbed into the neutrino

fields [31] than in the Dirac case. Majorana neutrinos would be required for neutrinoless double beta decay. The existence of sterile neutrinos [32], neutrinos that do not couple to the electroweak force, would extend the PMNS matrix from 3×3 to $(3 + N) \times (3 + N)$ for N sterile neutrinos, parametrising such a unitary matrix would mean $\frac{((3+N)-1)((3+N)-2)}{2}$ CP violating phases are produced [33] instead of one. Such neutrinos are required if the seesaw mechanism [34], where a very large element in the mass matrix drives one of the mass eigenvalues down, is the correct explanation of the small masses of the neutrinos when compared to other Standard Model particles.

Models with more than one Higgs-doublet [35], multi-Higgs theories, also contain interesting sources of CP violation through complex Yukawa couplings, the couplings between the Higgs and fermion fields, and scalar-pseudoscalar Higgs mixing [36][37]. Additionally multi-Higgs theories can break CP spontaneously [38], meaning that the Lagrangian is CP symmetric but the resulting vacuum is not, and so can produce large intrinsic CP violation. Multi-Higgs systems are a requirement for supersymmetric models.

There are models, known as left-right symmetric models [39], which add an extra SU(2) symmetry as a right handed equivalent of the weak interaction and predict right handed W' and Z' bosons, analogues of the left handed W and Z bosons, which are currently unobserved [40][41]. This introduces 6 CP violating phases in the quark sector and a CP violating phase in the Higgs sector [42] through extra complex couplings.

However, perhaps the largest new CP violating contributions come from supersymmetric theories [43]. Indeed limits on the electric dipole moments have already ruled out supersymmetric theories that permit CP violation at tree level [44].

Supersymmetry (SUSY) is an operation on a wave function that replaces bosons with fermions, and fermions with bosons, such that the two states differ only by their spin. As particles with equal mass but different spins have not been observed SUSY must be a broken symmetry, giving the SUSY partners of Standard Model particles significantly higher masses. This effectively doubles the number of particles in the Standard Model. Since SUSY theories introduce extra fields with their associated couplings there are many ways they can introduce

CP violating effects. Indeed in the general formulation there are 44 CP violating phases [45], enough to produce levels of CP violation much greater than experimental limits allow, as mentioned above. For a full review of the CP implications of SUSY theories the reader is directed to [46]. Supersymmetry is a promising solution to the hierarchy problem [47], where quadratically divergent loop contributions to the Higgs mass push its value towards the Planck mass, due to the partial cancellation of otherwise identical bosonic and fermionic loop contributions.

With the many speculations of extra sources of CP violation, it would be useful to have a way of measuring the scale directly without relying on model dependent processes. This will prove to be the forte of particle electric dipole moments.

1.4 Electric dipole moments

An electric dipole moment is a description of an electromagnetic interaction with a non-spherical charge distribution, specifically a distribution with two equal and oppositely charged poles. A system of two charges $\pm q$, separated by a distance r, gives an electric dipole moment of $\boldsymbol{d} = q\boldsymbol{r}$. Electric dipole moments (EDMs) are a useful tool to get a handle on CP violation. Consider a particle, the only direction associated with it is its spin, therefore the EDM may only be resolved either parallel or antiparallel to the spin. Seen in figure 1.1 is an illustration of how an EDM violates T and P symmetries. Assuming a particle's spin and EDM are aligned performing a parity transformation will flip the axes $\boldsymbol{r} \rightarrow -\boldsymbol{r}$ causing $\boldsymbol{d} \rightarrow -\boldsymbol{d}$ but not affecting the spin. The EDM and spin are then anti-aligned so the states before and after the parity transformation are not equivalent unless $|\boldsymbol{d}| = 0$, ergo a non-zero EDM would violate parity. Similarly a time-reversal transformation would flip just the spin leaving the EDM anti-aligned, again requiring $|\boldsymbol{d}| = 0$ to avoid T-violation.

Assuming CPT is conserved, which is a mathematical consequence of Lorentz invariant quantum field theories [48], the violation in T must be compensated for exactly by violation in CP to keep the combination CPT unviolated. Thus any measurement of an EDM is a

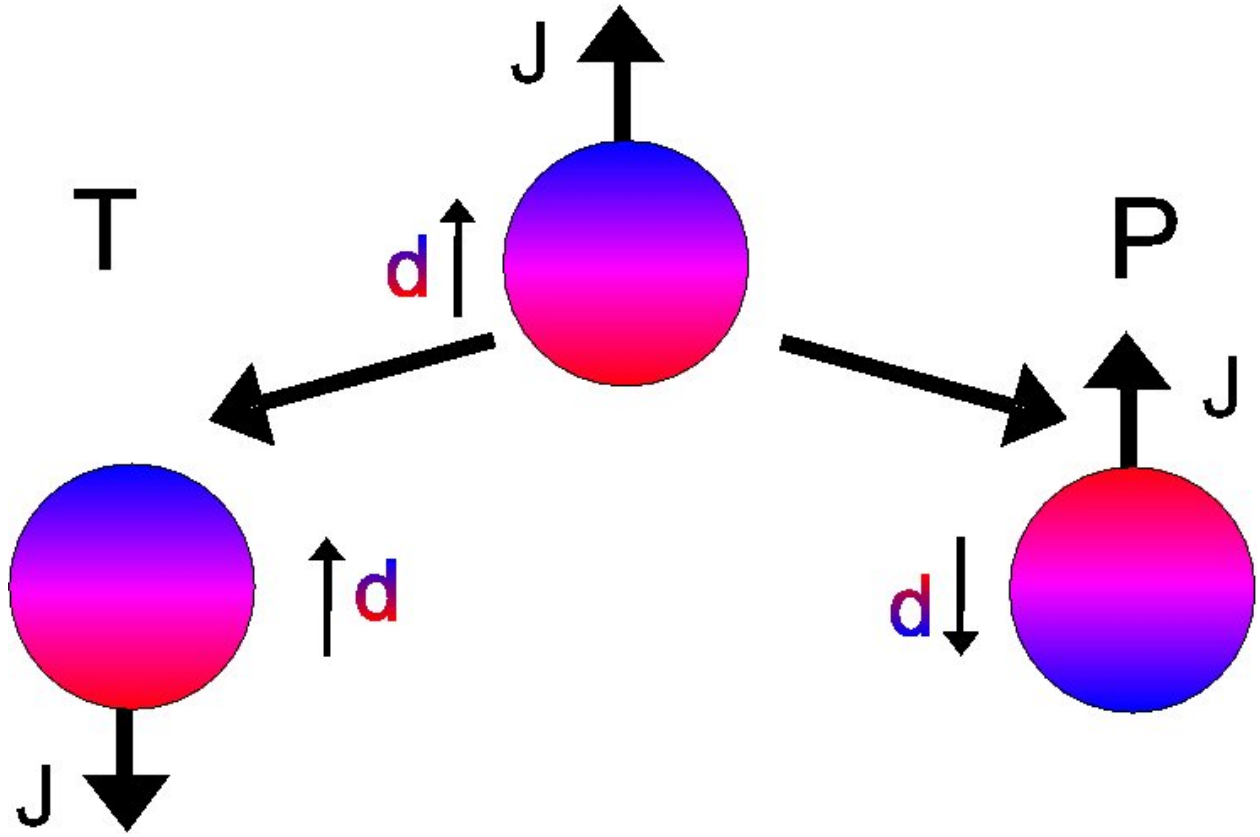

Figure 1.1: Illustration of an EDM violating T and P symmetries. Initially the EDM and spin are aligned, after either a T or P transformation the spin or charge distribution respectively flips leaving the EDM and spin anti-aligned. Neither operations preserve the initial configuration unless $|\boldsymbol{d}| = 0$.

measure of CP violation. For example, a Feynman diagram producing a quark EDM within the Standard Model is shown in figure 1.2. Diagrams with other flavours combinations in the quark loop must also be considered, but for simplicity are ignored here. The diagram is clearly dependent on the CKM matrix parameters $V_{cb}V_{ub}^{*}$, which is different from the CP inverted diagram's term $V_{ub}V_{cb}^{*}$ if there is a non-zero complex phase in the CKM matrix. These theoretical EDM results are a simple illustration, a more detailed treatment of the above is given in [49] and [50]. The focus of this thesis is on the experimental aspects of EDM searches.

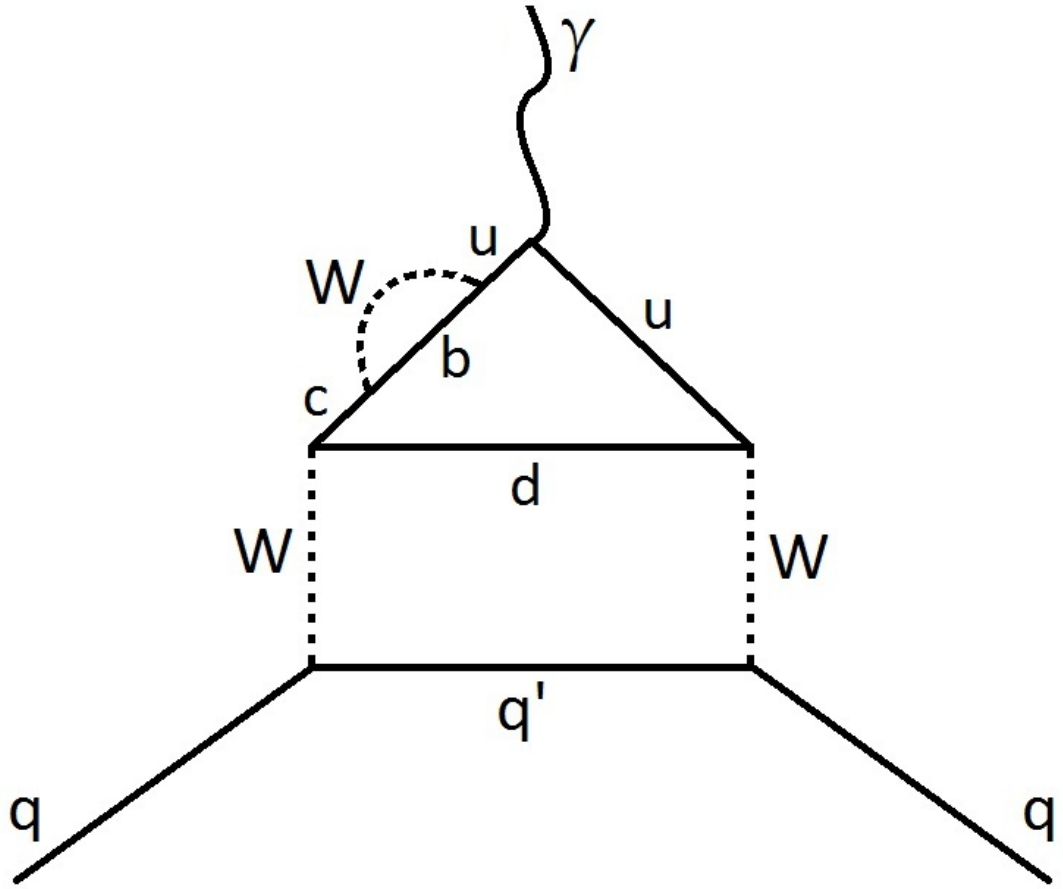

Figure 1.2: An example Feynman diagram showing a process producing an EDM for a quark in the Standard Model to the lowest possible order. The cb and bu W-vertices introduce a term sensitive to the phase in the CKM matrix, $V_{cb}V_{ub}^{*}$, that is not identical to the CP transformed term, $V_{ub}V_{cb}^{*}$, hence this diagram violates CP.

It shall finally be noted that the various sources of beyond Standard Model CP violation give different values to the electric dipole moments of different particle species. So while a single EDM measurement is valuable as a probe into CP violation, a collection of EDM measurements will give a method for model differentiation [51]. Figure 1.3 shows the expected ranges for both electron and neutron electric dipole moments in various popular theories. It is clear that certain theories have already been ruled out, and others give differing ranges for the EDM values for both particles. Table 1.1 shows the most recent limits on the EDM of the leptons and common baryons and notes on how they were obtained. Only the neutron, electron, and mercury EDMs can be measured using a small scale experiments and the mercury EDM already has a very good limit.

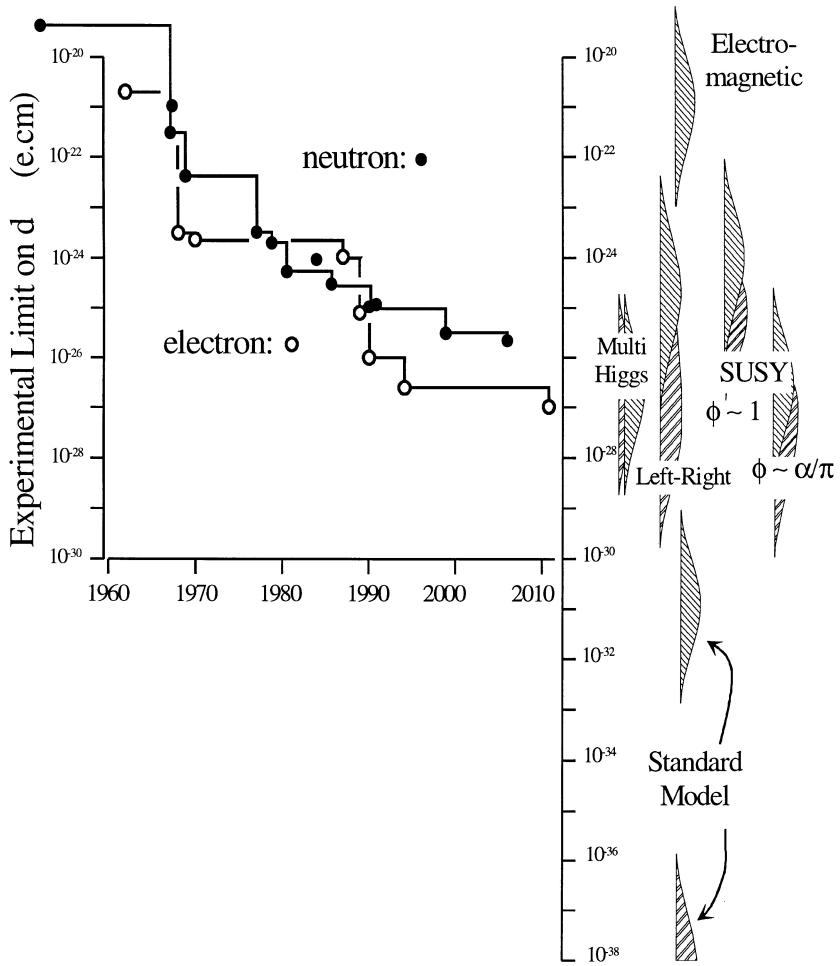

Figure 1.3: Shown are the experimental limits for neutron and electron EDMs with theoretical predictions for their values in certain beyond Standard Model models, shown in [52] updated with recent results. Neutron EDM predictions are double-hatched, electron EDM predictions are single-hatched.

Particle	Experiment/institution(Date)	Method	Limit / e cm
Neutron	nEDM (2003)	Stored neutrons	3.0×10^{-26} [53]
Proton	BINP (2003)	Inferred from Hg	5.4×10^{-23} [54]
e	Imperial College London (2011)	YbF beam	1.1×10^{-27} [55]
μ	muon (g-2)/Brookhaven (2009)	Storage ring	1.8×10^{-19} [56]
τ	BELLE (2003)	$e^+e^- \to \tau\tau$ events	4.5×10^{-17} [57]
Hg^{199}	University of Washington (2009)	Absorption spectrum	3.1×10^{-29} [58]
ν_e,ν_μ	Universidad de Granada (1990)	Astrophysical	2×10^{-21} [59]
ν_τ	Universidad de Granada (1990)	Astrophysical	1.6×10^{-16} [59]

Table 1.1: Table showing the current best experimental limits on the EDM of various particles.

The measurements of the EDMs of the electron and neutron have similar principles (described in section 2.1) but have various technical distinctions. Electron EDMs experience a theoretically determinable enhancement factor by being coupled into atoms or molecules and can be produced in higher densities than free neutrons, however the Standard Model prediction for its EDM is up to seven orders of magnitude lower than that of the neutron. Neutrons, being electrically neutral, can be harder to manipulate but conversely are more stable when experimental electromagnetic fields are applied. Additionally neutrons can be confined in traps for long periods of time where as electrons are typically measured in molecular beams, meaning that much longer measurements may be performed on neutrons than on electrons. cryoEDM is concerned with the measurement of the electric dipole moment of the neutron.

Chapter 2

cryoEDM

EDMs are predicted to be very small, the standard model value of the neutron EDM is predicted to be $|d_n| \sim (10^{-33} - 10^{-31})$ e cm [60] (or $|d_n| \sim 10^{-29}\,\mathrm{GeV}^{-1}$ in natural units). It therefore requires very sensitive measurements to approach this value. A description of the general measurement principle shall be given, followed by a brief overview of previous EDM experiments before focusing on the new generation of experiments and cryoEDM.

2.1 EDM measurement principle

The earliest measurements of electric dipole moments were carried out in electron-neutron scattering by looking for a dipole interaction directly, such as that performed by [61] and interpreted by [62]. The limit in sensitivity of this technique was quickly reached and so a more sensitive, if less direct, method was devised. This method has been the principle for all subsequent neutron EDM experiments.

Analogously to the fact that in the presence of a magnetic field a neutron's magnetic moment will precess around the axis formed by the magnetic field, a neutron's EDM in an electric field will precess around the axis formed by the electric field. As the EDM and spin are aligned, in an applied electric and magnetic field the precession frequency will be the sum of the precession terms from the magnetic moment and EDM. In the case of the fields

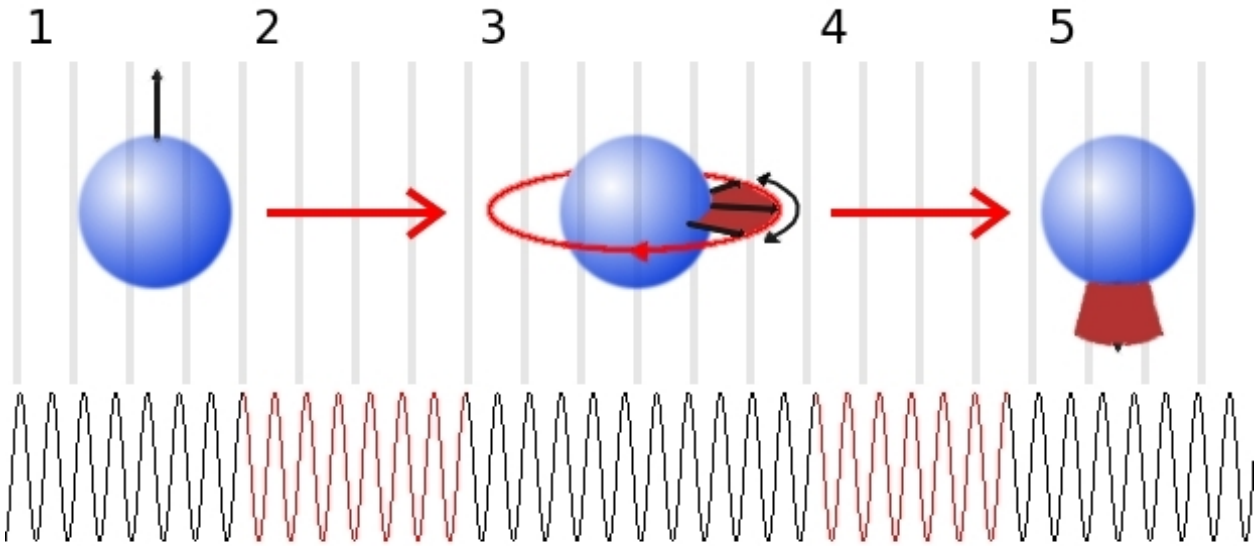

Figure 2.1: Shown is an illustration of the measurement procedure for a Ramsey resonance: 1. A spin polarised (spin-up) neutron precesses in a magnetic field; 2. An RF pulse is applied to the neutron to rotate the spin into an intermediate state; 3. The neutron has a period of free precession during which a phase shift relative to the RF signal may accumulate; 4. The RF pulse is applied a second time to complete the rotation of the spin; 5. The neutrons' spin-state is now dependent on the phase shift introduced during free precession, and is ready for spin analysis.

having a common axis this yields a precession frequency of $\hbar\omega = |2\mu_n B \pm 2d_n E|$, where the sign is given by the relative alignment $\hat{\boldsymbol{B}} \cdot \hat{\boldsymbol{E}}$. Clearly, measuring a change in this precession frequency correlated with a change in the applied electric field is a signal of an EDM.

The method most popular for measuring precise changes in the precession frequency is the method of separated oscillatory fields [63] developed by Ramsey in 1950. An Illustration of the process can be seen in figure 2.1. The measurement starts with a polarised bunch of particles precessing in a magnetic field. A rotating magnetic field is applied to the particles perpendicular to the magnetic field and at a frequency resonant with the precession. In the spin's rotating reference frame this field will appear static and perpendicular to the spin-polarising field, causing the spin to precess around it. This results in the particle's spin rotating between spin-up and spin-down states. The rotating field can be applied until the spins are in an intermediate spin state. Free precession is allowed in this state until the rotating field is applied to rotate the spin back to a pure spin polarisation. However, if

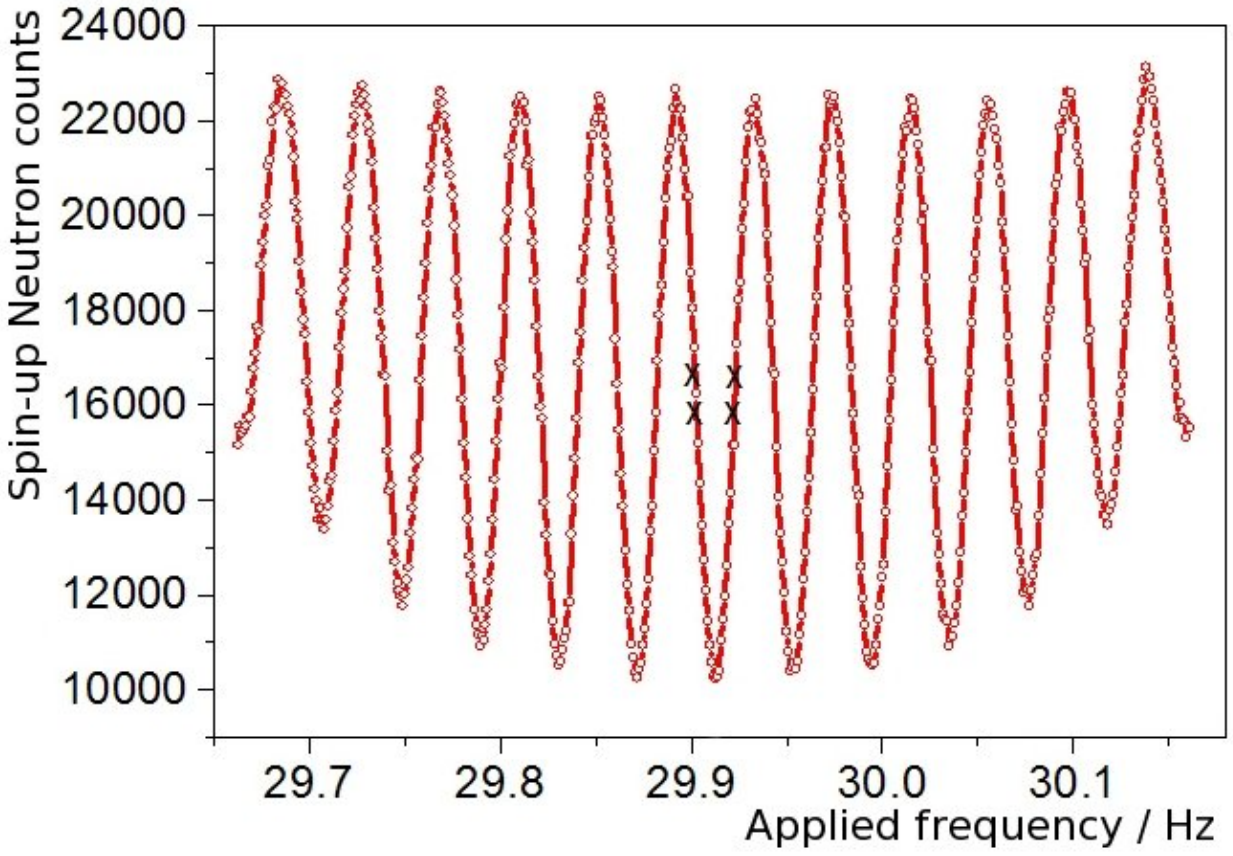

Figure 2.2: Shown is an example experimental curve showing how the neutron spin-up count varies with the frequency of an applied AC magnetic field. Clearly visible are Ramsey fringes giving a resonance frequency of 29.91 Hz. The regions of the curve the experiment operates in are labelled with 'X'.

the frequency of the rotating magnetic field is not resonant with the precession frequency during the period of free precession, a phase will be introduced between the applications of the rotating field relative to the spin. The subsequent rotation will result in a superposition of up and down spin states, which is dependent on the phase shift. Measuring the relative populations of spin-up and spin-down particles will yield the value of the phase shift and so the precession frequency. Figure 2.2 shows an experimental curve giving the number of spin-up counts from a polarised neutron source after such a cycle as a function of the applied field frequency. Clearly visible are the fringes in polarisation caused by the Ramsey resonance. The outer envelope is that of a Rabi resonance [64], the probability of a non-resonant field rotating the spin between the initial and final state with no period of free precession. The central minimum in counts is at the neutron's precession frequency.

2.2 Previous experiments

The first neutron EDM measurements using this resonance technique were performed by Ramsey himself in 1957 [65] using a neutron beam. An arrangement of magnetic fields as described above was applied to a polarised beam of neutrons. The intensity of neutrons passing through a spin analyser was measured in response to an electric field applied to the beam. A limit of $|d_n| < 4.0 \times 10^{-20}$ e cm was inferred from the measurements. Experiments continued in this manner until 1980 [66] when ultracold neutrons (UCN) were first used. The kinetic energy of these neutrons, typically $E_K < 1\,\mu eV$, is low enough to reflect at any angle from certain materials [67], e.g. Beryllium. This allows the neutrons to be stored for a significant length of time, increasing exposure to the electric field and removing large systematic effects. A summary of notable previous neutron EDM results can be seen in table 2.1.

Experiment/institution	Date	Limit / e cm	Reference
ANL	1950	3.0×10^{-18}	[62]
ORNL	1957	4.0×10^{-20}	[65]
ORNL	1967	7.0×10^{-22}	[68]
ORNL	1973	1.0×10^{-23}	[69]
ILL	1977	3.0×10^{-24}	[70]
PNPI	1980	1.6×10^{-24}	[66]
PNPI	1981	6.0×10^{-25}	[71]
PNPI	1986	2.6×10^{-25}	[72]
ILL	1990	1.3×10^{-25}	[73]
PNPI	1992	9.7×10^{-26}	[74]
nEDM	1999	6.9×10^{-26}	[75]
nEDM	2003	3.0×10^{-26}	[53]

Table 2.1: Table showing the experimental limits on the neutron EDM obtained by previous experiments.

This group has been involved in many EDM measurements and currently holds the best limit on the neutron EDM. The previous generation of the experiment produced a limit of $|d_n| < 3.0 \times 10^{-26}$ e cm. It is the aim of the cryoEDM group to improve on this limit by two orders of magnitude.

2.3 cryoEDM

cryoEDM is one of the next generation of neutron electric dipole moment searches [76] based at the Institut Laue-Langevin (ILL) in Grenoble, France. The major improvement over the previous incarnation of the experiment is its cryogenic nature, the experiment is conducted in superfluid helium-4 at approximately 0.5 K. The advantages of this become clear with reference to the equation governing the sensitivity of the determination of the EDM from the polarisation measurements, a standard result derived in appendix A. The value of the electric dipole is given by

$$d_n = \frac{\hbar(\delta N_{\uparrow\uparrow} - \delta N_{\uparrow\downarrow})}{2EN\alpha T}, \tag{2.1}$$

where δN is the difference between spin-up and spin-down counts and arrows represent the relative alignment of electric and magnetic field. The uncertainty on this value is given by

$$\sigma(d_n) = \frac{\hbar}{2E\alpha T\sqrt{N}}, \tag{2.2}$$

where E is the applied electric field, α is the product of the spin analyser efficiency and neutron polarisation, T is the neutron storage time (the period of free precession during a Ramsey measurement), and N is the total number of neutrons used to compute the result. Filling the experimental volume with liquid helium improves each of these parameters over the room temperature experiment.

Most importantly, a new UCN creation channel is available [77]. The number of UCN available directly from the reactor is quite small, more numerous higher energy neutrons (a kinetic energy of 1.4 meV) can be converted to UCN by scattering from phonons in superfluid

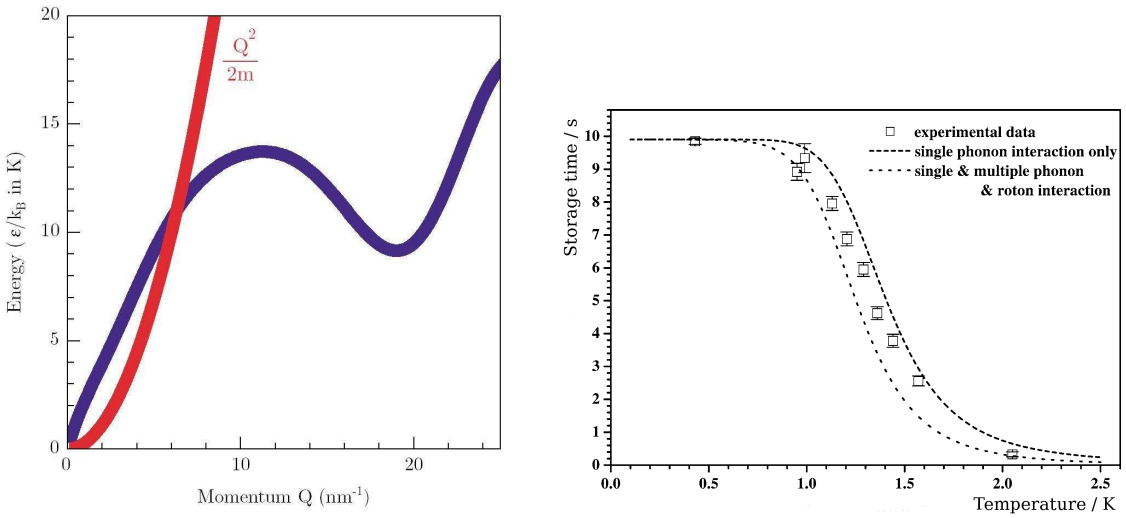

Figure 2.3: *(Left)* An illustration of the neutron (red) and phonon (blue) dispersion curves in liquid helium. At the temperature and pressure of cryoEDM these intersect at the point where a neutron's wavelength is approximately 9 Å. Therefore it is possible that neutrons with this momentum can produce a phonon with the same energy and momentum, scattering to very low energies, approximately 2×10^{-7} eV. *(Right)* The storage time of UCN as a function of temperature. Below 0.7 K the storage time is flat, above this temperature the storage time drops to zero. This is the reason cryoEDM must be operated below 0.7 K.

helium. At this energy the superfluid phonon dispersion curve intersects that of the neutron, so a neutron can resonantly scatter to an energy of approximately 2×10^{-7} eV . Figure 2.3 shows the dispersion curves for both neutrons and phonons in liquid helium, the intersection occurs at a wavelength of approximately 9 Å. To make use of this method the volume must be cooled below 0.7 K, this ensures that the thermal population of phonons is not large enough to scatter a significant number to UCN back to higher energies as shown by the neutron storage times in figure 2.3. The expectation is that this will improve the sensitivity of the experiment by a factor of 6.

A higher electric field can be supported in liquid helium than in vacuum [78]. As the EDM resolution is linearly dependent on the electric field applied this directly affects the sensitivity. The electric field is expected to be increase to at least 40 kV cm^{-1} and it is expected that this will improve the sensitivity by a factor of at least 4.

Impurities on surfaces that encounter neutrons, which can scatter the ultracold neutrons to higher energies, are cold and so less likely to scatter the neutrons significantly [79]. This

both relaxes the surface contamination requirements and also improves the lifetime of the neutrons when stored. Again, this affects the sensitivity linearly and it is expected that this will improve the sensitivity by a factor of at least 2.

The low temperature allows superconducting shielding to be used, increasing the stability of the magnetic environment, which can depolarise the neutrons. The expected improvement in sensitivity from this and other magnetic improvements is expected to be nearly a factor of 2.

2.4 Apparatus

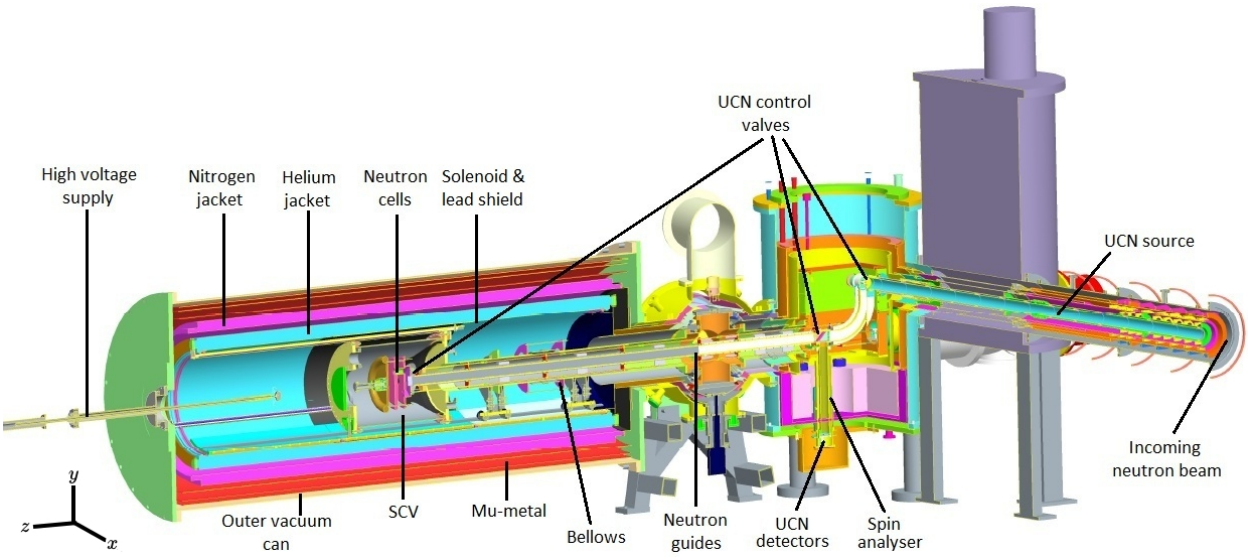

Figure 2.4: Shown is an Illustration of the cryoEDM apparatus, showing the important aspects discussed within this thesis. Neutron flow control valves, cryogenic systems and the neutron source are at the right of the image. The measurement region is centre-left within the SCV, solenoid, and lead and mu metal shields. High voltage is applied from room temperature at the extreme left.

cryoEDM consists of a Helium-3 cryostat cooling approximately 50 litres of isotopically pure superfluid helium-4. A schematic of the experiment can be seen in figure 2.4. Polarised cold neutrons enter the superfluid volume through an aperture to the neutron source where UCN are created. The UCN are trapped by copper-beryllium guide tubes, which steer the UCN through the experiment. At the far end of the experiment are neutron cells composed

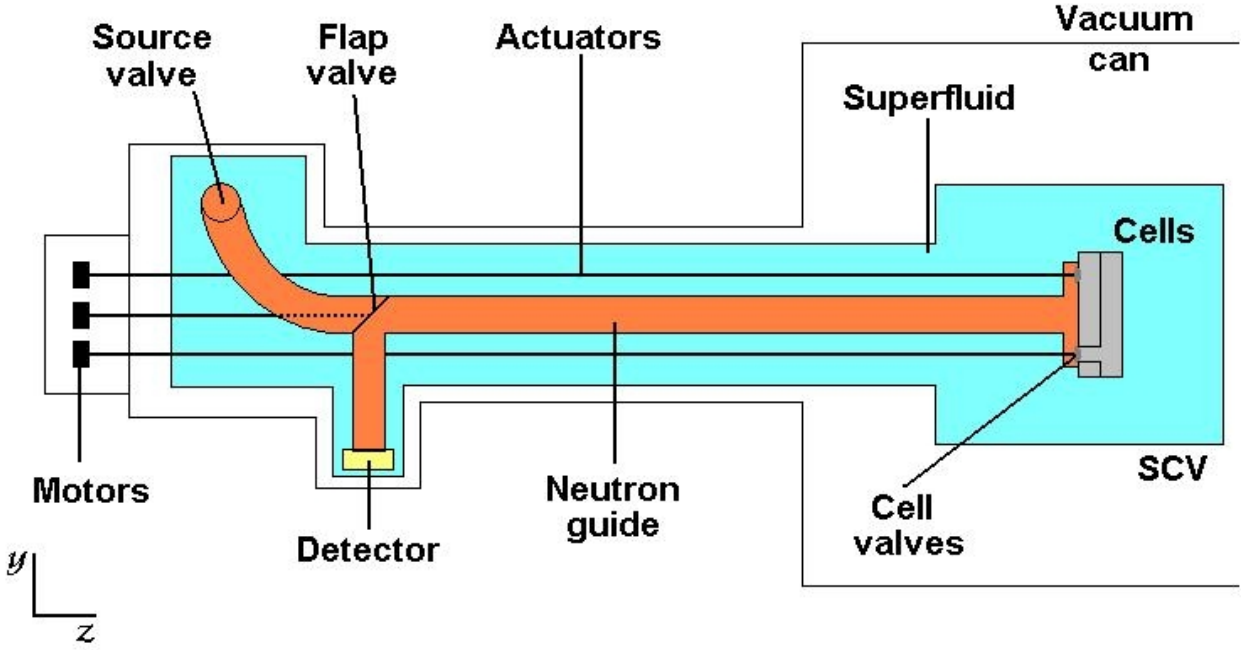

Figure 2.5: Shown is a diagram of the neutron guides, valves, actuators and motors in cryoEDM. Each cell has a plug acting as a valve, connected to stainless steel actuator rods and driven by external motors. A flap valve allows three paths for the neutrons: source to cells; source to detector; cells to detector. The flap valve is controlled in the same way as the cell valves. The source valve releases neutrons from the source tube into the guides and is operated by a separate internal motor.

of 250 mm diameter ceramic cylinders with pure beryllium electrodes acting as end caps, separated by 45 mm. In these cells Ramsey measurements are performed. Each cell is capable of storing the ultracold neutrons for of the order of hundreds of seconds during a measurement. The measurement assembly is held in a container known as the superfluid containment vessel (SCV) that is housed inside a superconducting solenoid. The solenoid is cooled by concentric liquid helium and liquid nitrogen canisters hung from the outer vacuum can by steel hangers. An electric field of 40 kV cm^{-1} is applied between the electrodes from a room temperature high voltage supply using a retractable feed-through.

A branch of the guide tube is terminated with a spin flipping solenoid, a spin analyser and neutron detectors. The analyser consists of a magnetised iron layer, which decreases the material's potential barrier for aligned spins allowing them to pass through to the detector. The potential barrier for anti-aligned spins is increased by the analyser, and so these are reflected. The detectors are purpose built lithium coated cryogenic silicon detectors [80]. An

incoming neutron reacts with the lithium layer through the reaction $^{6}\mathrm{Li} + \mathrm{n} \rightarrow {}^{4}\mathrm{He} + {}^{3}\mathrm{H}$. Either the alpha particle or triton can enter the silicon causing an ionisation current to flow. This current, the output of the detector, is read through a multi-channel analyser.

Various control values exist to contain neutrons in specific parts of the volume, e.g. the source tube while building UCN numbers, the neutron cells during a measurement, or to allow access to the detector tube. These are operated by actuator rods, which are housed in the superfluid volume, driven by room temperature motors housed outside the experiment. Figure 2.5 shows the arrangement of the neutron guides, cells, valves, actuators, motors and detectors.

2.4.1 Magnetic perturbation reduction

As cryoEDM is housed in a neutron scattering facility, there are many experiments that use magnetic fields in their operation. If large, these fields can alter the magnetic field over the measurement volume. To reduce the experiment's susceptibility to external magnetic fields three methods of reducing the magnetic perturbations are employed, mu-metal shields, a superconducting shield, and a compensation system. Figure 2.6 shows the arrangement of these within the experiment to scale. A detailed description of each method will now be given;

1. Mu-metal shields - a close to 4π coverage of the experimental volume with three layers of a high magnetic permeability metal, offering a more favourable path for stray magnetic field lines through the material than through the centre of the volume. The shield is formed from three cylinders co-axial with the experimental volume, each with a length of 3.2 m and diameters 1.018 m, 1.098 m, and 1.178 m. The outer cylinder is fitted with end caps to close the shields at either end. The cylinders have holes to allow the steel hangers supporting the nitrogen and helium tanks to attach to the outer vacuum can. The end caps have holes for the neutron guide and high voltage supply to pass through the shield. The HV end cap has an additional top-hat shaped cap to

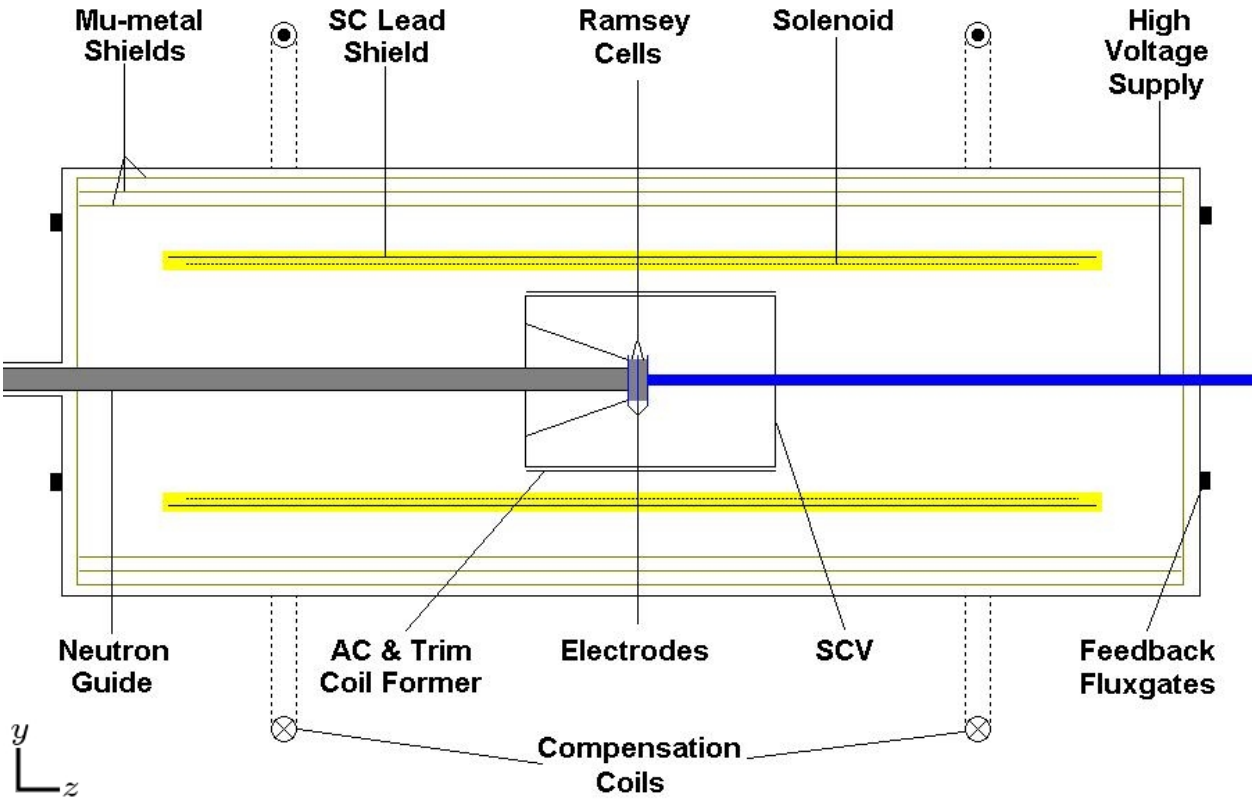

Figure 2.6: Shown is a scale diagram of the shielding elements in cryoEDM. The three concentric shielding systems are shown coaxially with the SCV and neutron cells. The neutron guide and solenoid are held in a helium tank, which is free to move. Also shown is a carbon fibre former that holds the AC field coil and trim coils.

close the shields further. Each layer of the mu-metal has the ability to be demagnetised individually.

The mu-metal is estimated from simulations to have a dynamic shielding factor of order 10 against axial perturbations and of order 10^2 against transverse perturbations.

2. Superconducting lead shield - a cylinder of lead, with a superconducting transition temperature of 7.2 K, thermally anchored to a liquid helium bath. Flux penetration through the lead is zero due to the Meissner effect [81]. The cylinder is co-axial with the experimental volume, 2.68 m long, 0.724 m in diameter, and 1 mm in thickness.

   The shielding factor of a superconducting cylinder is given by $\exp\left(-3.83\frac{z}{R}\right)$ and $\exp\left(-1.84\frac{z}{R}\right)$ for the axial and transverse components respectively [82], and so the shielding is estimated to be 1×10^6 against axial perturbations and 1×10^3 against

transverse perturbations.

3. Compensation coils - two dipole coils, octagonal in shape with side length 0.6 m, wound externally around the experimental volume and 2 m apart. The coils are co-axial with the experimental volume, have 78 turns and are powered to minimise the axial component of the error flux over the volume as read by six external fluxgates. The cancellation of much of the external perturbation produces an effective shielding.

   The shielding factor is estimated to be approximately 10.

The design of shielding is such that the combination of all three methods should produce a total shielding factor that is the product of the individual shields [83], giving shielding of 10^8 against axial perturbations and 10^5 against transverse perturbations. Figure 2.7 shows a photo of the cryoEDM apparatus with the vacuum can open and the mu-metal cap removed giving a view of the concentric shielding.

2.4.2 The solenoid, AC & trim coils

Figure 2.6 also shows the position of the solenoid, and a carbon fibre coil former that supports the magnetic coils used in the operation of the experiment.

The superconducting solenoid is of length 2.627 m, has a diameter of 0.684 m, and is coaxial with the neutron cells. It is formed from 10800 turns in two layers of 0.48 mm niobium-titanium wire. The two ends of the solenoid are joined with a superconducting link connected to a heater that acts as a switch to enable running of the solenoid in a persistent mode. The solenoid has a transition temperature of 9.3 K.

The carbon fibre former holds multiple coils. There is a saddle shaped coil designed to produce a uniform magnetic field vertically along the experimental y-axis over the neutron cells, which is used to apply the AC fields required in a Ramsey cycle. The former also supports 19 “trim” coils of various designs and orientations used to improve the homogeneity of the experimental field.

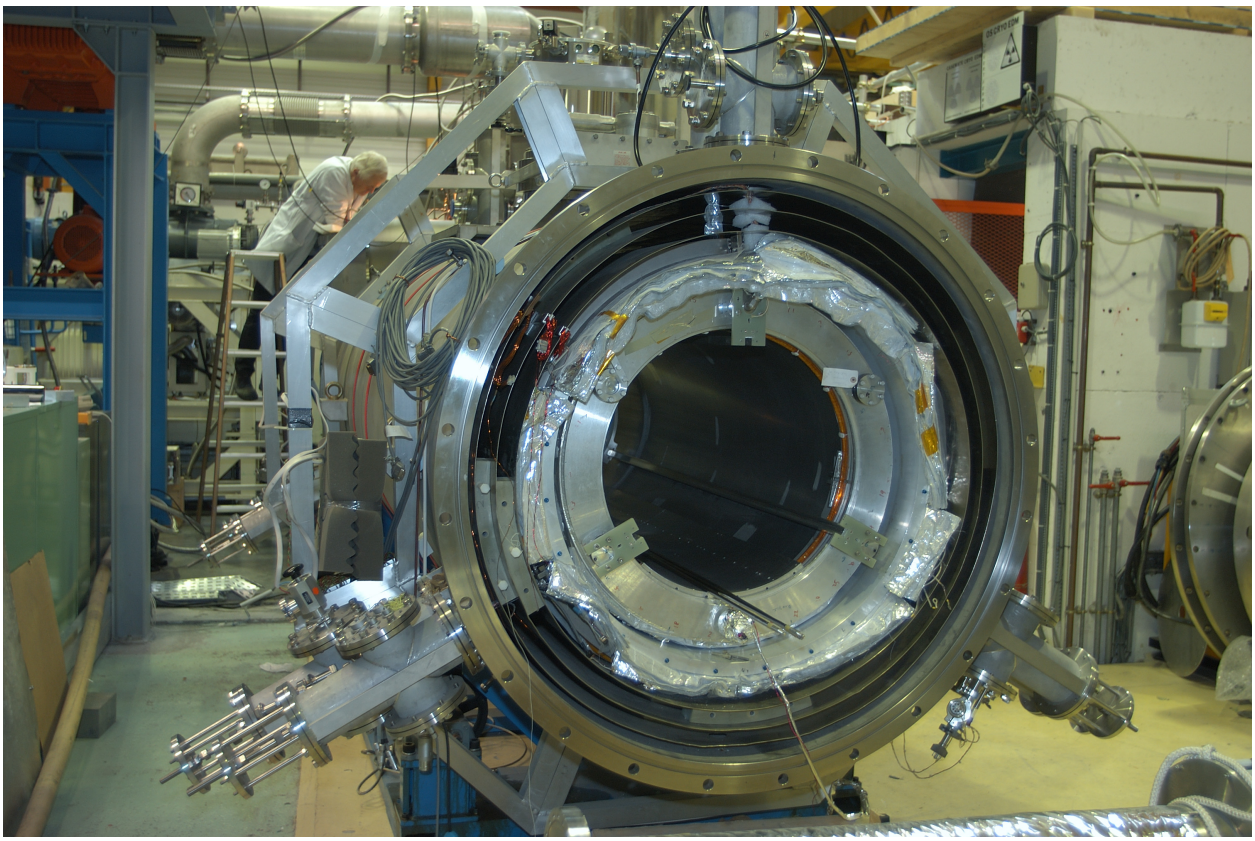

Figure 2.7: Shown is a photo of the cryoEDM magnetic shields. The cylinders from the centre out are the helium tank containing the superconducting shield and solenoid, the nitrogen tank, three layers of mu-metal, and the outer vacuum can. Outside of the vacuum can an octagonal compensation coil can be seen.

2.5 Measurement cycle

Polarised ultracold neutrons are created in the superfluid Helium-4 source volume. Once enough UCN have collected they are transferred to the neutron cells in the SCV by opening the source valve and cell valves. The neutrons then diffuse analogously to a gas through the guide tubes. Once the neutron density in the cells is maximal, the cell valves are closed trapping the neutrons within. The source begins refilling immediately to minimise the dead time between measurements.

The cells are exposed to the same magnetic field, but different electric fields are applied across them (typically $+E_0$, 0, and $-E_0$) to reduce systematic uncertainties. A resonance cycle is then applied to the neutrons and a cell valve opened releasing the neutrons. The valve to the detector tube is also opened and the neutrons diffuse to the analyser. The

numbers of neutrons in the spin-up polarisation state is measured directly by the detector response. Once this state is measured, the anti-aligned spin state is then flipped using the spin flipping solenoid and passed through the analyser again to count their numbers. This process is repeated for each neutron cell.

By this point, the number of neutrons in the source volume should have built up enough that a measurement on the next batch of neutrons can begin.

2.5.1 EDM determination

Once the relative populations of spin-up and spin-down neutrons have been determined it is possible to obtain a value of the neutron EDM. A brief description of the method of determining this value will now be given. Many constant systematic uncertainties can be removed if the EDM is deduced from the quantity

$$\hbar(\omega_{\uparrow\uparrow} - \omega_{\uparrow\downarrow}) = |2\mu_n|(B_{0\uparrow\uparrow} - B_{0\uparrow\downarrow}) - 4d_nE, \tag{2.3}$$

where $\omega_{\uparrow\uparrow}$ and $\omega_{\uparrow\downarrow}$ are the precession frequencies of the neutrons observed with the magnetic and electric fields aligned and anti-aligned respectively, and $B_{0\uparrow\uparrow}$ and $B_{0\uparrow\downarrow}$ are the modulus of magnetic fields applied to the neutrons when the magnetic and electric fields aligned and anti-aligned respectively. In the notation above the relative alignment of the magnetic and electric fields is averaged over both combinations of field orientations that produce it, i.e. $\uparrow\uparrow$ is the average of $\uparrow\uparrow$ and $\downarrow\downarrow$, and $\uparrow\downarrow$ is the average of $\uparrow\downarrow$ and $\downarrow\uparrow$. So the EDM is determined from a shift in the precession frequency between measurements of aligned and anti-aligned field configurations, this eliminates systematics that are uncorrelated to the electric field direction. At this stage any corrections for changes in the magnetic field must be applied. When using the theoretical curve to relate neutron counts to precession frequency equation 2.3 gives a value of the EDM and uncertainty stated in equations 2.1 & 2.2.

The neutron EDM and its uncertainty can now be determined from sets of spin statistics with correction for magnetic field changes.

2.5.2 Operating points

While the position of the central minimum of the resonance curve will yield the neutron precession frequency, it is impractical to obtain a full resonance curve in each configuration. Once a resonance has been measured the applied frequency is set such that the experiment is conducted in a region of the resonance where $\frac{\partial N_{\uparrow}}{\partial \omega}$ is maximum, thus producing the largest change in neutron counts for change in precession frequency. As the curve is symmetric, by changing the frequency of the applied field by a small amount the operating point can be moved to the point on the resonance curve where the gradient is exactly opposite. Measurements taken at these points can be combined to cancel certain systematic effects, which are associated with changing the polarity of the electric field. For example if the signal from the neutron detector is affected by the field polarity to cause a shift in neutron counts, the opposing gradients would produce a false EDM in opposite directions, which can be cancelled upon summation. These points are marked with an 'X' in figure 2.2.

2.5.3 Systematic effects

A number of systematic effects can affect the results of an EDM measurement. The most serious are those which are correlated with the direction of the electric field and can escape the precautions used in sections 2.5.1 and 2.5.2. The systematic effects have been well studied in the lifetime of this experiment and are considered manageable. A short list of the major systematics will be given.

Hysteresis of magnetic shields Magnetic shielding is typically constructed out of a magnetisable material and so can have a remnant magnetisation. This is caused due to the hysteresis of the magnetisation in an applied field. The magnetisation can distort an applied field, which, if the magnetisation is dependent on the electric field direction, can appear as a false EDM signal.

In the room temperature experiment, changes in the electric field induced currents in the magnetic field coil causing disturbances in the stabilised current. This was fed back

to the coil producing an unstable magnetic field through the mu-metal surrounding the experiment. These fields induced a hysteresis in the mu-metal shielding correlated with the electric field direction. This varied the magnetic field with position and, due to its highly localised nature, the magnetometry could not be used to fully compensate for the magnetic field changes. The effect gave a false EDM signal of about 1×10^{-25} e cm when using a 12 kV cm^{-1} electric field.

Such fields have been reduced significantly in the cryogenic experiment by the inclusion of a superconducting solenoid to be used in persistent mode. This bypasses the need for a current stabiliser, which caused the unstable fields. Further, the inclusion of a superconducting shield reduces the effect of any hysteresis in the mu-metal to a point where false EDM signals of 1×10^{-30} e cm are expected, well below the sensitivity aim of the experiment.

Cell displacement No magnetic field can be perfectly homogeneous. The neutron cells are subjected to a magnetic field that depends upon their position. If the position changes then the value of the magnetic field changes also. If the changes correlate with the electric field direction a false EDM is induced.

The cells are subjected to large electric fields in the region of 50 kV cm^{-1}. This can exert a large force on the cells, of order 1 Newton, and could in principle move the cells. If the geometry was symmetric then the forces should not vary with the sign of $\boldsymbol{E}$, however asymmetries are expected and must be allowed for.

The apparatus is designed such that asymmetries at the 5 % level will produce displacement of no more than 10 nm. The rigidity of the system required to overcome this effect must be weighed against the damage that could be caused due to thermal contraction if the system is too rigid. With this design, and anticipated radial magnetic gradients of approximately 3 nTm^{-1}, the false EDM signals induced by such movements are estimated to be at the 1×10^{-28} e cm level.

High voltage ripple The high voltage used to create the electric field across the experiment is produced by a Cockroft-Walton [84] generator. This requires an alternating

driving voltage, which is produced at 10 kHz. The driving voltage is generated from the mains supply, which operates at 50 Hz. It is possible that one, or both, of these frequencies will be transmitted through the generator in small amounts, causing a ripple on the otherwise DC voltage.

This small AC component will produce an alternating displacement current and hence an alternating magnetic field in the x-y plane, which will create a shift in the Larmor frequency of the neutrons [85]. The 10 kHz component shifts the Larmor frequency down. A 0.6 % ripple in the high voltage at this frequency would cause a false EDM signal of 1×10^{-28} e cm to be observed. The specifications of the generator limit the ripple from both frequencies to 0.04 % and so this frequency should not concern the experiment. The 50 Hz component will shift the Larmor frequency up, a 0.08 % ripple in the high voltage at this frequency would cause a false EDM signal of 1×10^{-28} e cm to be observed. Again the specification of the supply is better than this requirement.

$\boldsymbol{E} \times \boldsymbol{v}$ Particles moving with a velocity $\boldsymbol{v}$ through an electric field $\boldsymbol{E}$ will experience a magnetic field $\frac{\boldsymbol{E} \times \boldsymbol{v}}{c^2}$. As this field will reverse with $\boldsymbol{E}$ it is liable to produce a false EDM signal.

The measurement is not performed on a single neutron but on a batch of neutrons, and so the average motion through the electric field is the concern. This motion can be broken down into the motion of the centre of mass of the neutrons, and motion around the centre of mass. Motion of the centre of mass is expected to be caused by warming of the neutrons counteracting their displacement under gravity, and will be vertically upward. Movement in the x-y plane produces a magnetic field also in the x-y plane, which gives a change in magnitude of the magnetic field that does not depend on the sign of the electric field. This will give a false EDM of approximately 2×10^{-29} e cm but can be removed by the procedure in section 2.5.1.

Motion around the centre of mass is likely to be caused due to off axis filling of the neutron cells and the torque around the centre of the cell due to gravity. This motion decays quickly after filling. Simulations show that the false EDM produced by such a

motion would be 2.7×10^{-27} e cm, this is far too large. However, a 2.5 second delay between filling and measurement allows the motion to decay to the extent that a false EDM would be 2.7×10^{-28} e cm, which is far more acceptable.

Geometric phase In addition to the field above, created by movement through the electric field, there is a magnetic field in the x-y plane caused by field gradients expected in the axial direction. The product of these two fields create a field that the neutrons experience as rotating. The interaction of the neutrons with this field is complicated and is treated fully in [86], here the analytic result is quoted. The resulting false EDM is given by

$$d_e = \frac{\hbar}{6}\frac{dB_0}{dz}\frac{1}{B_0^2}\frac{\langle v^2 \rangle}{c^2}, \tag{2.4}$$

where B_0 is the applied magnetic field, and v is the velocity distribution of the neutrons.

With the magnetic gradients expected, $3\,\mathrm{nT\ m^{-1}}$, a false EDM of 2.4×10^{-27} e cm would be produced. This is unacceptably large and sets of coils are placed around the experimental volume designed to trim the field such that the gradients are as low as $0.1\,\mathrm{nT\ m^{-1}}$. This leads to a false EDM of 8×10^{-29} e cm with the magnetic field trimmed.

Additionally, as the false EDM from this effect scales as $\frac{1}{B_0^2}$, a larger applied magnetic field is beneficial. Although this is true only up to the point where the growing magnetic field instabilities dominate the systematics.

2.5.4 Magnetic specification summary

The false EDM induced by a change in the magnetic field during a measurement considering both electric field directions is

$$\sigma(d_n) = \frac{\mu_n \sigma(B)}{2E}, \tag{2.5}$$

where μ_n is the magnetic moment of the neutron and $\sigma(B)$ are the fluctuations in the magnetic field. This will dominate the counting statistics when

$$\sigma(B) > \frac{\hbar}{\mu_n \alpha T \sqrt{N}}. \tag{2.6}$$

For the expected values $\alpha = 0.9$, $T = 300\,\mathrm{s}$, and $N = 4.5 \times 10^5$ this is approximately 60 fT. The magnetic field stability must be significantly below this for the experiment to be limited by neutron statistics.

This, and systematics described in the previous section, give the various magnetic specifications required to meet the aimed for experimental sensitivity. The following will be a short summary of the necessary magnetic properties required in order to reach the limit of 10^{-28} e cm:

- Fluctuations of the magnitude of the magnetic field during a 300 second measurement cycle must be below 60 fT.
- Axial magnetic gradients must be below 0.1 nTm^{-1}
- Radial magnetic gradients must be below 3 nTm^{-1}

With such strict magnetic requirements a suitable magnetometer is required to ensure they are upheld.

2.6 Choice of magnetometer

As the magnetic properties need to be monitored carefully, the correct choice of magnetometers to accomplish this is important. The previous generation of the experiment, operated at room temperature, used a mercury co-magnetometer [87] where polarised ^{199}Hg atoms occupied the same volume of neutrons and underwent the same measurement cycle. The polarisation of the ^{199}Hg atoms was continuously measured through the absorption of light from a ^{204}Hg lamp passing through the cell, and read with a PMT. The EDM of mercury

having already been measured to be far below the measurement aim of the neutron EDM [58], there would be no significant electric contribution to the Ramsey curve. This system could not operate in the cryogenic environment of the new experiment and so an alternative magnetometer is required. Several options were available for a cryogenic magnetometer. The options considered are:

Fluxgates Their operation described below, fluxgates are magnetometers that can operate in cryogenic temperatures and have resolutions of the order nanotesla, which is insufficient for the needs of the experiment. They also introduce magnetic field perturbations in the regions around the device due to magnetic materials in the sensor and modulation currents applied to them during read-out. These issues make fluxgates unsuitable as the primary magnetometer in this experiment.

UCN If not experiencing an electric field, the measurement procedure described above on UCN will provide a measure of the magnetic field. This will automatically be at the precision required to correct for fields experienced by the neutrons. However, this arrangement, by definition, can not be used as a co-magnetometer as they must not experience the applied electric fields required for a measurement, and so would have to be housed in separate cells. To produce a good picture of the magnetic field from isolated measurements requires frequent sampling. Having many neutrons cells devoted to magnetometry is infeasible, and so the spacial resolution would be low. Additionally, UCN measurements do not provide a real time measure of the magnetic field, only the volume and time average after each measurement is available. This means certain diagnostics would be difficult to perform using UCN alone.

^{3}He co-magnetometer This is intended for use in a proposed EDM measurement [88]. This is operated in a similar manner to the mercury magnetometer, however the polarisation of the ^{3}He atoms is monitored by a separate magnetometer, a SQUID system. The inclusion of a second magnetometer for read-out adds significant extra complication to the system. A source of polarised ^{3}He would also need to be developed, again adding complications. Additionally, since ^{3}He is a strong absorber of neutrons through

the reaction $n + {}^3\mathrm{He} \rightarrow {}^4\mathrm{He}$, careful control of the amounts of ^{3}He would need to be put into place to ensure that both neutrons and ^{3}He survive long enough for the measurement to be performed.

SQUIDs Their operation described below, SQUIDs are magnetometers that have been shown to have resolutions as good as 0.3 fT $\mathrm{Hz}^{-\frac{1}{2}}$ [89] and to have measured magnetic field changes as low as $5 \times 10^{-18}\,\mathrm{T}$ over several days [90]. The downside of using a SQUID sensor is that it measures the flux only through its sensing region. This makes it an isolated measurement and so a number of SQUIDs are required to sample the field enough to give a good spacial resolution. As SQUIDs are commercial products this causes no problem, as an arbitrary number of SQUIDs can be installed. The sensing regions are loops of wire and so many can be accommodated in the experimental volume.

Due to its advantages, a SQUID system has been chosen to be used as the primary magnetometer. To complement this, fluxgates are installed to give a rough estimate of the field, and two magnetometry UCN cells are planned for fine adjustments.

2.7 Magnetometry equipment

The experiment has two types of magnetometers available for diagnostics, SQUIDs and fluxgates. Their basic operation is given below and their arrangement in the cryostat follows.

2.7.1 Fluxgates

A fluxgate consists of a core of high permeability metal wound with two coils. An alternating driving field is applied through one coil, and an induced response measured on the second coil. The magnetisation curve of a material is not globally linear, so the induced response will be asymmetric unless the material has zero net magnetisation. An offset in the driving field is adjusted until the response is symmetric. At this point the offset field is exactly

cancelling any applied field, and so its negative is used as the output of the magnetometer.

The fluxgates used in cryoEDM have a maximum sensitivity of 0.1 nT, a maximum range of $\pm 20\,\mu$T, and a maximum bandwidth of 10 Hz [91]. Additionally the driving field would be a significant magnetic perturbation and so fluxgates can not be placed in the immediate vicinity of the neutron volumes.

The fluxgates use both off-the-shelf sensors and readout electronics, and so have tabulated calibrations.

2.7.2 SQUIDs

A superconducting quantum interference device (SQUID) is a sensitive flux-to-voltage converter. It operates due to the interference of two currents passing around a superconducting loop, where each branch of the loop contains thin insulating barriers, Josephson junctions [92]. Figure 2.8 shows the basic layout of a SQUID loop and the currents involved. A supercurrent can tunnel across a thin barrier between two superconductors with the electrons pairs remaining correlated. The resulting current is modulated sinusoidally by the phase difference between the states either side of the junction, i.e $I(t) = I_c \sin(\phi(t))$, where I_c is the critical current of the junction, the lowest current through the junction that will not produce a potential across the junction and depends on the construction of the junction, and $\phi(t)$ is the phase difference across the junction. In a loop of two junctions the total current through the loop would be the sum of the two junctions' currents,

$$I = I_c \sin(\phi_1) + I_c \sin(\phi_2) \tag{2.7}$$

$$I = I_c \sin(\phi_0 - \frac{\delta\phi}{2}) + I_c \sin(\phi_0 + \frac{\delta\phi}{2}) \tag{2.8}$$

$$I = 2I_c \sin(\phi_0) \cos(\frac{\delta\phi}{2}) \tag{2.9}$$

where ϕ_0 is the average phase across the two junctions and $\delta\phi$ is the phase around the loop.

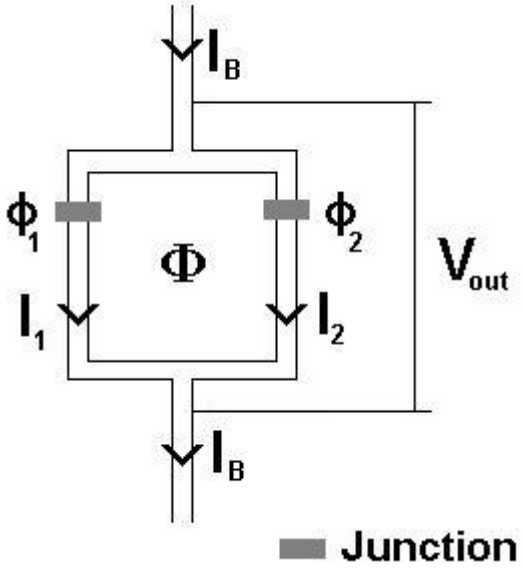

Figure 2.8: Shown is a schematic of a DC SQUID, with a bias current passing through a superconducting loop, with each branch containing a Josephson Junction that the supercurrent can tunnel through. The current through each branch depends sinusoidally on the phase across the junction ϕ. The total current through the loop is modulated by the phase difference around the loop. This phase difference is proportional to the magnetic flux through the loop, Φ. A voltage is produced across the loop, which is taken as the SQUID output, this voltage has a sinusoidal response to the flux through the loop.

In the presence of a magnetic field a phase is introduced around the loop that is proportional to the flux through the loop [93], an effect which leads to flux quantisation in the absence of the junctions. The difference in phase around the loop is $\delta\phi = 2\pi\frac{\Phi}{\Phi_0}$, and the current through the loop modulates as $I = 2I_c \sin(\phi_0)\cos(\pi\frac{\Phi}{\Phi_0})$. The SQUID is operated so that the bias current is larger than the critical current and so a potential difference is present across the loop, given by $V = R(I_B^2 - I^2)^{\frac{1}{2}}$. Due to the modulation of the current the voltage across the loop is modulated by the flux passing through the loop, this modulation is sinusoidal with a period of one flux quantum. This periodic nature of the SQUID means that a static SQUID cannot measure the absolute field, there will be an unobserved offset equal to an integer number of flux quanta from the real field. Figure 2.9 shows an illustration of how this can be used to produce a magnetometer.

This sinusoidal output can cause difficulty in extracting a meaningful magnetic signal. This can be rectified by use of a feedback coil wound with the SQUID coil to keep the flux

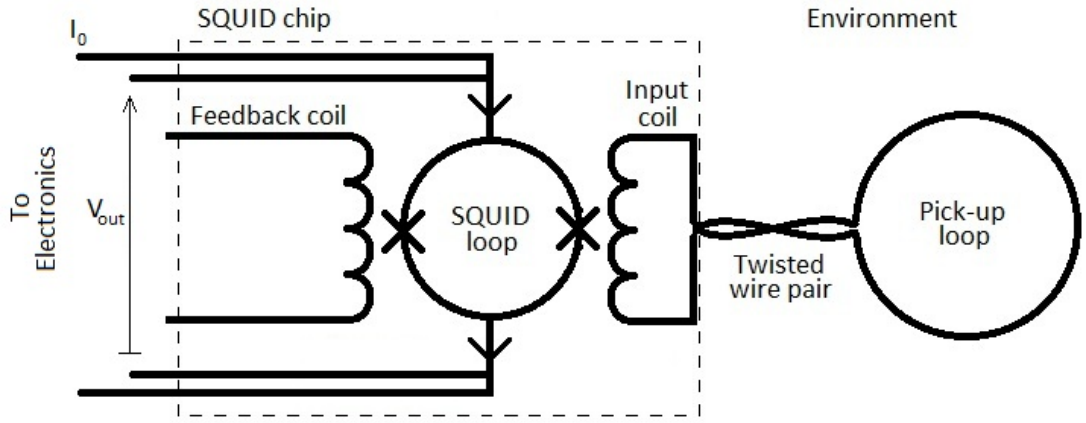

Figure 2.9: Schematic of the components of a SQUID used as a magnetometer. The SQUID loop is powered by a bias current, and the output voltage is measured across the SQUID loop. This output is modulated by the flux through the SQUID loop with a period of one flux quantum. Flux is applied to the SQUID loop through an input coil, which is typically connected through a twisted-wire pair to a pick-up loop, all of which is superconducting. The pick-up loop acts as a flux transformer coupling flux from the environment into the SQUID loop. A feedback coil allows offsets to be applied to the SQUID loop to keep the field through the SQUID loop constant.

through the SQUID coil constant, and in a region where the SQUID response is linear in applied field. This region is referred to as a "working point". The change in feedback field will be linear in the applied field and is taken as the magnetometry signal. Figure 2.10 shows an illustration of this feedback procedure. Limitations in the application of the feedback and its effects on the SQUID signal will be discussed in section 2.7.2.1.

A small modulation current of 100 kHz is also sent through feedback coil and so coupled into the SQUID. A lock-in amplifier attached to the SQUID output then extracts the component of the output with the same modulation.

The size of the SQUIDs mean it would be impractical to measure the magnetic field applied directly through the loop. The SQUIDs measure the field applied to them through a superconducting multi-turn input coil wound with the SQUID coil. This input coil is attached to a superconducting "pick-up" loop, which acts as a flux transformer to couple magnetic field observed in the environment on to the SQUID chip.

The SQUIDs used in cryoEDM can have a sensitivity in the region of 0.1 fT and a

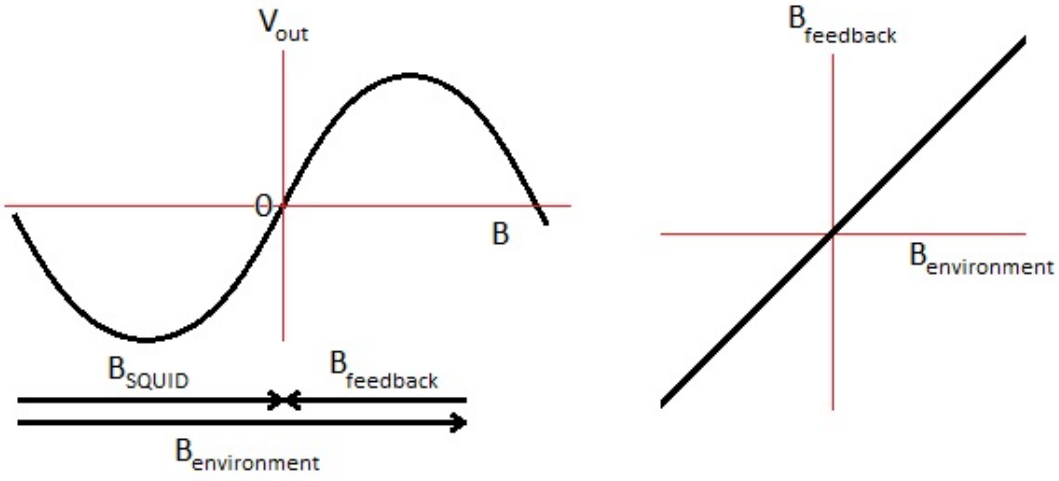

Figure 2.10: An illustration of the operation of the feedback system. *(Left)* The field through the SQUID loop B_{SQUID} is maintained at a level to produce no output. The feedback field $B_{feedback}$ is adjusted to offset the environmental field $B_{environment}$ such that B_{SQUID} is at this point. Any small changes in $B_{environment}$ will cause a change in output of $\frac{dV}{dB}\delta B$, the feedback electronics induces a field change of opposite magnitude through the feedback coil to return the output to zero. Provided the field changes are small $\frac{dV}{dB}$ can be taken as constant. This produces a linear relationship between $B_{environment}$ and $B_{feedback}$ *(Right)*. The feedback voltage is taken as the output of the magnetometer.

bandwidth of 100 kHz. They too suffer from creating significant magnetic perturbations due to modulation currents, but can be effectively isolated from the neutron volume by the pick-up loops.

The SQUIDs themselves are commercially available but some components have been modified to ensure correct operation at 0.5 K. The readout electronics for the SQUIDs is a combination of off-the-shelf systems and electronics and cabling designed and built by the Oxford group. Each pick-up loop requires careful calibration.

2.7.2.1 Feedback artefacts

The finite voltage range and non-zero slew rate of the SQUID feedback system gives rise to three types of artefacts in the data coming directly from the hardware. These are resets, flux jumps, and loss of lock.

Resets The simplest feedback artefact introduced by the above system is a reset. The

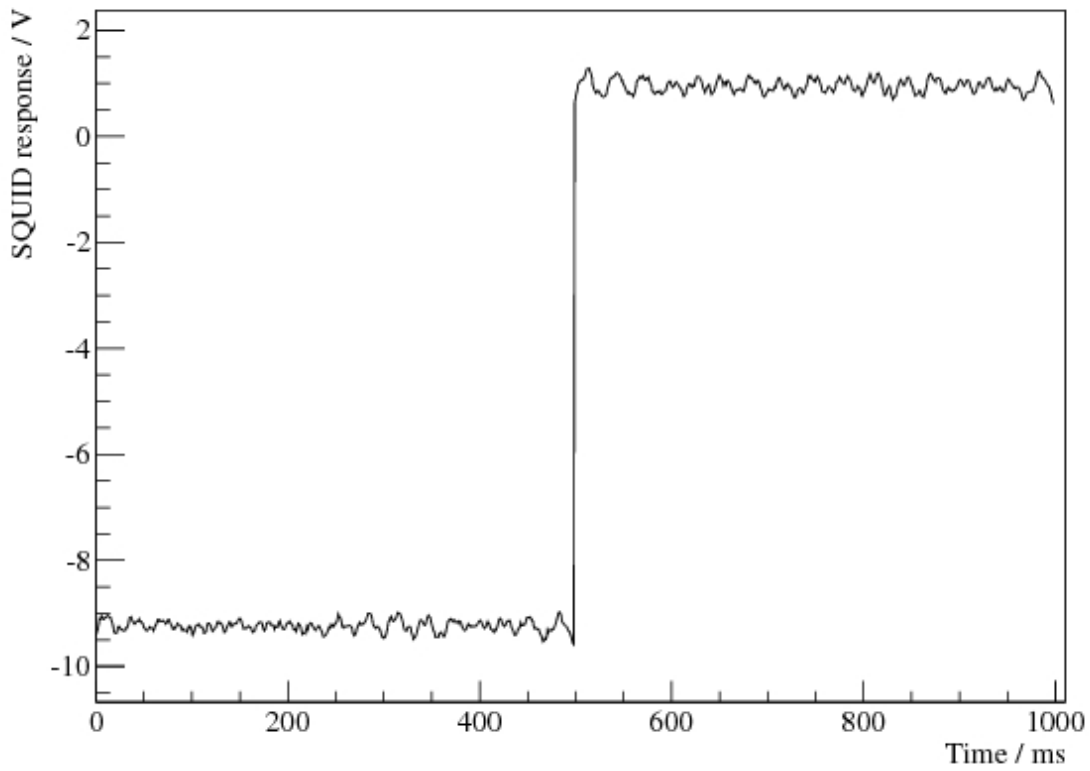

Figure 2.11: Example of a reset with a discrete change in voltage of 10.2 V seen on a SQUID. The FWHM of the noise on this SQUID is 0.4 V, so the reset is easily distinguished from noise.

feedback system can apply voltages to the feedback coil in the range -10 V to +10 V, and so the magnitude of the feedback field is limited. Once this is at maximum any further increase in the applied field can not be compensated for and the SQUID output will leave the linear region. To overcome this, as the feedback signal approaches the limiting value the feedback is zeroed and the system attempts to find a linear region of the SQUID response requiring a lower feedback field. In normal operation the feedback range covers more than one flux quantum, so a linear region with a lower feedback field can be found. The result of this resetting is a large discontinuity in the SQUID signal. The change in voltage at the discontinuity corresponds to an integer number of flux quanta.

Figure 2.11 shows an example of a reset in the feedback system as the output tries to exceed -10 V.

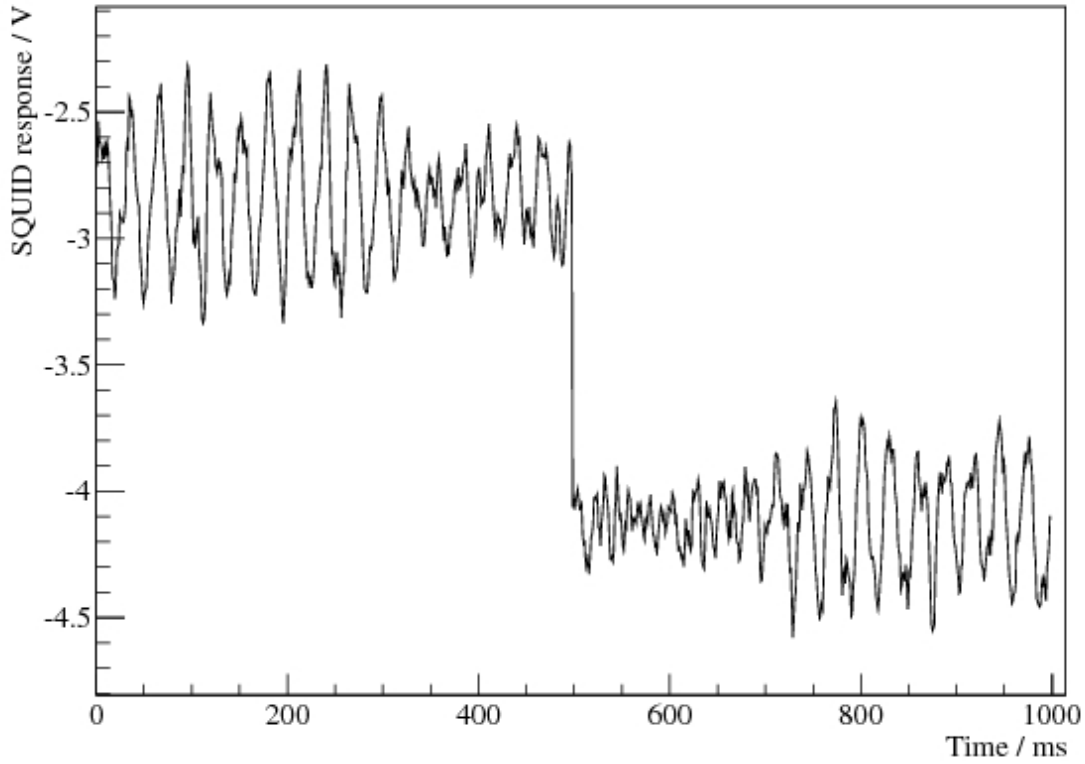

Figure 2.12: Example of a single flux quantum loss with a discrete change in voltage of -1.3 V seen on the same SQUID as in figure 2.11. The FWHM of the noise on this SQUID is 0.4 V, so the flux jump can be distinguished from noise. However there are occasions when the noise is larger than the size of change corresponding to a flux quantum, and so it can be difficult to distinguish artefacts from the noise.

Flux jumps A more problematic artefact is introduced by the non-zero response time of the feedback. If the rate of change of applied field exceeds the slew rate of the system, the feedback momentarily "loses lock" and can settle on a different working point. Specifically, if the difference between the applied and feedback fields is greater than half a flux quantum [94], then the gradient of the response curve tends to move the feedback away from the original working point. The result is a smaller discontinuity, again the voltage change corresponds to an integer number of flux quanta, although typically only one flux quantum. Flux jumps are typically caused by very rapidly changing signals, or poor signal-to-noise ratios.

Figure 2.12 shows an example of a flux jump as the noise temporarily exceeds the slew rate in the feedback system.

Losing lock A serious issue is caused if the applied field exceeds the slew rate of the feedback system for an extended period of time. The feedback system can not compensate for changes in the magnetic field but still attempts to. This causes the SQUID to become unstable and vary wildly. As the magnetometer output no longer has a direct relation to the applied field, there is no way to compensate for this effect and retrieve the magnetic signal seen during this period.

2.7.2.2 SQUID noise

The power spectrum of magnetometer noise tends to take two distinct forms: a white noise region, constant in frequency, which is the fundamental noise of the magnetometer; and a noise that scales as $\frac{1}{f}$. Noise of the latter type appears in many condensed matter systems [95][96] and SQUIDs are no exception [97], much work is done by manufactures to reduce this and can push the region where the $\frac{1}{f}$ contribution dominates to below about 1 Hz. Noise of the form $\frac{1}{f^{\alpha}}$ also enters from the environment which, if large, can swamp the noise produced by the SQUID itself. These noise features are in addition to any noise peaks from specific sources. An example noise spectra, here defined as the square root of the power spectra, displaying all these features can be seen for 5 SQUIDs without pick-up loops attached in figure 2.13 as measured by S. Ingleby [98].

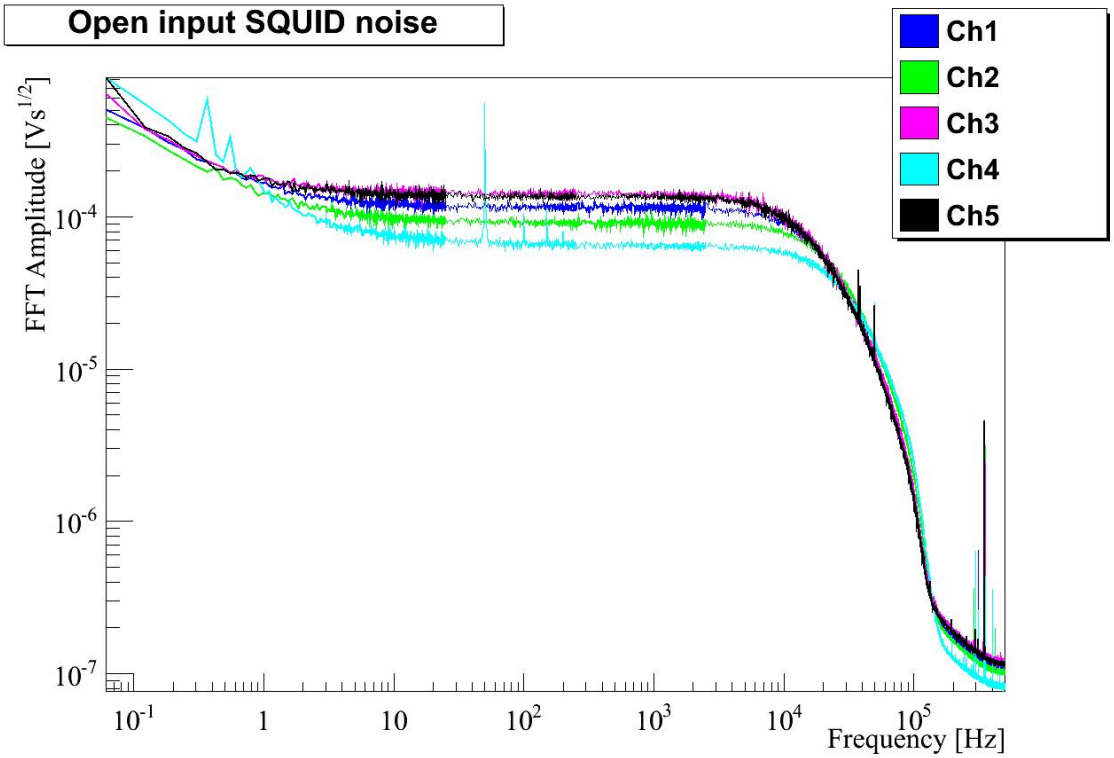

Figure 2.13: An example of SQUID noise spectra as measured by S. Ingleby [98] on 5 SQUIDs without pick-up loops attached. Clearly visible is a $\frac{1}{f}$ region below 1-10 Hz, and a white noise region above. Also peaks at particular frequencies can be seen, e.g. at 50 Hz. The measurement was taken through a 10 kHz low pass filter.

2.7.3 Arrangement

Of the two types, the experiment has the following magnetometers. 12 SQUIDs, 5 cryogenic fluxgates, 3 external fluxgates arranged in a triple axis set-up, and 6 fluxgates used in a feedback system to control the compensation coils. Figure 2.14 shows a schematic of how the magnetometers are generally arranged inside the cryostat.

Much of the work in this thesis concerns data taken during a 2010 data run. Figure 2.15 and table 2.2 gives the position and orientations of the magnetometers installed during the data run. While other SQUIDs were installed in this run, either no calibration coils were installed with them or they could not be operated, so they are not useful as magnetometers and shall not be listed. It is worth noting that the pick-up loops of SQUIDs 1, 4, and 11 are of the same dimensions, the pick-up loop attached to SQUID 5 has approximately one fifth of the area. SQUID 5 has a diameter of 16 mm and a sensing area of 2×10^{-4} m^2, SQUIDs

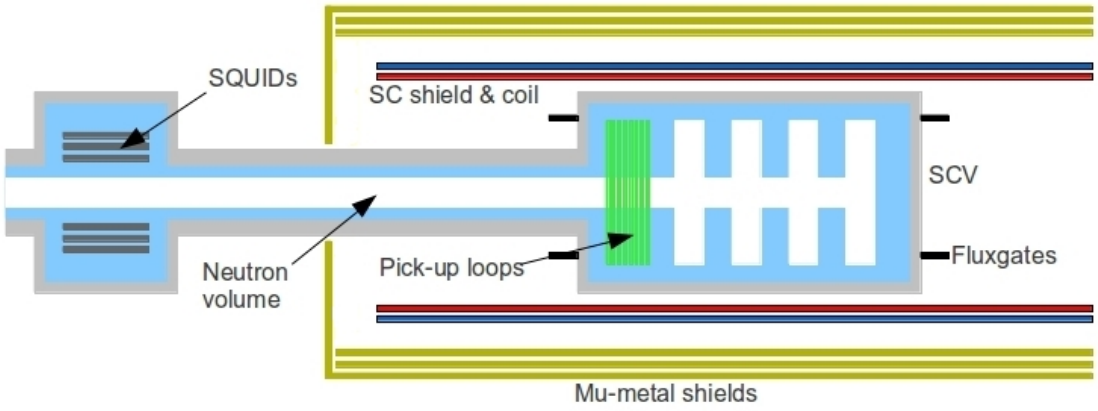

Figure 2.14: Schematic of the magnetic components of cryoEDM. Shown are the SQUIDs located on the far left outside of the shielding, and the pick-up loops coupling magnetic flux into the SQUIDs near the neutron cells. The fluxgates are shown surrounding but outside the SCV.

1, 4, & 11 each have a diameter of 35mm and a sensing area of 9.6×10^{-4} m^2. Figure 2.16 shows a photo of the pick-up loops installed in the experiment for the 2010 data run.

Magnetometer	Position			Aligned
	Radial / cm	Angular / degrees	Axial / cm	
SQUID 4	11	-5	-18	x
SQUID 1	11	-5	-18	y
SQUID 11	11	-5	-18	z
SQUID 5	11	75	-11	z
Fluxgate 1	20	45	-25	z
Fluxgate 2	10	45	75	z
Fluxgate 3	10	315	75	z
Fluxgate X	200	0	0	x
Fluxgate Y	200	0	0	y
Fluxgate Z	200	0	0	z

Table 2.2: Table showing the approximate positions and orientations of the magnetometers installed in cryoEDM during the 2010 data run. The angular position is measured from vertical up and axial position measured from the wall of the first neutron cell.

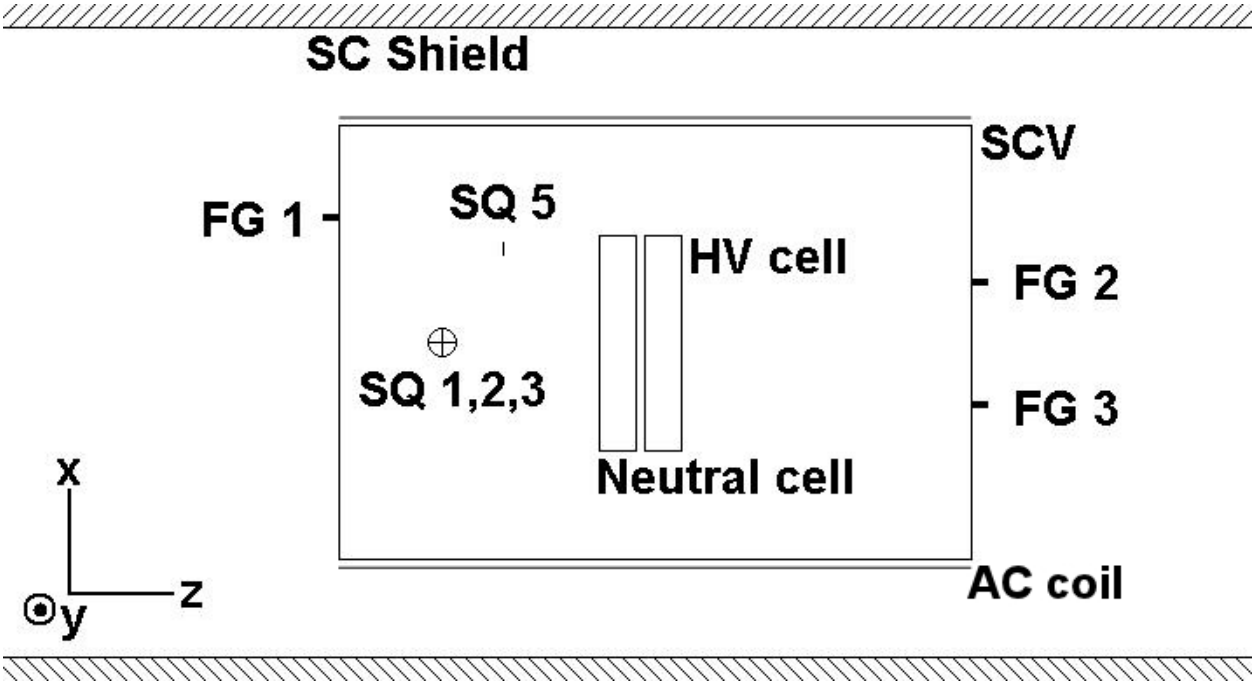

Figure 2.15: Shown is a scale schematic of the size and positions of the magnetometers in and around the SCV viewed from above.

2.7.3.1 Calibration coils

In order to calibrate the SQUIDs, dedicated calibration coils are installed on the same supports as the pick-up loops. Three such coils can be seen in figure 2.16, to the left of centre on the lower pick-up loop and to the extreme left and right of the upper loop. These coils are short solenoids wound on a PTFE former, they have a length of 5 mm, and an inner diameter of 2 mm. They typically have of order 200 turns. They are connected through filtered twisted wire pairs to DACs controlled by the magnetometry software. Table 2.3 gives the position and orientation of the calibration coils installed in the 2010 run.

Coil	Position			Aligned
	Radial / cm	Angular / degrees	Axial / cm	
cc 5	10	250	-8	z
cc 6	10	185	-8	z
cc 7	10	80	-14	z

Table 2.3: Table showing the approximate positions and orientations of the calibration coils installed in cryoEDM during the 2010 data run. The angular position is measured from vertical up and axial position measured from the wall of the first neutron cell.

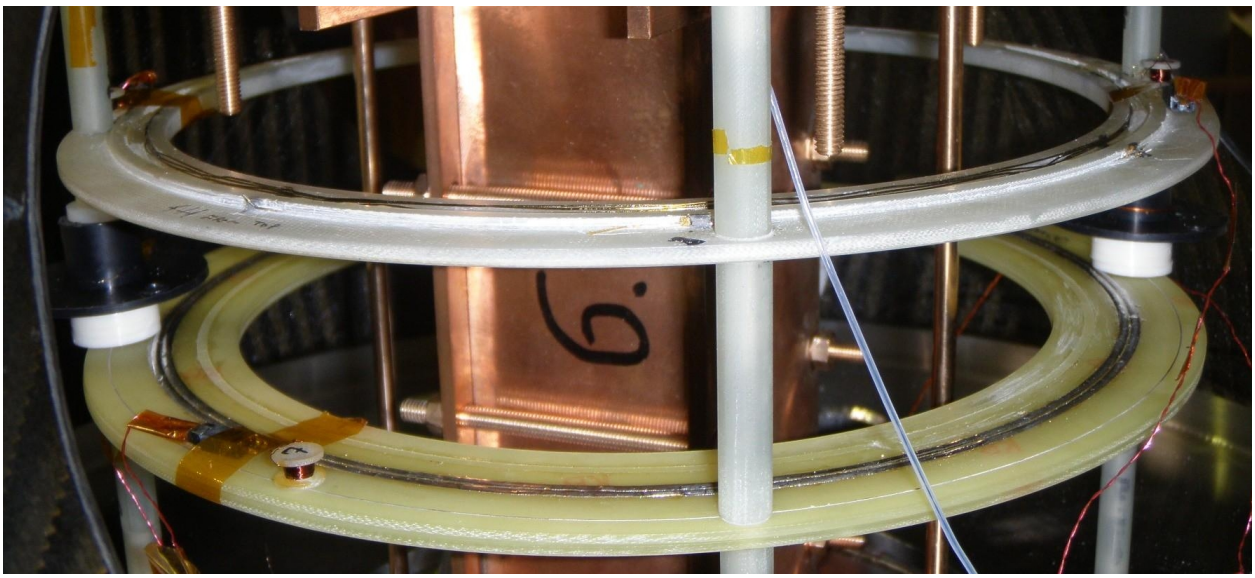

Figure 2.16: Shown are two of the large SQUID pick-up loops, and two small off-axis pick-up loops installed in cryoEDM surrounding the neutron guide, the image is rotated 90° such that the experimental z-axis is from the bottom to the top of the image. The upper pick-up loop is made from niobium-titanium wire glued to the G10 former with stycast. The lower pick-up loop is formed from a lead-tin coated copper track on a PCB. The off-axis loops are niobium titanium wire wound around the black formers on the supports between the large loops. To the left of centre on the lower pick-up loop is a SQUID calibration coil mounted so the solenoid's axis is in the z direction, two more are mounted on the extremes of the upper loop. SQUID 5 is the off-axis loop on the right, the triple axis arrangement is behind the neutron guide.

Figure 2.17 shows a typical installation of the SQUIDs' pick-up loops rotated 90°, which are mounted on the green G10 circular supports near the bottom of the picture. The neutron storage cells are at the top of the picture, isolated from the pick-up loops.

It is using this arrangement of magnetometers that corrections for magnetic field changes and monitoring of the experimental operating points must be performed.

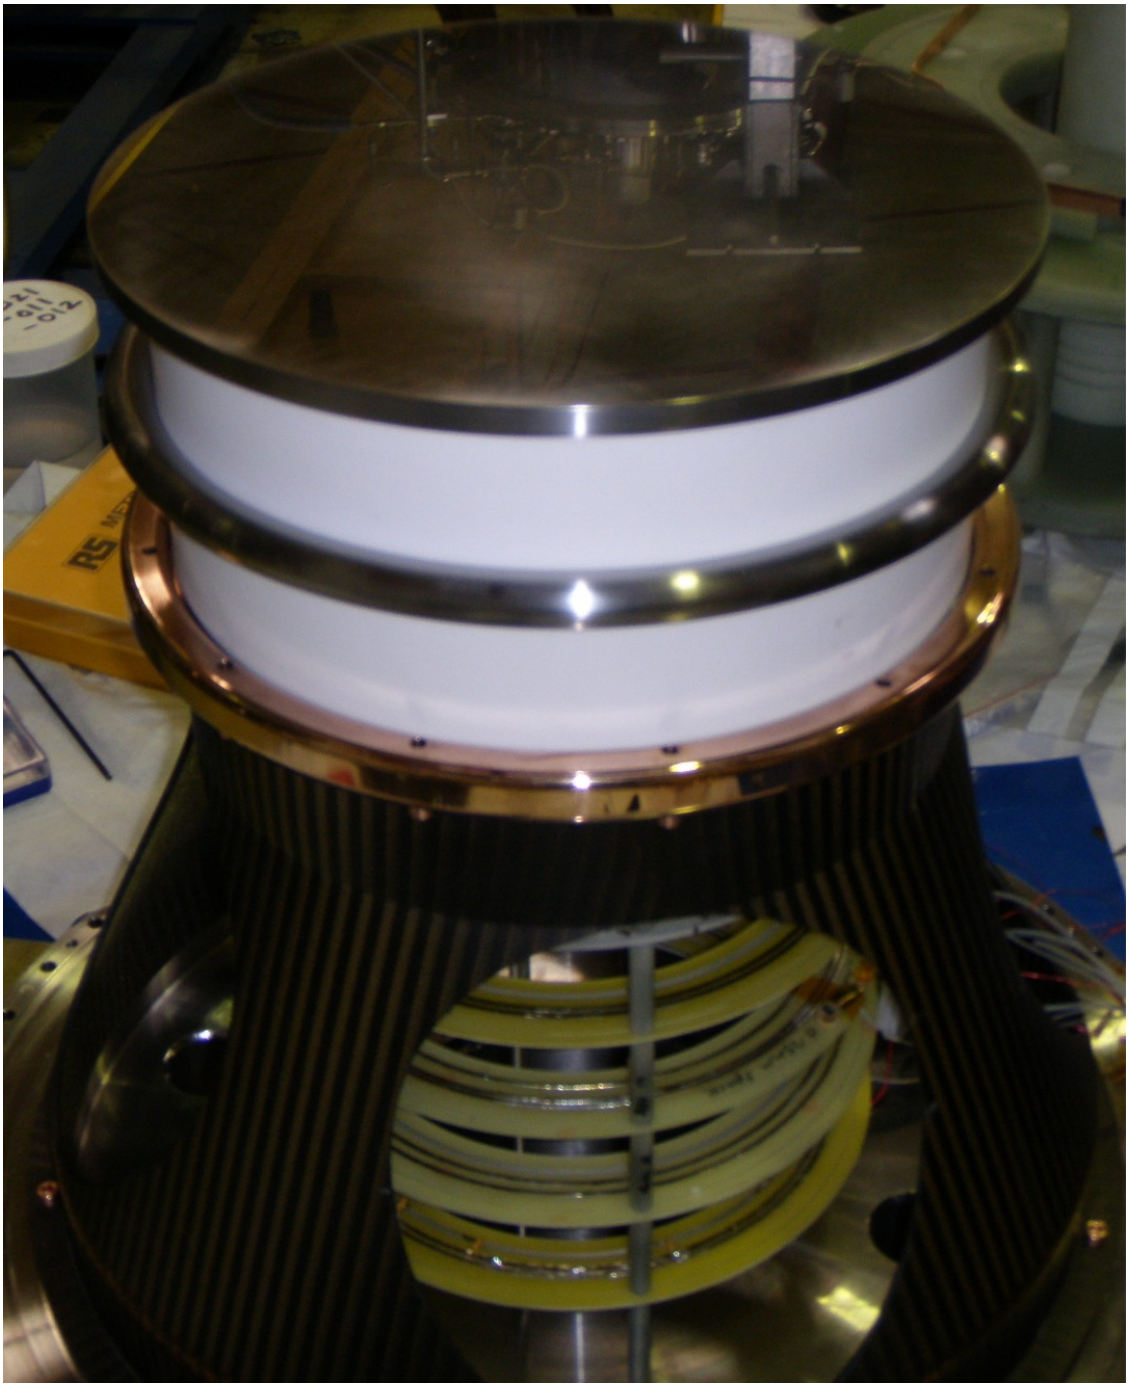

Figure 2.17: Shown is a typical installation of the SQUID pick-up loops rotated 90° so that the experimental z-axis is vertical. The pick-up loops are mounted on the G10 supports (green) visible at the bottom of the picture. The neutron cells and electrodes are near the top the the picture (white and brown respectively).

Chapter 3

Developing a magnetometry system

The goal of this chapter is to produce a method for extracting magnetometry data of use to the experiment from measurements taken by the available magnetometers. These data will be used to compensate for false EDM signals, and to operate as part of a field-stability system keeping the experimental operating point as described in section 2.5.4.

3.1 Mathematical framework

The first task in accomplishing this is to see how individual magnetometer measurements can be used to produce the desired magnetometry output for the rest of the experiment. A short mathematical framework for incorporating these measurements into a final value will be discussed.

3.1.1 Desired quantity

As described in the previous chapter, the measurement procedure of cryoEDM requires determination of the Larmor precession frequency of the neutron. This frequency depends

on the magnitude of the field the neutron experiences. Namely,

$$\omega_B = \frac{2\mu|\boldsymbol{B}|}{\hbar}, \tag{3.1}$$

where ω_B is the precession frequency in a magnetic field of $\boldsymbol{B}$, and μ is the magnetic moment of the neutron.

However, the measurement is made not on a single neutron but over a batch of neutrons and the desired quantity is the average precession frequency of the batch. As the neutrons have such low energies, the distribution of neutrons is slightly asymmetric due to the neutrons' interaction with gravity. However, since each batch of neutrons will have the same centre of gravity, the relation between fields experienced by a uniform distribution and the true distribution will be constant and will cancel during the analysis [99]. The quantity required from the magnetometry is then the value of the precession frequency averaged over both the neutron volume and the storage time, giving;

$$\langle\omega_B\rangle = \left\langle \frac{2\mu|\boldsymbol{B}|}{\hbar} \right\rangle \tag{3.2}$$

$$\langle\omega_B\rangle = \frac{2\mu\,\langle|\boldsymbol{B}|\rangle}{\hbar} \tag{3.3}$$

The requirement of the magnetometry is then to find the volume average of the modulus of the magnetic field within the neutron volume.

3.1.2 Field description

The magnetic field used for the neutron precession is applied axially using a superconducting solenoid, which is coaxial with the neutron volume and approximately ten times longer than the volume. The field produced by the solenoid when in persistent mode should be highly stable, but it is expected that this field will have very small variations in both time and space, either from current instabilities in the solenoid or from external sources penetrating

the magnetic shielding.

The field experienced by the neutrons can be written as the sum of the constant and desired axial field, plus a perturbation field of arbitrary position and time dependency as follows:

$$\begin{aligned} \boldsymbol{B} &= \boldsymbol{B}_{applied} + \boldsymbol{B}_{perturb} && (3.4)\\ &= \begin{pmatrix} 0 \\ 0 \\ B_0 \end{pmatrix} + \begin{pmatrix} B_x(\boldsymbol{r},t) \\ B_y(\boldsymbol{r},t) \\ B_z(\boldsymbol{r},t) \end{pmatrix} && (3.5)\\ &= \begin{pmatrix} B_x(\boldsymbol{r},t) \\ B_y(\boldsymbol{r},t) \\ B_0 + B_z(\boldsymbol{r},t) \end{pmatrix} && (3.6) \end{aligned}$$

where B_0 is the field applied by the solenoid at $r = 0$ and $t = 0$, and $B_{x,y,z}$ are time and positional dependent perturbations to the field in the direction x, y, z respectively. The modulus of this total field can then be taken,

$$\begin{aligned} |\boldsymbol{B}| &= \left({B_x}^2 + {B_y}^2 + (B_0 + B_z)^2\right)^{\frac{1}{2}} && (3.7)\\ &= |B_0| \left(1 + \frac{2B_z}{B_0} + \frac{{B_x}^2 + {B_y}^2 + {B_z}^2}{{B_0}^2}\right)^{\frac{1}{2}} && (3.8) \end{aligned}$$

It is assumed that the perturbations B_x, B_y, and B_z are of a similar size, and $|B_x(\boldsymbol{r},t)|$, $|B_y(\boldsymbol{r},t)|$, and $|B_z(\boldsymbol{r},t)| \ll |B_0|$. The design specifications for the experiment expect $\frac{B_z}{B_0} < 10^{-6}$ to achieve the desired limit, so this should be a valid assumption. Using this assumption it is possible to expand equation 3.8, giving

$$
\begin{aligned}
|\boldsymbol{B}| &\approx |B_0|\left(1+\frac{1}{2}\left[\frac{2B_z}{B_0}+\frac{{B_x}^2+{B_y}^2+{B_z}^2}{{B_0}^2}\right]-\frac{1}{8}\left[\frac{2B_z}{B_0}+\ldots\right]^2\right) && (3.9)\\
&= |B_0|+\frac{|B_0|}{B_0}\left(B_z+\frac{{B_x}^2+{B_y}^2}{2B_0}+\mathcal{O}\left(\frac{{B_z}^3}{{B_0}^2}\right)\right) && (3.10)\\
&= |B_0|+\mathrm{sgn}(B_0)\left(B_z+\frac{{B_x}^2+{B_y}^2}{2B_0}+\mathcal{O}\left(\frac{{B_z}^3}{{B_0}^2}\right)\right) && (3.11)
\end{aligned}
$$

Given this large difference between terms, only the first two terms of the expansion need be considered for a good level of field resolution. Therefore the modulus of the field can be characterised through measurements of only the field's axial component to a first order approximation.

From now on it will be convenient to work in cylindrical polar co-ordinates due to the cylindrical geometry of the region of interest, but the above argument is unaffected. Integrating the modulus of the field over the neutron volume yields

$$
\begin{aligned}
\langle|\boldsymbol{B}|\rangle &= \frac{\int_{z_1}^{z_2}\int_0^{2\pi}\int_0^R |\boldsymbol{B}| r dr d\theta dz}{\int_{z_1}^{z_2}\int_0^{2\pi}\int_0^R r dr d\theta dz} && (3.12)\\
&= \frac{1}{\pi R^2(z_2-z_1)}\int\left(|B_0|+\mathrm{sgn}(B_0)B_z+\mathrm{sgn}(B_0)\frac{B_r^2}{2B_0}+\ldots\right) r dr d\theta dz && (3.13)\\
&= |B_0|+\frac{\mathrm{sgn}(B_0)}{\pi R^2(z_2-z_1)}\int B_z r dr d\theta dz+\ldots && (3.14)
\end{aligned}
$$

The total field expansion can be written as a series of increasingly higher order correction

terms to the field applied by the solenoid, these terms are

$$\langle|\boldsymbol{B}|\rangle_0 = |B_0| \tag{3.15}$$

$$\langle|\boldsymbol{B}|\rangle_1 = \mathrm{sgn}(B_0)\langle B_z\rangle \tag{3.16}$$

$$\langle|\boldsymbol{B}|\rangle_2 = \mathrm{sgn}(B_0)\frac{\langle B_r^2\rangle}{2B_0} \tag{3.17}$$

$$\langle|\boldsymbol{B}|\rangle_3 = -\mathrm{sgn}(B_0)\frac{\langle B_z B_r^2\rangle}{2B_0^2} \tag{3.18}$$

for the zeroth, first, second, and third orders respectively. Given the expected size of perturbations, $\frac{B_z}{B_0} < 10^{-6}$, each correction term will be approximately 6 orders of magnitude smaller than the preceding term.

Now only a parameterisation of the components of the field is required to complete the model.

3.1.3 Form of the field

The majority of the field will be a relatively large compared to the fluctuations that need correction, and will be produced by a current in the solenoid. Perturbations to the field will be from sources outside of the measurement volume, or they could be eliminated, and so there are no other local currents. Further, as the experiment is conducted within a superconducting shield, which is much longer than the neutron volume, it is a reasonable assumption that any axial perturbations have a linear axial dependence, therefore we assume that

$$B_z(\boldsymbol{r},t) = \left\langle \frac{\partial B_z}{\partial z} \right\rangle (t)z + \delta B_0(t) \tag{3.19}$$

where $\left\langle \frac{\partial B_z}{\partial z} \right\rangle$ is the magnetic field gradient through the volume and $\delta B_0(t)$ accounts for any time dependent changes in the solenoid field. Integrating this over the neutron volume to find the average yields;

$$\langle B_z \rangle = \frac{1}{\pi R^2 (z_2 - z_1)} \int_{z_1}^{z_2} \int_0^{2\pi} \int_0^R \left(\left\langle \frac{\partial B_z}{\partial z} \right\rangle z + \delta B_0(t) \right) r dr d\theta dz \tag{3.20}$$

$$= \frac{{z_2}^2 - {z_1}^2}{2(z_2 - z_1)} \left\langle \frac{\partial B_z}{\partial z} \right\rangle + \delta B_0(t) \tag{3.21}$$

$$= \frac{z_2 + z_1}{2} \left\langle \frac{\partial B_z}{\partial z} \right\rangle + \delta B_0(t) \tag{3.22}$$

Substituting for B_z from equation 3.22 into equation 3.14 gives,

$$\langle |\boldsymbol{B}| \rangle = |B_0| + \mathrm{sgn}(B_0) \langle B_z \rangle \tag{3.23}$$

$$= |B_0| + \mathrm{sgn}(B_0) \frac{z_2 + z_1}{2} \left\langle \frac{\partial B_z}{\partial z} \right\rangle + \mathrm{sgn}(B_0) \delta B_0(t) \tag{3.24}$$

We must now determine how the measurements from magnetometers can be used to determine the parameters of the field.

3.1.4 Determining parameters

It must be remembered that neither the magnetic field nor its gradient can be measured by the SQUIDs, only changes to their value may be determined. At the levels of field stability we are hoping to achieve, the fluxgates will not have the resolution to measure a these adequately. This leaves only the neutrons themselves as a suitable magnetometer to establish a baseline (discussed in section 3.4) and to use the SQUIDs for shifts in this value.

Once a reasonable idea of the field has been established the above parameterisation can be fit to the measurements through any suitable algorithm to find these changes.

3.1.4.1 Restructuring

Before proceeding it will be useful to rewrite equation 3.24 in a more convenient form, and make use of some known geometrical properties. For simplicity we set the origin of the co-ordinate system at the wall of the neutron volume and coaxial with it, we shall also note that the neutron cells are identical and call their length is L. This makes equation 3.24,

$$\langle |\boldsymbol{B}| \rangle = |B_0| + \mathrm{sgn}(B_0)\frac{L}{2}\left\langle \frac{\partial B_z}{\partial z} \right\rangle + \mathrm{sgn}(B_0)\delta B_0(t). \tag{3.25}$$

cryoEDM has multiple cells arranged in a row. This introduces a factor of $C_c = 2c - 1$ for cell number $c = 1$, 2, 3, etc. So

$$\langle |\boldsymbol{B}| \rangle_c = |B_0| + \mathrm{sgn}(B_0)\frac{C_c L}{2}\left\langle \frac{\partial B_z}{\partial z} \right\rangle + \mathrm{sgn}(B_0)\delta B_0(t) \tag{3.26}$$

but as this is just a cell specific numerical factor it shall be ignored here.

We shall also multiply the right hand side of equation 3.25 by $\frac{\mathrm{sgn}(B_0)}{\mathrm{sgn}(B_0)}$ to give

$$\begin{aligned}
\langle |\boldsymbol{B}| \rangle &= \frac{\mathrm{sgn}(B_0)|B_0| + \mathrm{sgn}^2(B_0)\frac{L}{2}\left\langle \frac{\partial B_z}{\partial z} \right\rangle + \mathrm{sgn}^2(B_0)\delta B_0(t)}{\mathrm{sgn}(B_0)} && (3.27)\\
&= \frac{B_0 + \frac{L}{2}\left\langle \frac{\partial B_z}{\partial z} \right\rangle + \delta B_0(t)}{\mathrm{sgn}(B_0)} && (3.28)\\
&\approx \left| B_0 + \frac{L}{2}\left\langle \frac{\partial B_z}{\partial z} \right\rangle + \delta B_0(t) \right|, && (3.29)
\end{aligned}$$

where in the last line we remember the fact that the perturbation to the field over the neutron volume is very small compared to the applied field, so the sign of the applied field is the same as the sign of the total field experienced by the neutrons.

We now have an equation giving the field experienced by the neutrons in terms of only two unknown parameters, the applied axial field $B_0 + \delta B_0(t)$, and the planar gradient $\left\langle \frac{\partial B_z}{\partial z} \right\rangle$. Both of these values must be determined from the measurements made using the SQUIDs, which will now be inserted into this framework.

3.1.4.2 Fitting

For the following discussion we are going to restrict ourselves to pick-up loops that sample the axial component of the field. Then the measurements from the SQUIDs give the value of the total field at a particular point as

$$B_S = B_0 + \delta B_0(t) + \left\langle \frac{\partial B_z}{\partial z} \right\rangle z. \tag{3.30}$$

Given N measurements of the field from pick-up loops at various z positions we can use a standard least squares fit to produce the following equations,

$$B_0 + \delta B_0(t) = \frac{\sum_i^N B_{S_i} \sum_i^N z_i^2 - \sum_i^N z_i \sum_i^N z_i B_{S_i}}{N \sum_i^N z_i^2 - \left(\sum_i^N z_i \right)^2} \tag{3.31}$$

$$\left\langle \frac{\partial B_z}{\partial z} \right\rangle = \frac{N \sum_i^N z_i B_{S_i} - \sum_i^N z_i \sum_i^N B_{S_i}}{N \sum_i^N z_i^2 - \left(\sum_i^N z_i \right)^2}, \tag{3.32}$$

where z_i is a position of a contributing pick-up loop, and B_{S_i} is the measurement of the field at the point z_i. It is also possible to calculate the associated uncertainty on these values, the equations shall not be presented here for brevity.

It is now possible to combine equations 3.29, 3.31 & 3.32 and insert the magnetometer measurements into the framework giving,

$$\langle |\boldsymbol{B}| \rangle = \left| B_0 + \delta B_0(t) + \frac{L}{2} \left\langle \frac{\partial B_z}{\partial z} \right\rangle \right| \tag{3.33}$$

$$\langle|\boldsymbol{B}|\rangle = \left|\sum_i^N B_{S_i}\left(\frac{\sum_i^N z_i^2 - z_i\sum_i^N z_i}{N\sum_i^N z_i^2 - \left(\sum_i^N z_i\right)^2} + \frac{L}{2}\frac{Nz_i - \sum_i^N z_i}{N\sum_i^N z_i^2 - \left(\sum_i^N z_i\right)^2}\right)\right| \tag{3.34}$$

$$= \left|\sum_i^N B_{S_i}[\alpha + \beta z_i]\right| \tag{3.35}$$

where α and β are constants, which depend only on the geometry;

$$\alpha = \frac{\sum_i^N z_i^2 - \frac{L}{2}\sum_i^N z_i}{N\sum_i^N z_i^2 - \left(\sum_i^N z_i\right)^2} \tag{3.36}$$

$$\beta = \frac{\frac{NL}{2} - \sum_i^N z_i}{N\sum_i^N z_i^2 - \left(\sum_i^N z_i\right)^2} \tag{3.37}$$

3.1.5 Time average

In addition to the volume average, the time average of this magnetic field must be taken during a measurement cycle. Both functions of the magnetometry system, false EDM correction and field stabilisation, only require knowledge of changes in the time average of the field between certain points. This can be used to remove further unknown quantities. The desired magnetometry output is

$$dB_{2,1} = \langle\langle|\boldsymbol{B}|\rangle\rangle_{t_2} - \langle\langle|\boldsymbol{B}|\rangle\rangle_{t_1} \tag{3.38}$$

$$= \frac{1}{\delta t_2}\int_{t2-\delta t_2/2}^{t2+\delta t_2/2}\langle|\boldsymbol{B}|\rangle\,dt - \frac{1}{\delta t_1}\int_{t1-\delta t_1/2}^{t1+\delta t_1/2}\langle|\boldsymbol{B}|\rangle\,dt \tag{3.39}$$

We can insert the values obtained in equation 3.35

$$\begin{aligned} dB_{2,1} &= \frac{1}{\delta t_2}\int_{t2-\delta t_2/2}^{t2+\delta t_2/2}\left|\sum_i^N B_{S_i}[\alpha+\beta z_i]\right|dt \\ &\quad -\frac{1}{\delta t_1}\int_{t1-\delta t_1/2}^{t1+\delta t_1/2}\left|\sum_i^N B_{S_i}[\alpha+\beta z_i]\right|dt \qquad (3.40) \\ &= \mathrm{sgn}(B_0)\sum_i^N[\alpha+\beta z_i]\left(\frac{1}{\delta t_2}\int_{t2-\delta t_2/2}^{t2+\delta t_2/2}B_{S_i}dt-\frac{1}{\delta t_1}\int_{t1-\delta t_1/2}^{t1+\delta t_1/2}B_{S_i}dt\right), \qquad (3.41) \end{aligned}$$

where it has been assumed that the sign of B_0 is not changed between measurements, which will be the case in normal operation. This equation gives the desired quantity in terms of measurements of the magnetic field experienced by SQUIDs, B_{S_i}, and the sign of B_0, but not its magnitude, obtainable from the fluxgates.

Of course the time integral will be performed on a digitised signal and so equation 3.41 reduces to the following discrete sum

$$\begin{aligned} dB_{2,1} &= \mathrm{sgn}(B_0)\sum_i^N[\alpha+\beta z_i]\left[\frac{1}{M_1\delta t}\sum_{n=\frac{M_1}{2}}^{\frac{M_1}{2}-1}B_{S_i}(t_2+n\delta t)\delta t\right. \\ &\quad \left.-\frac{1}{M_2\delta t}\sum_{m=\frac{M_2}{2}}^{\frac{M_2}{2}-1}B_{S_i}(t_1+m\delta t)\delta t\right] \qquad (3.42) \\ &= \mathrm{sgn}(B_0)\sum_i^N[\alpha+\beta z_i]\left[\frac{1}{M_1}\sum_{n=\frac{M_1}{2}}^{\frac{M_1}{2}-1}B_{S_i}(t_2+n\delta t)-\frac{1}{M_2}\sum_{m=\frac{M_2}{2}}^{\frac{M_2}{2}-1}B_{S_i}(t_1+m\delta t)\right] \qquad (3.43) \end{aligned}$$

This gives the full algorithm for finding the changes in the time- and volume-averaged magnitude of the magnetic field in the regions of interest, which is the desired magnetometry output of the SQUID system. The algorithm still requires good knowledge of the calibration of the magnetometers, a problem which will be addressed in section 3.2.

3.2 Calibration

The method for extracting the desired magnetometry signal described in section 3.1 is dependent on knowing the magnetic field measured through each pick-up loop. This requires a good knowledge of the calibration of the magnetometers and pick-up loop system to convert the raw voltage into the measured flux. In this section, a method for calibrating and cross-calibrating the SQUID's raw output will be discussed: the application of a known field using specially installed calibration coils; and the measurement of an arbitrary field with another sufficiently precise magnetometer. A discussion of the limitations in the ability to apply precise fields will be presented, and a calibration scheme using the only sufficiently precise magnetometer available, the neutrons themselves, provided.

Measurements of the calibration of SQUIDs installed in cryoEDM is presented using the calibration coils and compared to calibrations obtained from independent methods. Calibrations intended to be produced by using the neutrons were unobtainable due to external problems in making neutron measurements, instead results based on a Monte Carlo simulation and from a testing rig are presented for this method.

3.3 Calibration via coils

To calibrate the magnetometers a known magnetic field must be applied and their response measured. A set of short solenoids, calibration coils, are installed close to the SQUID pick-up loops. These can be used to apply magnetic fields to the pick-up loops, and so to calibrate them if a the field produced at the pick-up loop is known.

The magnetic field produced by a calibration coil could be measured using the fluxgates, but as a pick-up loop could be any surface, and allowing for future designs changes, a full map of the field at all possible positions and orientations would be required. As some pick-up loop designs call for sensing regions of millimetre dimensions, the fluxgate spatial resolution of approximately 1 cm^3 means measuring such a map would not be practical. To find the

magnetic field produced by the calibration coils a calculation based a coil's geometry and using standard electromagnetism relations is performed.

A method for performing this kind of calculation was available, but with severe limitations in the calculations that could be performed. It was also a labour intensive operation with certain values left to the interpretation of the user. Signals would be applied to each calibration coil manually using a function generator, the response would have to be extracted from the recorded data by hand. Basic calculations were performed on the data using the approximation that the source was dipole-like, and the relative orientations of the sources and the pick-up loops were limited to be parallel or antiparallel. Additionally only specific shapes of pick-up loop were handled, either planar circles or rings. This process would often take several days to complete and would not provide a robust measure of the uncertainty in the calibration.

This section details the work done to create a new framework that would make automating the calibration possible, applying signals and measuring responses without need for intervention. The aim was to not limit the system with geometric assumptions that were likely to change in the future, and to provide a reasonable estimate of uncertainty in the measurement. The aim was to make this form of calibration take minutes rather than days.

3.3.1 Calculating applied current-to-flux conversions

The first step to performing the calibration is to apply a known flux through the pick-up loops. This requires knowledge of the field produced by various calibration coils and integrating these fields over the sensing area of each of the pick-up loops.

3.3.1.1 Calculating the field

The calibration coils used in cryoEDM are short solenoids, with a small radius and numerous layers. A number of ways of modelling these coils were created, ranging from a simple analytic solution using a dipole approximation to calculating the field from line segments, forming a

set of inter-wound helices, directly by applying a discretised version of Biot-Savart's law.

The line segment model gives superior results in the cases where the coils are close to the pick-up loops, and so the dipole approximation is not valid. This can often be the case as the calibration coils are mounted on the same supports as the pick-up loops, and the nearest pick-up loop to the calibration coil will be within a short distance of the coil. One drawback of the line segment model is that it requires five orders of magnitude more computing time to calculate the fields, 0.16 μs vs 14 ms for the dipole and line segment model respectively to calculate the field at one location on the available hardware, however this can be performed offline and so does not effect the calibration. Another issue is the requirement for knowing the shape of the wire in the coil. As the coils are manufactured by the Oxford group using a hand operated winding machine that places the wire on specially designed formers, and great care is taken to maintain the wire's alignment, their structure should be highly predictable.

The models were tested to ensure that they gave the analytically calculable results for simple geometric cases, e.g. the on axis field of long solenoids, and also that the field from the solenoid model limits to the dipole model's field at large distances.

3.3.1.2 Numerical integration

While all currently installed pick-up loops are planar and somewhat circular, this will not necessarily be the case for future loop designs. To ensure the calibration system can cope with future improvements a pick-up loop is defined as an arbitrary surface in 3-dimensions. This is done by forming the surface using arbitrary triangles.

Once a rough outline of the surface has been provided by a user in terms of large triangles, these triangles can be bisected along their longest side, producing triangles with half the surface area. This process is repeated until the area of all triangles is below a certain threshold value. In this way it is possible to create a map of the surface with arbitrarily small integral regions from a rough outline provided. An example is shown in figure 3.1 where a pick-up loop, defined roughly in terms of 20 small triangles, is split into integral

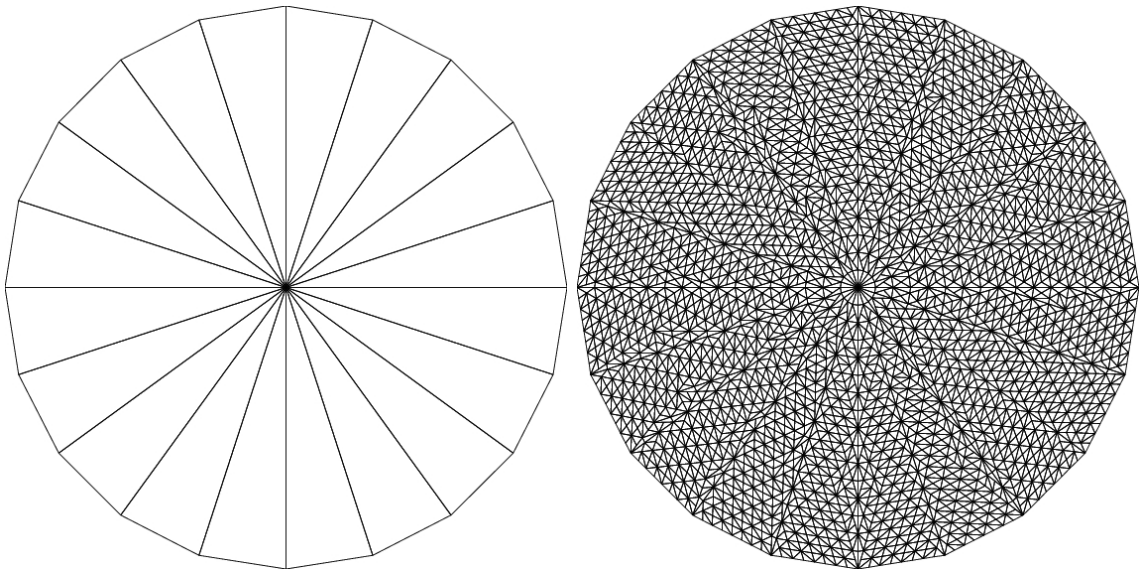

Figure 3.1: *(Left)* An example of a rough outline of a circular pick-up loop defined in terms of 20 triangles. *(Right)* An example of the calculated integral regions each with an area over 3000 times smaller than the original description, formed by recursively bisecting each triangle along its longest side to reach a threshold area.

regions with an area over 3000 times smaller than the originally defined pick-up loop.

The field from each calibration coil is then calculated at the centroid of each triangle. The flux through the triangle is approximated as

$$\Phi = (\boldsymbol{B_c} \cdot \hat{\boldsymbol{n}})\, A \tag{3.44}$$

where $\boldsymbol{B_c}$ is the field at the triangle's centroid, $\hat{\boldsymbol{n}}$ is the normal to the plane of the triangle, and A is the area of the triangle. The total flux is then calculated as the sum of the flux through each triangle.

As Maxwell's equations are linear it is sufficient to calculate the value of the field at a single current and scale the result as necessary. The values are therefore calculated as flux-per-unit-current, $\frac{\partial \Phi}{\partial I}$, in weber per amp.

3.3.1.3 Uncertainty in flux

The major source of uncertainty in the applied flux-per-unit-current value is from the relative positions of the calibration coils and pick-up loops. Due to the complicated positional dependencies of the method above, determining these uncertainties is not possible analytically. A Monte Carlo simulation was written using the uncertainties in position, orientation, and shape of pick-up loops and calibrations coils as input distributions. By performing and collating the field calculations at each step of the Monte Carlo simulation, the distributions, and hence uncertainty, of the total flux received could then be computed for each calibration-coil/pick-up-loop pair. Figure 3.2 is a histogram showing the distribution obtained from the Monte Carlo for SQUID 5's pick-up loop from calibration coil 7, and table 3.1 shows the mean and RMS of the distributions for all pick-up loop/calibration coil pairings.

	Coil 5		Coil 6		Coil 7	
	Mean	RMS	Mean	RMS	Mean	RMS
SQUID 1	-1.0	0.56	-19	7.9	62	25
SQUID 4	0.26	4.7	-41	14	-46	21
SQUID 5	32	12	2100	15000	-6.9	1.3
SQUID 11	35	7.1	-46	12	-79	24

Table 3.1: The mean and RMS of the distributions of the flux-per-unit-current values in pWb A^{-1} for each pick-up loop/calibration coil pair.

3.3.2 Extracting calibration signal response

A method of reliably extracting the response of the magnetometer to the applied field was required. This needed to be done without the need for an operator to review the data. The calibration is performed in a magnetically noisy environment, and so a degree of tolerance of noise is desirable. The following is a description of the algorithm used to measure the magnetometer response and estimate the uncertainty in its value.

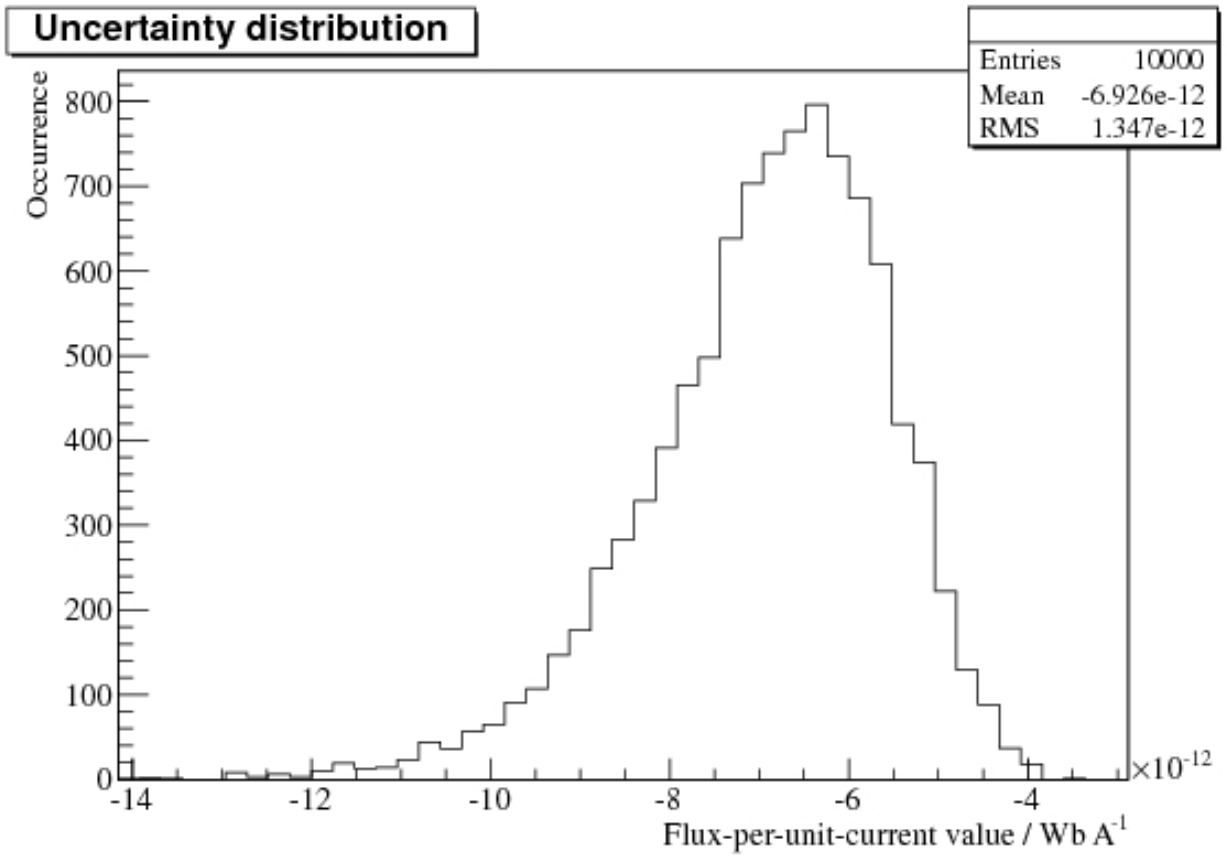

Figure 3.2: Shown is the distribution in flux-per-unit-current values for SQUID 5 and calibration coil 7 obtained from the Monte Carlo. Due to the non-linear dependence of the flux on position this is a slightly skewed gaussian.

3.3.2.1 Measuring SQUID response

The magnetometer output during the calibration signal will be the sum of the signal and background field changes

$$V(t) = S(t) + B(t), \tag{3.45}$$

where $V(t)$ is the raw magnetometer output, $S(t)$ is the response to the applied signal, and $B(t)$ is the background field. If it is assumed that the background drifts linearly then the integral of the applied signal component of the SQUID output can be determined from measurements of the SQUID output through the following equation

$$\int_{t_0}^{t_1} S(t)dt = (t_1 - t_0)\left(\langle V\rangle - \frac{\langle V\rangle_{\text{before}} + \langle V\rangle_{\text{after}}}{2}\right) \tag{3.46}$$

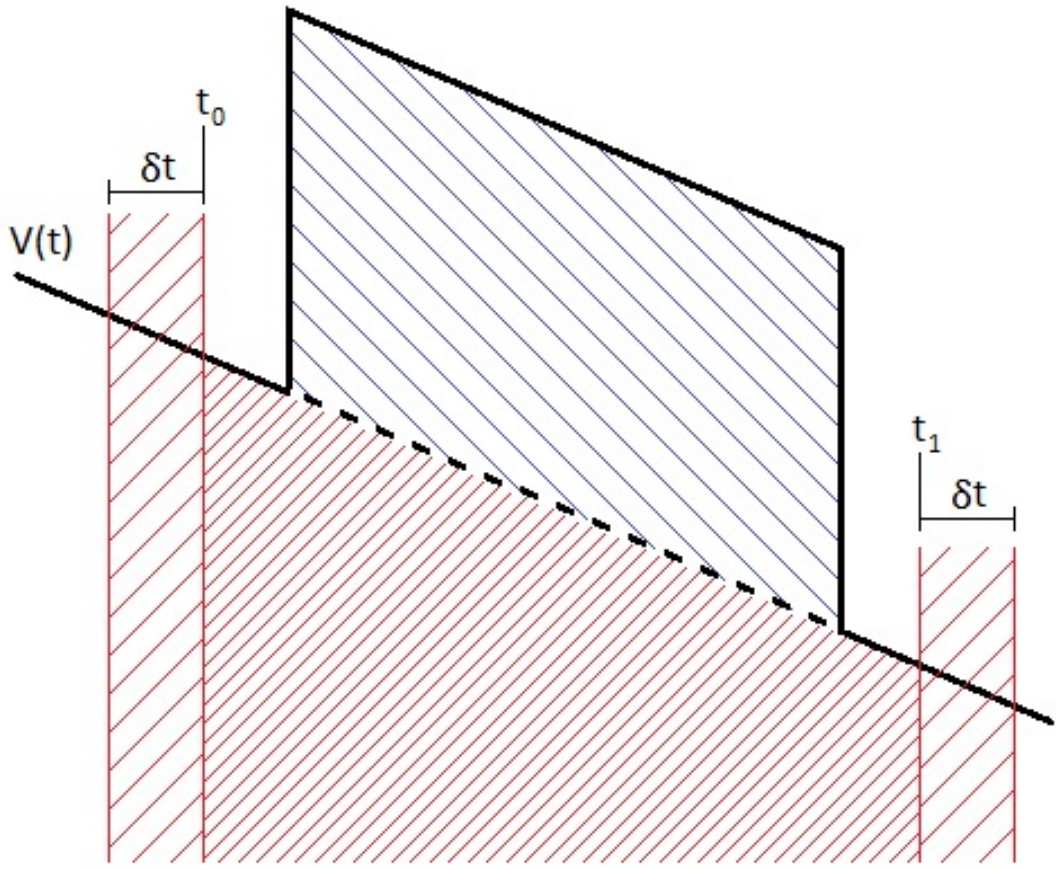

Figure 3.3: Shown is the principle behind algorithmically recovering the response to a calibration pulse. The red forward hatched region has the same time average as the red double hatched region. This can be used to recover the integral of the pulse (blue backward hatched region) by subtracting the integral of the average from the integral of $V(t)$ between t_0 and t_1.

where $\langle V \rangle$, $\langle V \rangle_{\text{before}}$, and $\langle V \rangle_{\text{after}}$ are the time averages of the output between t_0 and t_1, a point shortly before the pulse, and a point shortly after the pulse respectively. This relation is derived in appendix B.1.

This process can be seen illustrated in figure 3.3, which shows a square pulse on a decreasing background, the value of the background can be estimated from the regions either side of the pulse and the pulse can be extracted from this and the integral of the signal.

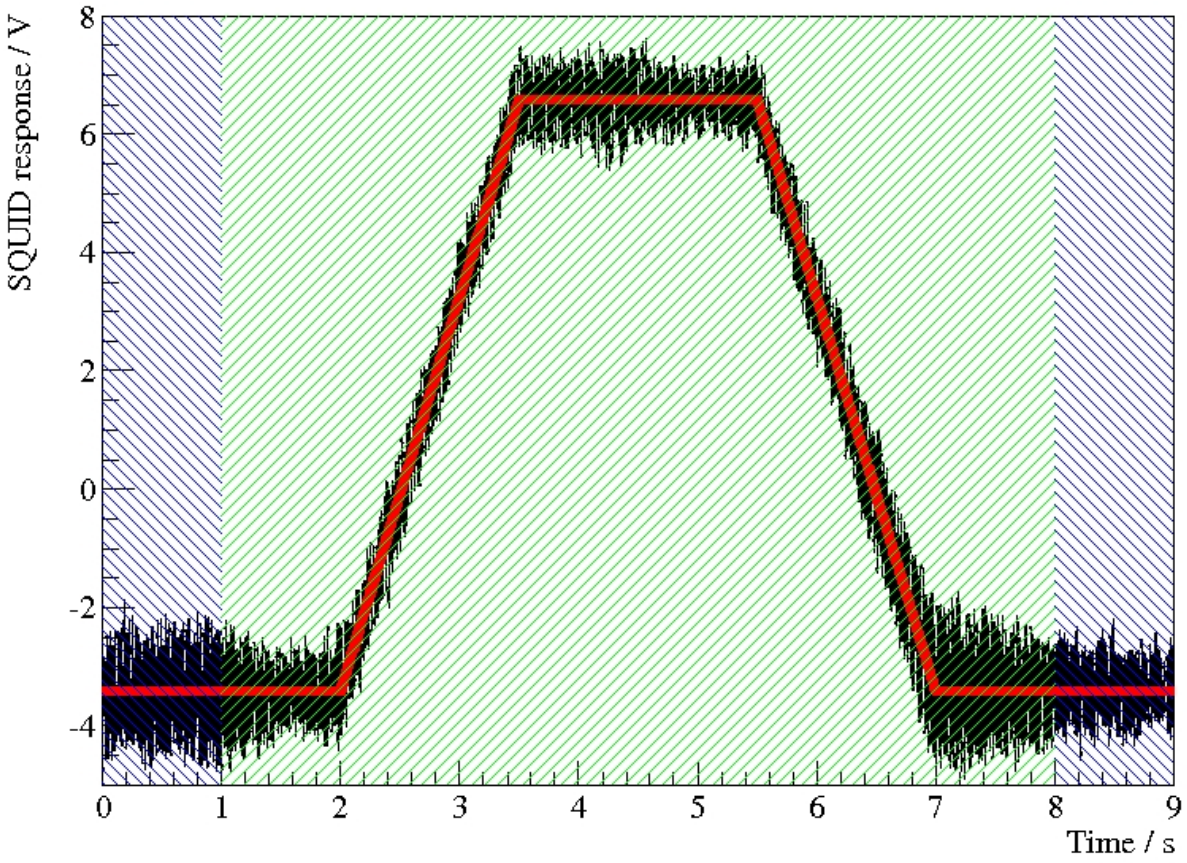

Figure 3.4: Example of a typical calibration pulse. Shown is the SQUID response to the pulse (black); the applied pulse shape (red); the region integrated to determine the pulse height (green, forward hatching); regions to be used to determine the gradient of the background field (blue, backward hatching). The data were recorded using a DAQ system with a bandwidth of 1 kHz. This was produced from a calibration coil with a current of 1.1 mA. A response of $(10.4 \pm 0.4)\,\mathrm{V}$ was extracted from the data. The magnetic flux through the pick-up loop was calculated to be $-2.23\,\mathrm{pWb}$, leading to a magnetic field of $-11.1\,\mathrm{nT}$ and a calibration coefficient of $-1.07\,\mathrm{nTV}^{-1}$.

3.3.2.2 Choice of calibration signal

The obvious choice for a calibration signal is a square pulse, as this gives a very sharp change in signal that should be easy to identify programmatically. However, due to the limited slew rate of the SQUIDs' feedback system (see section 2.7.2.1), rapid changes in field can cause instability in the output signal, it is therefore a requirement to keep field changes below a certain rate.

This effectively rules out a square pulse for all but the smallest of step changes and leads

to a trapezoidal pulse. For simplicity the trapezoidal pulse is chosen to be symmetric in time about the centre of the pulse, so having equal rise and fall times. This is not necessary for the following discussion, but generality does not provide any additional functionality in this case. Figure 3.4 shows the pulse shape used, and the response on the SQUIDs with an obvious correlation. Also shown are the regions used to determine the pulse and the background fields.

We can now find the integral of the expected signal as

$$\begin{aligned}\int_{t_0}^{t_1} S(t)dt &= \int_{t_0}^{t_0+t_{\text{rise}}} V_{\text{pulse}}\frac{t-t_0}{t_{\text{rise}}}dt + \int_{t_0+t_{\text{rise}}}^{t_1-t_{\text{rise}}} V_{\text{pulse}}dt + \int_{t_1-t_{\text{rise}}}^{t_1} V_{\text{pulse}}\frac{t_1-t}{t_{\text{rise}}}dt \quad (3.47)\\ &= \frac{V_{\text{pulse}}t_{\text{rise}}}{2} + V_{\text{pulse}}(t_1-t_0-2t_{\text{rise}}) + \frac{V_{\text{pulse}}t_{\text{rise}}}{2} \quad (3.48)\\ &= V_{\text{pulse}}(t_1-t_0-t_{\text{rise}}), \quad (3.49)\end{aligned}$$

where t_{rise} is the time taken for the pulse to rise from zero to the maximum value V_{pulse} assuming the pulse starts at t_0 and ends at t_1. Of course this would require good knowledge of the position of the pulse to identify t_0 and t_1, which is not always available. It is simple to extend the model to allow for the pulse to be merely between t_0 and t_1, and not delimited by them, by introducing the time the pulse is at its maximum level, t_{length} and noting the signal is zero at other times, yielding

$$\int_{t_0}^{t_1} S(t)dt = V_{\text{pulse}}(t_{\text{length}}+t_{\text{rise}}). \quad (3.50)$$

Combining equations 3.50 and 3.46 it is possible to recover the height of the pulse observed by the magnetometer

$$V_{\text{pulse}} = \frac{t_1-t_0}{t_{\text{length}}+t_{\text{rise}}}\left(\langle V\rangle - \frac{\langle V\rangle_{\text{before}} + \langle V\rangle_{\text{after}}}{2}\right). \quad (3.51)$$

The response of a magnetometer to a pulse can therefore be calculated from integrals of the magnetometer output before, during, and after an expected pulse, with knowledge of the timing structure of the applied pulse.

3.3.2.3 Noise immunity

The integral treatment above can extract the value of the calibration even if the response to the signal is lower than the noise on the magnetometer. As much of the noise is high frequency it will be averaged out in the integral. However, the above response finding algorithm will produce a value even if the signal is deformed due to some magnetometer artefact or, in the case that no signal is present, a very small response will be reported due to inexact noise cancellation.

To counter this, the integral region was split into a number of averages of the signal over short periods, this would not affect the value of the integral, but would provided a level of uncertainty in the integral's value. In this reduced dataset, a judgement of how well the observed pulse shape matches the known applied pulse shape was needed. Due to the requirement to operate in real time a χ^2-test could not be performed as fitting the pulse shape to 12 channels could not be accomplished fast enough on the DAQ hardware, requiring approximately 2 orders of magnitude longer than was available. Instead the known and observed pulse shapes were plotted against each other, in the presence of a pulse the relation between the two variables should be linear and so the correlation between the two signals can be taken, this value is defined as standard as

$$r_{xy} = \frac{\sum(x_i - \bar{x})(y_i - \bar{y})}{\sqrt{\sum(x_i - \bar{x})^2 \sum(y_i - \bar{y})^2}} \tag{3.52}$$

where x_i and y_i are coincident data points from the two channels being compared. The modulus of this was taken to provide a correlation parameter that ran from 0 to 1, with 1 being well matching. Figure 3.5 shows the distribution of pulses found by the algorithm with and without calibration pulses present. The real pulses are found in the gaussian peak above the decaying noise. Pulses with a value of the correlation parameter high enough to be out of the exponentially decaying noise were chosen as good pulses and used to determine the calibration factors.

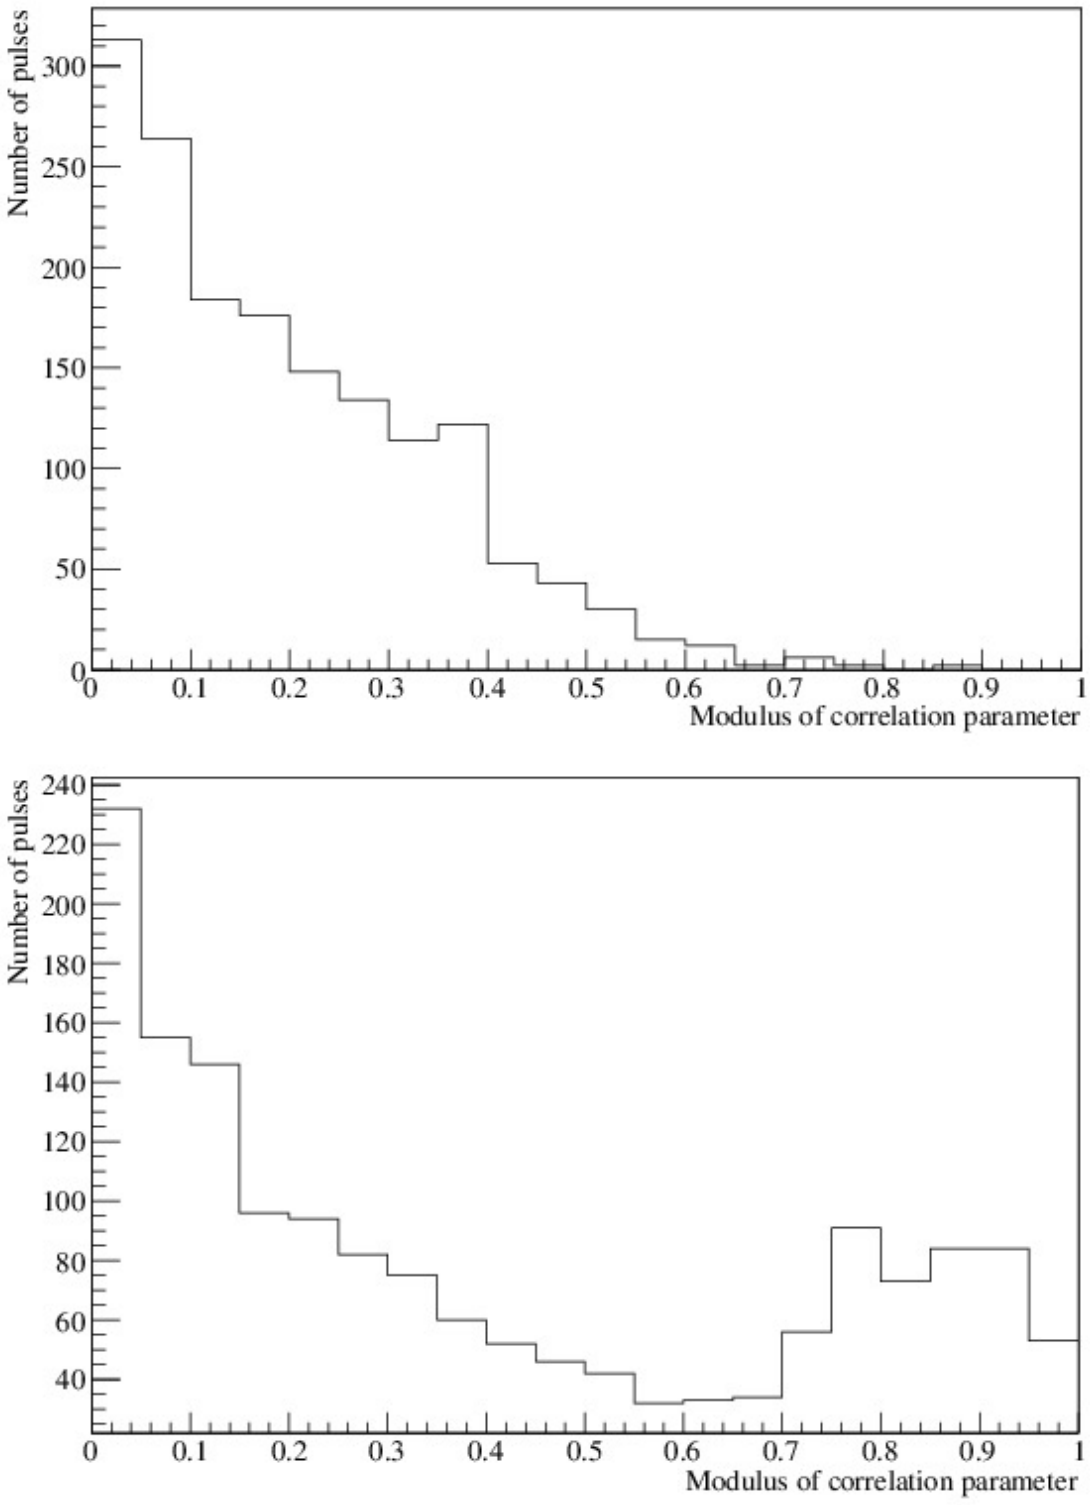

Figure 3.5: Histograms of the numbers of pulses extracted using the above algorithm having certain correlation parameters without *(Top)*, and with *(Bottom)* calibration pulses present in the data. Without pulses a decaying exponential distribution is observed caused by fitting to noise. With pulses a gaussian peak is visible in the high correlation region. This is used to set the threshold value for the detection algorithm.

Shown in figure 3.6 is an example of a calibration pulse with response significantly below the noise level, shown as raw data, as recorded from the magnetometer, and with the signal smoothed over 200 data points. The correlation with the known pulse shape can be compared between the two cases. This method clearly finds pulses with significantly smaller magnitudes than can by found by eye.

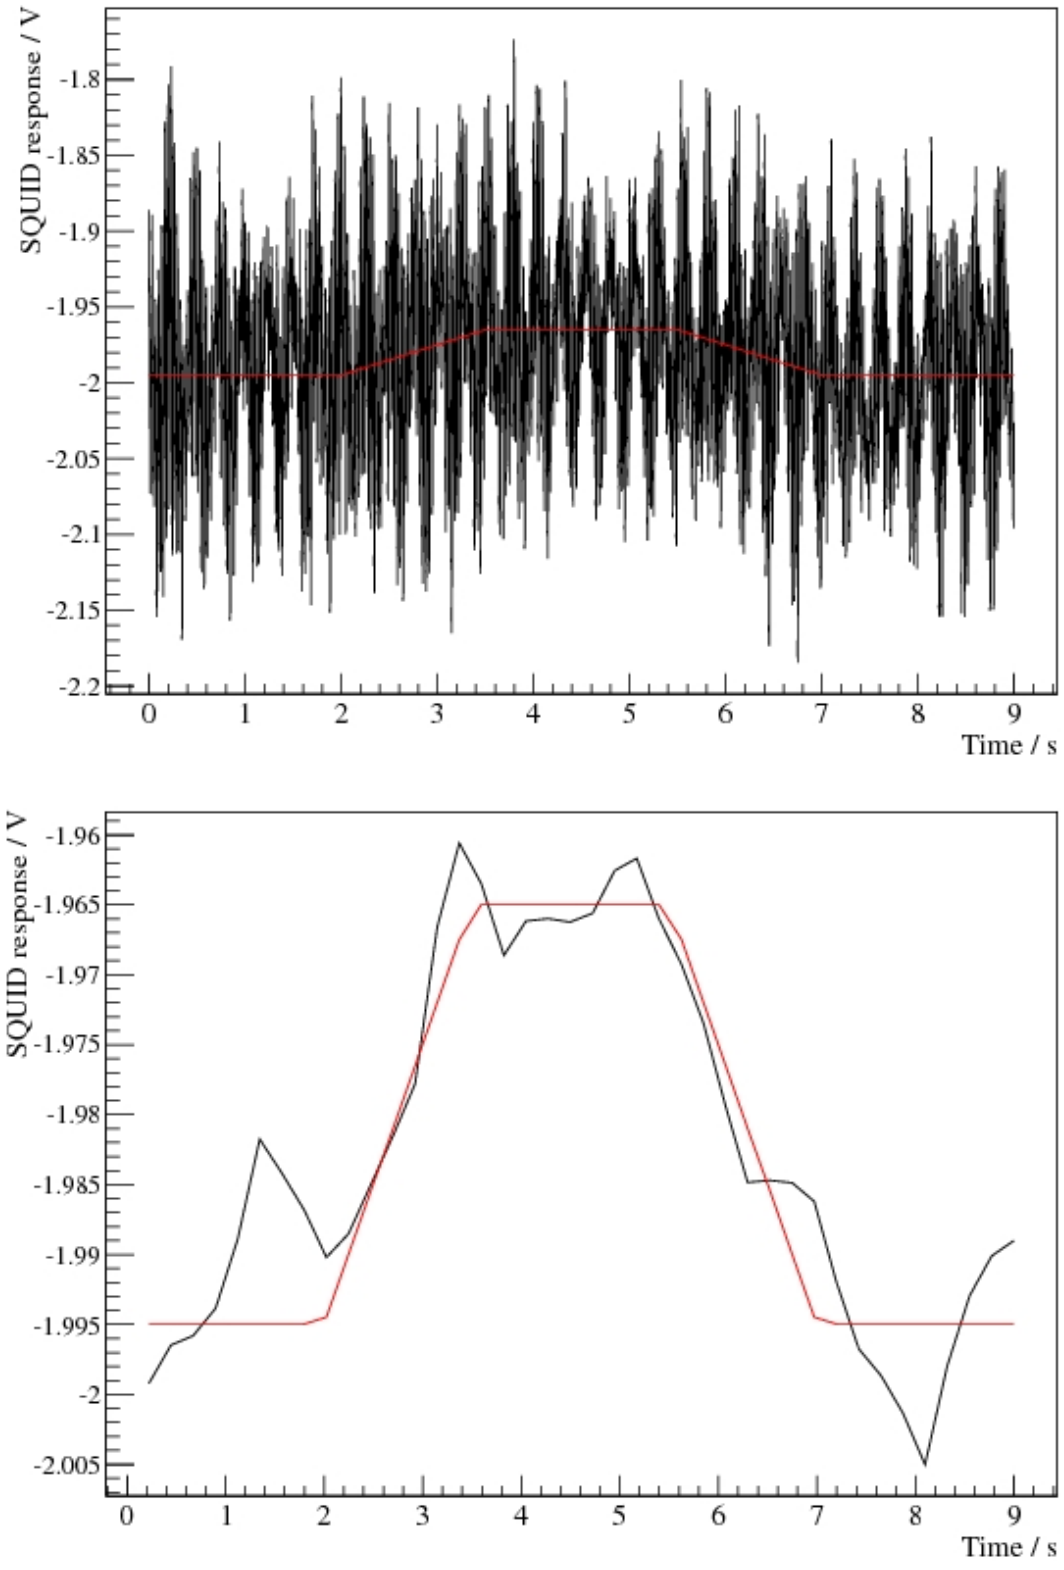

Figure 3.6: *(Top)* A calibration pulse as observed by a magnetometer, *(Bottom)* the same calibration pulse shown after smoothing the data over 200 samples. The red line shows the known pulse shape applied to the magnetometer. A correlation of 0.925 is calculated for the smoothed case and 0.176 for the raw data. Even by eye the pulse shape is far more obvious in the bottom diagram.

3.3.3 Main sources of uncertainty

There is no fundamental reason for a large error on the calibrations using this method, the numerical stability of the integration is far above what is required for these purposes, and fields calculated from the models are in good agreement with analytic results in simplified geometries, e.g., the on-axis fields from the solenoids. The main sources of uncertainty in the calculation come from uncertainties in the positions of pick-up loops and coils, which affect the flux-per-unit-current value for a calibration-coil/pick-up-loop pair, and the uncertainty in the response to the signal from the magnetometer. The positional errors can be minimised by accurate measurement of the positions and orientation of the components of the magnetometry system. In the 2010 data run there was a last minute change to the arrangement during installation of the equipment. The new positions of the pick-up loops and calibration coils were not measured to a high degree of precision. This gave uncertainties of order 1 cm in the positions, and typically lead to uncertainties in calibration of around 10%. It is expected that in more typical conditions the positions should be known to better than 1 mm.

The response uncertainty can be minimised by measuring many identical pulses, and by measuring at many different current values while constraining values via the linear response of the magnetometer.

3.3.4 Determining calibration factors

Once the magnetic flux through a pick-up loop has been calculated, and the magnetometer response has been measured, all that remains is to compare the two to provide the calibration factor. This is done by taking the SQUID response to different field values and from different coils to produce a field vs response dataset. A straight line through zero is fit to the data and the gradient is taken as the calibration factor. As the uncertainty distributions in the applied field need not be gaussian, a Monte Carlo is used to compute the uncertainty in the calibration factor. Using the magnetic field distributions and SQUID response distributions

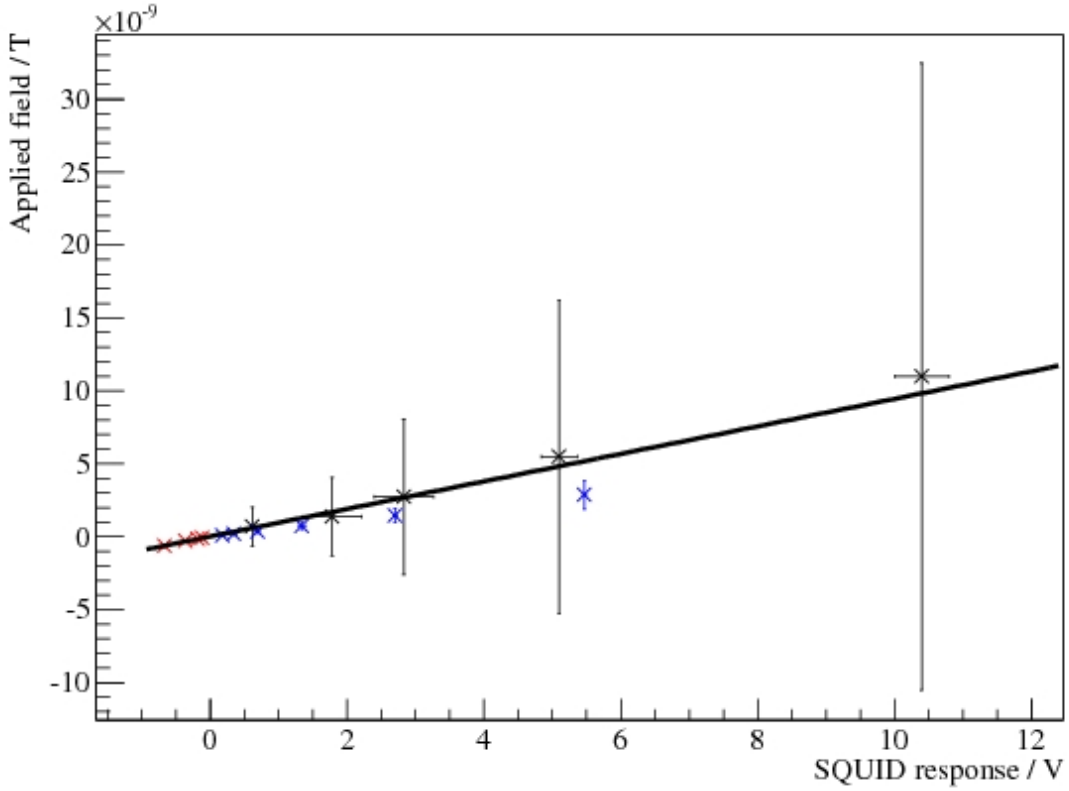

Figure 3.7: Field applied to a SQUID as a function of the response of SQUID 5. Differently coloured points come from different calibration coils. The solid line indicates the independently observed calibration using changes of field in the solenoid measured by the fluxgates. It is seen that data from two coils (displayed in black and red) are consistent with this calibration, but uncertainties in geometry make the data from the third coil (displayed in blue) inconsistent.

as inputs to the simulation, the distribution, and hence uncertainty in the calibration factors, can be computed for each pick-up loop.

The calibration routines were used in a data run in 2010, the positions of the magnetometers are given in figure 2.15. Figures 3.7 & 3.8 show two SQUIDs' response and applied field relationship from three calibration coils (black [cc6], red [cc7], and blue [cc5] crosses) with the uncertainties in the measured response and the applied field, along with an alternate calibration (black line) obtained in the next section. It is clear that the alternate calibrations are consistent with the calibration coils displayed in black and red, but the calibration coil displayed in blue is significantly different from the alternate calibration seen below in section 3.3.4.1. The electrical connections to this coil were checked, its impedance was as expected and no shorts to ground were apparent. This leaves the most likely explanation for

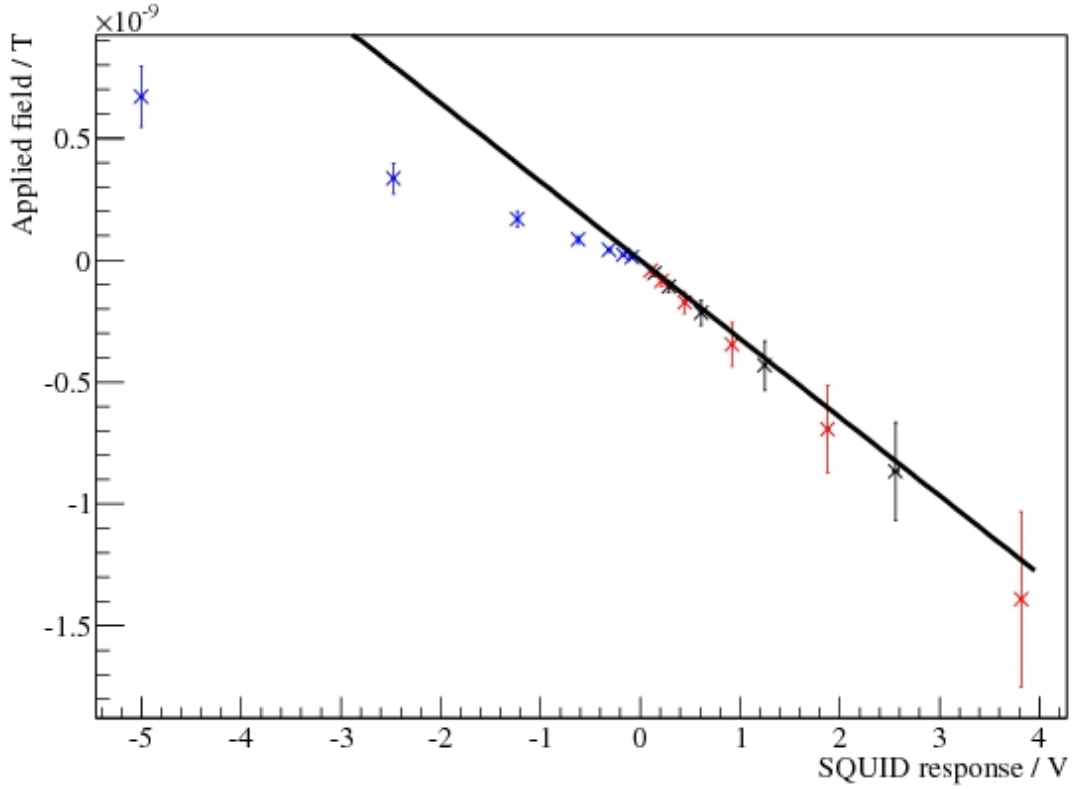

Figure 3.8: Field applied to a SQUID as a function of the response of SQUID 11. Differently coloured points come from different calibration coils. The solid line indicates the independently observed calibration using changes of field in the solenoid measured by the fluxgates. It is seen that data from two coils (displayed in black and red) are consistent with this calibration, but uncertainties in geometry make the data from the third coil (displayed in blue) inconsistent.

this discrepancy as incorrectly recorded data concerning the position and orientation of the calibration coil taken during installation. This is due to last minute changes in the pick-up loop assembly shortly before installation, given the time constraints the new positions were not measured as accurately as hoped. This highlights the importance of good positional measurements with this kind of calibration.

3.3.4.1 Comparison with known field changes

For axial, planar pick-up loops it is possible to compare calibrations obtained using the above method with those from the fluxgates installed. The field from the solenoid is axial, homogeneous enough that all magnetometers see, approximately, the same axial component

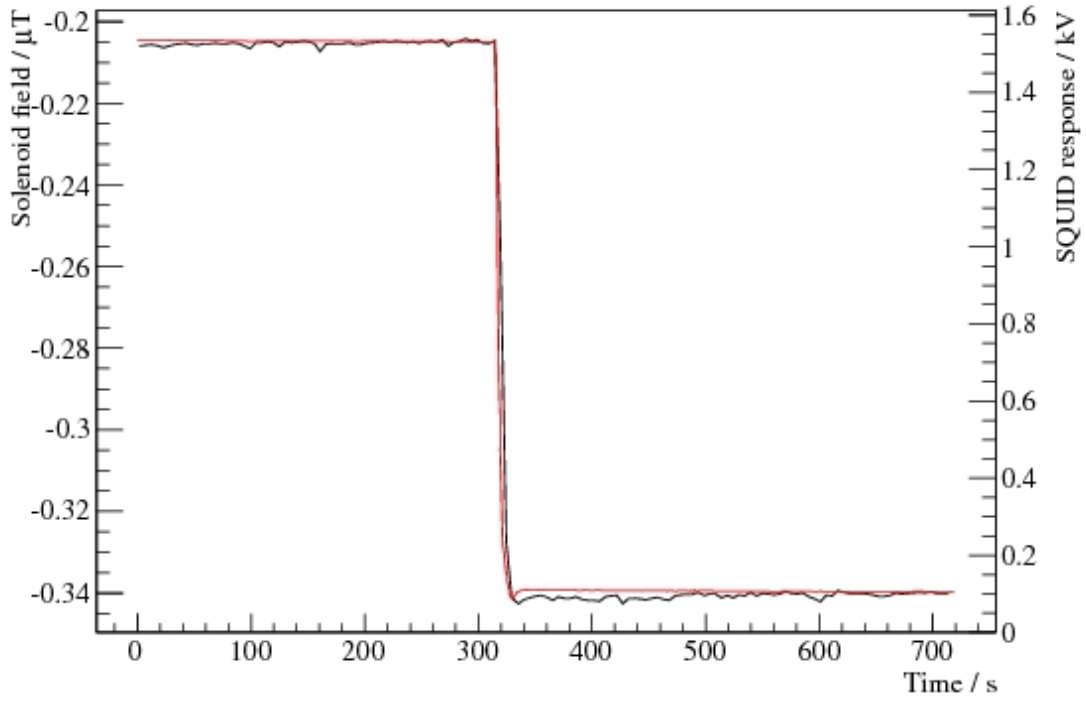

Figure 3.9: Example of an alternate calibration on SQUID 5. The black line shows the measured field changes using the fluxgates, this must be compared with the SQUID response (red). The best fit produces a calibration of approximately 1 nT V^{-1}.

of the field, and large enough to produce signals many orders of magnitude above the noise level. It is then possible to measure a large change in field from the solenoid using the fluxgates and compare this to the changes in the SQUID response. Figures 3.9 & 3.10 show an example of a large change in the solenoid field caused by a layer of superconducting shielding becoming normal. This change is observed on both the SQUIDs shown in figures 3.7 & 3.8 and internal fluxgates installed in cryoEDM. By fitting the SQUID's response to the fluxgate response and using the fluxgate's known calibration, the SQUIDs with axial pick-up loops can have their calibration checked.

Table 3.2 shows the best fit to the calibration pulses and the calibration observed from comparison with the fluxgates for the two axial SQUIDs, taken during the most recent data run. The values are clearly consistent but there is a large uncertainty on the value determined from the pulses due to uncertainties in the positions of pick-up loops and coils.

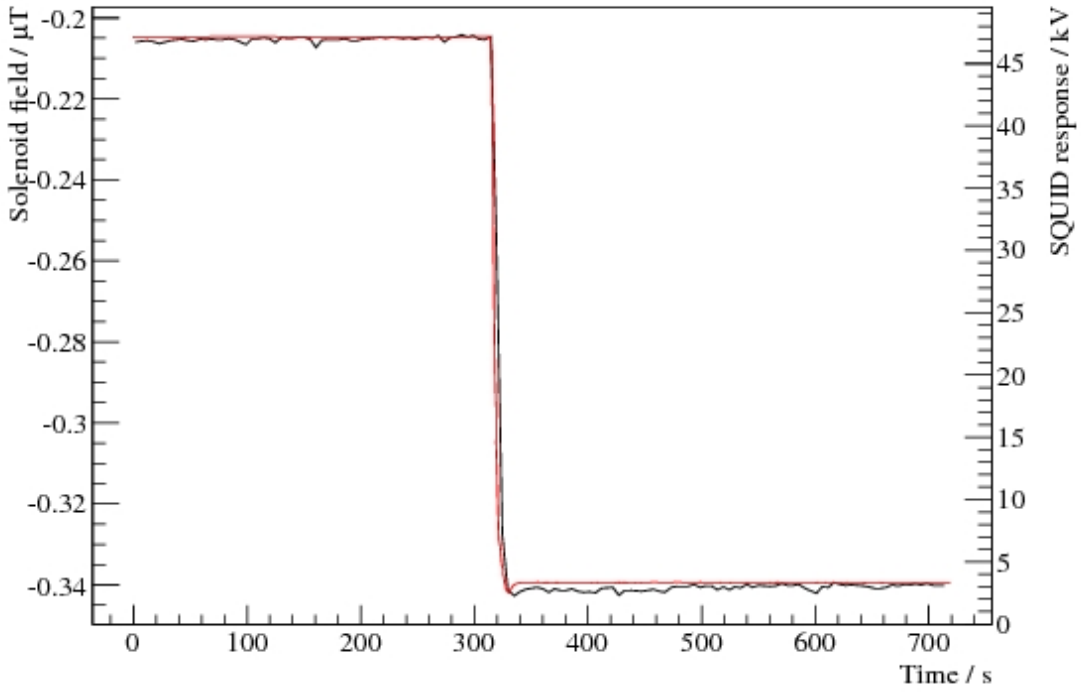

Figure 3.10: Example of an alternate calibration on SQUID 11. The black line shows the measured field changes using the fluxgates, this must be compared with the SQUID response (red). The best fit produces a calibration of approximately 0.03 nT V^{-1} respectively.

SQUID	Calibration from pulses / nTV^{-1}	Calibration from fluxgates / nTV^{-1}
5	-1.01 ± 0.28	-0.964 ± 0.009
11	0.0355 ± 0.0040	0.0314 ± 0.0003

Table 3.2: A comparison of the calibration values for axial pick-up loops measures from calibration coil responses and from comparison to the fluxgates.

3.3.5 Automation

Each of the above processes can be performed without interaction with a user, provided the geometry has been defined. This removes any subjectivity in determining responses and improves the number of usable calibration pulses from the tens an operator would be willing and able to locate to hundreds, thousands, or as many required to achieve an acceptable

uncertainty.

The total time to calculate the applied fluxes, and to take and analyse the response data is less than two hours when using of the order one thousand steps in the Monte Carlo simulations, of the order one thousand integral regions in the pick-up loops and of the order ten thousand line segments in the calibration coils. This means that a calibration for the SQUIDs would be available early in a data run, something not achievable before due to the complex operations required to obtain a calibration.

The calculation of applied fluxes can be performed offline, indeed this process would easily fit in the period between installation of the magnetometers when their positions are measured and the point where the experiment has cooled sufficiently to operate the SQUIDs, which can take several weeks. This would make the time taken to apply the calibration pulses the limiting factor, reducing the calibration time to several minutes. This process could easily be repeated after any changes to the SQUID system, which might change the calibration, allowing accurate calibrations to be achieved even in a changeable environment.

Additionally, many sources of failure are removed: there is no need to attach an external source to the correct coil; records of which coil is active, at what voltage and at what time are unnecessary; complications of controlling the DAQ systems is removed. These simplifications means that the calibration can be performed by someone unfamiliar with the magnetometry system, or even performed remotely.

3.4 Future calibration method

The second method of calibrating the magnetometers is to use the neutrons themselves as a way of measuring a reference field. Without an electric field applied the Ramsey resonance curve is just a function of the magnetic field the neutrons are experiencing. A formalism for determining calibration factors for the magnetometers using the data obtained during a neutron resonance will be presented.

3.4.1 Premise

After a Ramsey resonance measurement has been performed on the neutrons it is possible to reconstruct the magnetic field experienced by the neutrons during each measurement point. This is done by comparing the neutron polarisation to the theoretical curve in the same manner used to find a neutron EDM signal. In the presence of zero electric field there will be no neutron EDM contribution to the curve, and so a pure magnetic field measurement can be produced. A Ramsey, or Rabi, resonance curve can be plotted at a fixed applied AC frequency from a sequence of individual measurements taken at different solenoid fields. Performing such a resonance will yield a collection of precisely known magnetic field changes. If during this time the magnetometer response is also recorded then a set of field changes and magnetometer responses has been measured, and in principle a calibration may be performed.

To calibrate the magnetometers from measurements of the neutron resonance's field, all that is required is a relation between the field the neutrons experience and the field measured by the pick-up loops. This can be provided by using the analysis scheme developed in section 3.1, namely, equation 3.35.

$$\langle |\boldsymbol{B}| \rangle = \mathrm{sgn}(B_0) \sum_i^N B_{S_i}[\alpha + \beta z_i]$$

3.4.2 Calibration measurements

The field measured by a pick-up loop B_{S_i} is the output of the SQUID plus an arbitrary field offset, which remains constant during the SQUID's operation. This field may be written as

$$B_{S_i} = \frac{\partial B_z}{\partial V}_i S_i + O_i, \tag{3.53}$$

where $\frac{\partial B_z}{\partial V}_i$, is the calibration constant, S_i is the raw SQUID measurement in volts, and O_i the constant offset from the real field of that SQUID. Using this it is possible to rewrite

equation 3.35

$$B = \sum_{i}^{N} \frac{\partial B_z}{\partial V}_i S_i[\alpha + \beta z_i] + \sum_{j}^{N} O_j[\alpha + \beta z_j], \tag{3.54}$$

where B is a measurement of the field via the neutrons multiplied by the sign of the solenoid field, which is assumed to be constant for the calibration process.

The terms of the final sum in equation 3.54 are constant in time, these can be grouped together. Adding explicit time dependencies and letting the geometric terms $\alpha + \beta z = \mathrm{F}(z)$ we can simplify equation 3.54

$$\begin{aligned} B(t) &= B(0) + \sum_{i=1}^{N} \frac{\partial B}{\partial V_i} F(z_i) S_i(t) && (3.55) \\ \Delta B(t) &= B(t) - B(0) && (3.56) \\ &= \sum_{i=1}^{N} \frac{\partial B}{\partial V_i} F(z_i) S_i(t) && (3.57) \end{aligned}$$

The neutron resonance is performed over a time δt and so equation 3.57 must be averaged over that period

$$\begin{aligned} \int_{t_j}^{t_j+\delta t} \Delta B(t) dt &= \sum_{i=1}^{N} \frac{\partial B}{\partial V_i} F(z_i) \int_{t_j}^{t_j+\delta t} S_i(t) dt && (3.58) \\ \Delta B_j &= \kappa_i S_{ji} && (3.59) \end{aligned}$$

where ΔB_j is a measurement of the change in field as seen by the neutrons in zero electric field over a measurement cycle, and S_{ji} is the measurement from SQUID i over the same period.

At this point it is worth noting that as $F(z_i)$ is constant in time it can be absorbed by the $\frac{\partial B}{\partial V_i}$ calibration to produce an overall calibration factor, κ_i. For the purposes of this discussion only the overall calibration factor will be calculated, it is a simple matter to calculate the geometric factor and so to determine the calibration factor for an individual

pick-up loop, which might be useful for comparison with the calibration coils. However, as the calculation of the field experienced by the neutrons requires the product of the terms only, it is just as convenient to combine the two factors. This has the advantage of removing the requirement for precise knowledge of the magnetometer positions for calibrating the magnetometers, which was a limiting factor in the method using calibration coils. From now on it is taken that the geometric factor is part of the total calibration coefficient κ.

3.4.3 Determining calibration factors

Equation 3.59 define a set of linear equations, which, to recover the calibration coefficients, must be solved. There are numerous ways to solve such a set, the simplest is to form a matrix equation. This gives equation 3.59 the form

$$\boldsymbol{\Delta B} = S\boldsymbol{\kappa} \tag{3.60}$$

The simplest way to solve such an equation is to find the inverse of the matrix S, and apply to the vector $\boldsymbol{\Delta B}$. However the numerical stability of such an approach is not appropriate for our purposes, additionally a matrix inverse is only calculable for a square matrix, limiting the solution to include only as many data points as there are magnetometers to calibrate. Clearly to provide a good estimate of the calibration, and its uncertainty, it is desirable to use as many data points as possible in the calculation, and so this simplistic approach cannot be used.

For the sake of numerical stability and the desire to have an over-constrained set, and so an idea of the uncertainty, this equation is solved through a singular value decomposition (SVD) [100]. This is a matrix decomposition that is useful for finding the pseudoinverse of a non-square matrix, i.e. it can be used to find S^{-1} such that $\boldsymbol{\kappa} = S^{-1}\boldsymbol{\Delta B}$ for an over constrained set of equations. This is the value of the calibration that minimises $||\boldsymbol{\Delta B} - S\boldsymbol{\kappa}||$.

3.4.4 Uncertainty in fit

There are two ways to attribute an uncertainty to the result of the fitting. A method making use of the properties of the SVD method, and that of a Monte Carlo will be discussed.

3.4.4.1 Error ellipse

An advantage of the SVD method is that one of the components calculated during the decomposition represent a set of vectors that define the semi-axes of a normalised unit-less ellipsoid in the fitted-parameter space. This ellipse represents the region in which the reduced χ^2 of the fitting is within 1 unit of its minimum. This needs to be multiplied by the total uncertainty in the measurements, α, to produce a physical meaningful uncertainty. If the matrix equation is written in the form

$$\boldsymbol{\Delta B} = S\boldsymbol{\kappa} + \boldsymbol{v} \tag{3.61}$$

where $\boldsymbol{v}$ is a measure of the measurement noise, then α must satisfy $\alpha \geq |v|$. The simplest way of estimating this noise term is to sum in quadrature the standard deviations in the measurements i.e. $v_j^2 = \sigma_{\Delta B_j}^2 + \kappa_i^2 \sigma_{S_{ji}}^2$, this has the consequence that $|v|$ increases with each measurement. To scale the noise correctly this value must be divided by the degrees of freedom of the measurement [101], giving

$$\alpha^2 = \frac{\sum_{j=0}^{M-1} \left(\sigma_{\Delta B_j}^2 + \sum_{i=0}^{N-1} \kappa_i^2 \sigma_{S_{ji}}^2 \right)}{M - N} \tag{3.62}$$

Where M is the number of measurements performed and N the number of calibration factors to find.

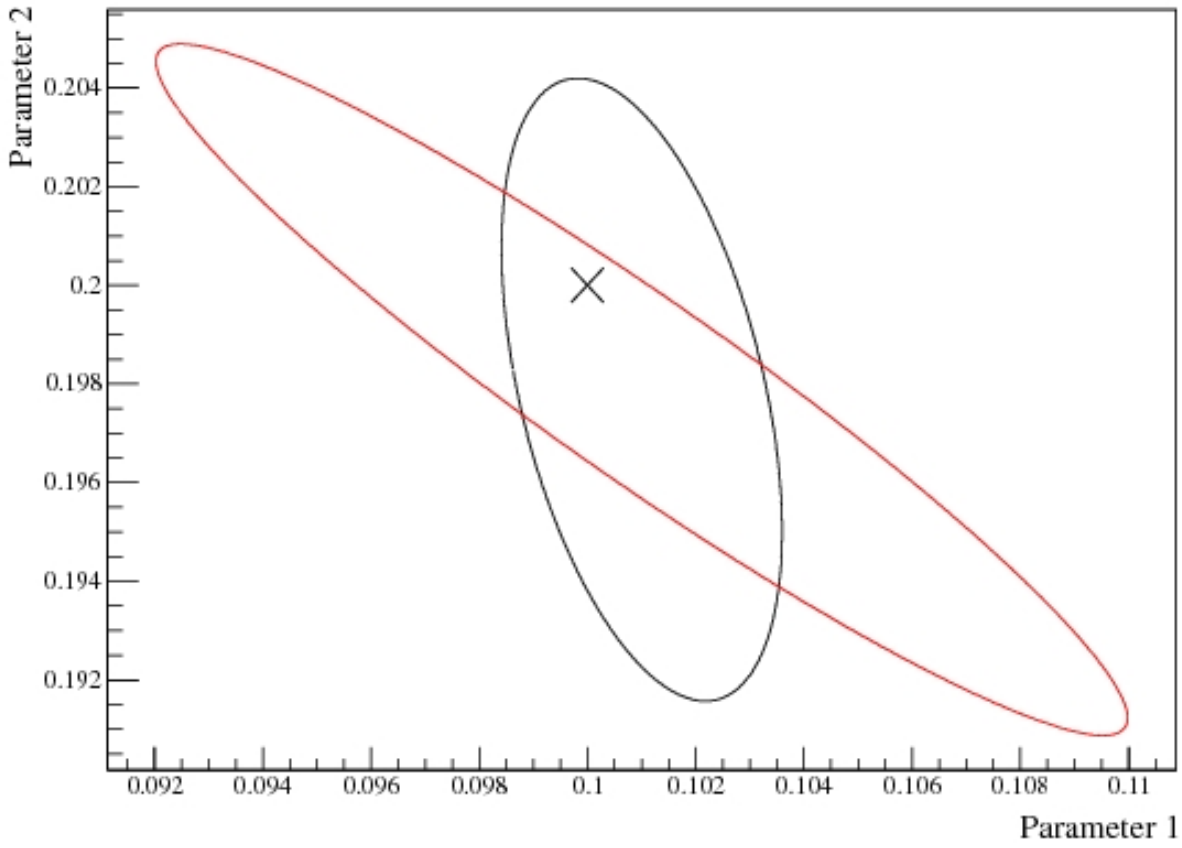

Figure 3.11: Shown is the 95 % confidence error ellipses produced by the Monte Carlo method (black), and as a by-product of the SVD method (red). The simulation uses two data points in the calibration and has inputs of 0.1 and 0.2 for parameter 1 and 2 respectively, shown with a cross. Clearly over this range the two methods are comparable.

3.4.4.2 Multiple fits

A more intuitive method is to perform multiple fits with elements of the matrix chosen randomly, using a Gaussian distribution with the measurements mean and standard deviation. A correlation matrix can be computed for each step of the Monte Carlo using the resulting fit as a new measurement. After sufficient steps, a distribution will emerge that will take into account the uncertainties in the fitting and those of the inputs to the fit. This distribution agrees well with statistically expected properties.

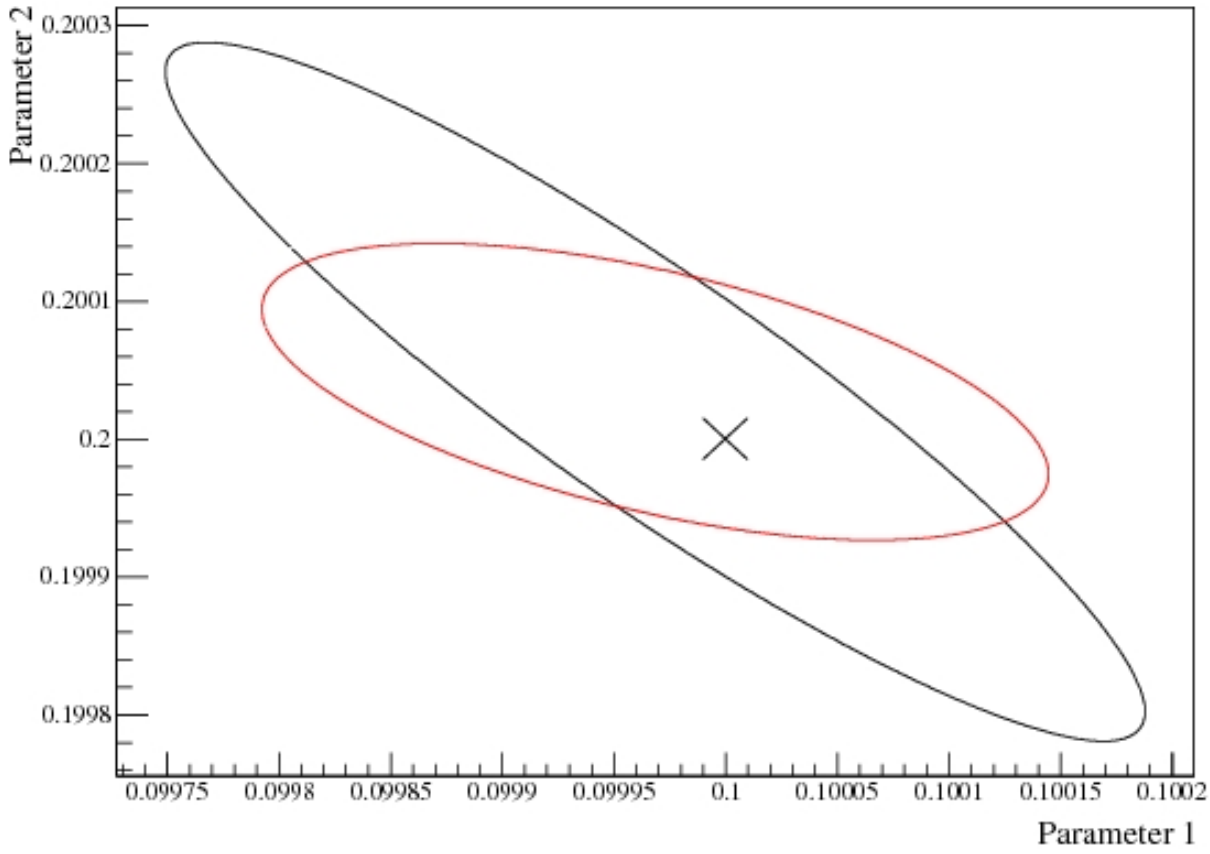

Figure 3.12: Shown is the 95 % confidence error ellipses produced by the Monte Carlo method (black), and as a by-product of the SVD method (red). The simulation uses 1000 points in the calibration, and has inputs of 0.1 and 0.2 for parameter 1 and 2 respectively, shown with a cross. As in figure 3.11 the two methods are comparable.

3.4.4.3 Comparison

Figures 3.11 & 3.12 show the results of a comparison of the error ellipses produced by both methods using a simulated data set in two dimensions, using set calibration coefficients and random changes in input signal a simulated output could be found. The algorithm was then used to retrieve the coefficients. The first simulation uses 2 points to determine the calibration coefficients and the second uses 1000 points. The two methods for obtaining an error ellipse compare well over this range, and both scale correctly as $\sim \frac{1}{\sqrt{N}}$. In general the Monte Carlo tends to produce a slightly larger ellipse than the SVD method and, for the sake of conservatism, shall be used from now on as the measure of uncertainty in calibrations.

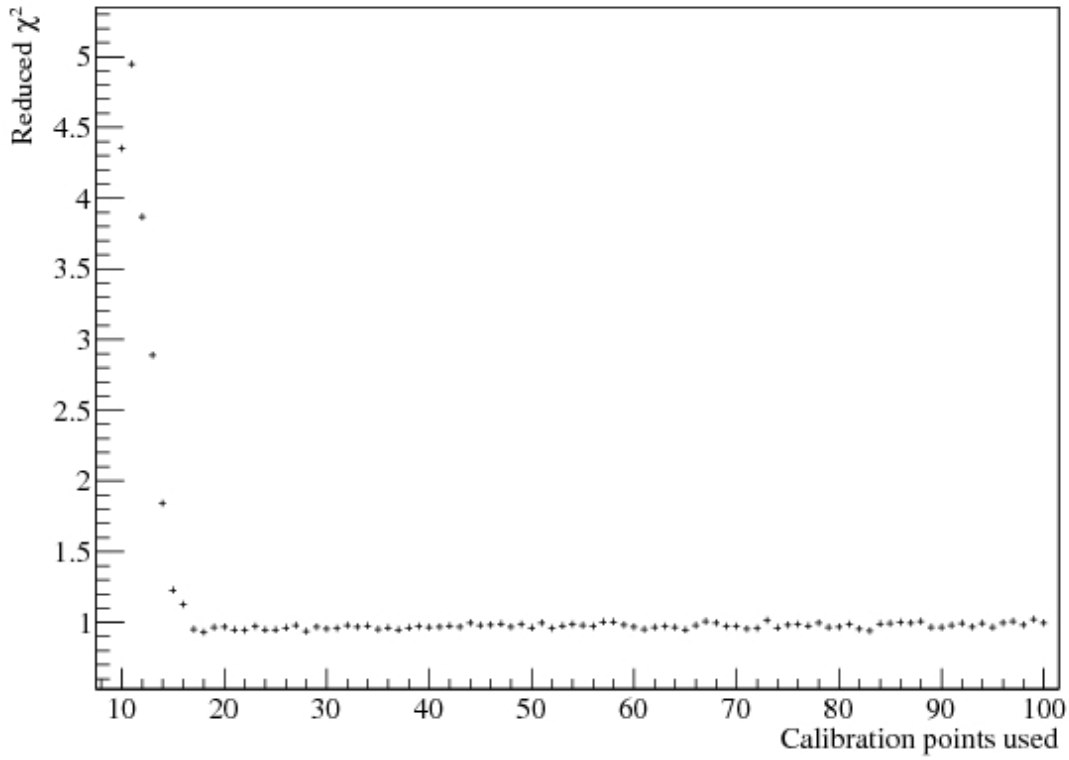

Figure 3.13: The results of determining the calibration in a simulation are shown, 10 calibration coefficients were to be determined. Plotted is the reduced χ^2 of the calculated calibration coefficients as a function of the number of points used to calculate them, it can be seen that this value drops and become stable around 1 after 20 points.

3.4.5 Simulation

The above algorithm was tested in a model as discussed above. A set of coefficients were predetermined and values for signal changes generated uniformly in the interval $[-1, 1]$. The reference field was then calculated using the known coefficients and signal changes. Simulating measured values for both the reference field and signal changes by computing a randomly generated value from a Gaussian distribution centred on the appropriate value, with standard deviation to simulate measurement errors. These values were fed into the calibration calculating algorithms to test their validity. Figure 3.13 shows the result of a simulation using 10 calibration coefficients. The reduced-χ^2 is plotted as a function of the number of reference points measured. This value quickly falls to 1 after a small number of additional data points. This demonstrates that the fitting is good, and the reported

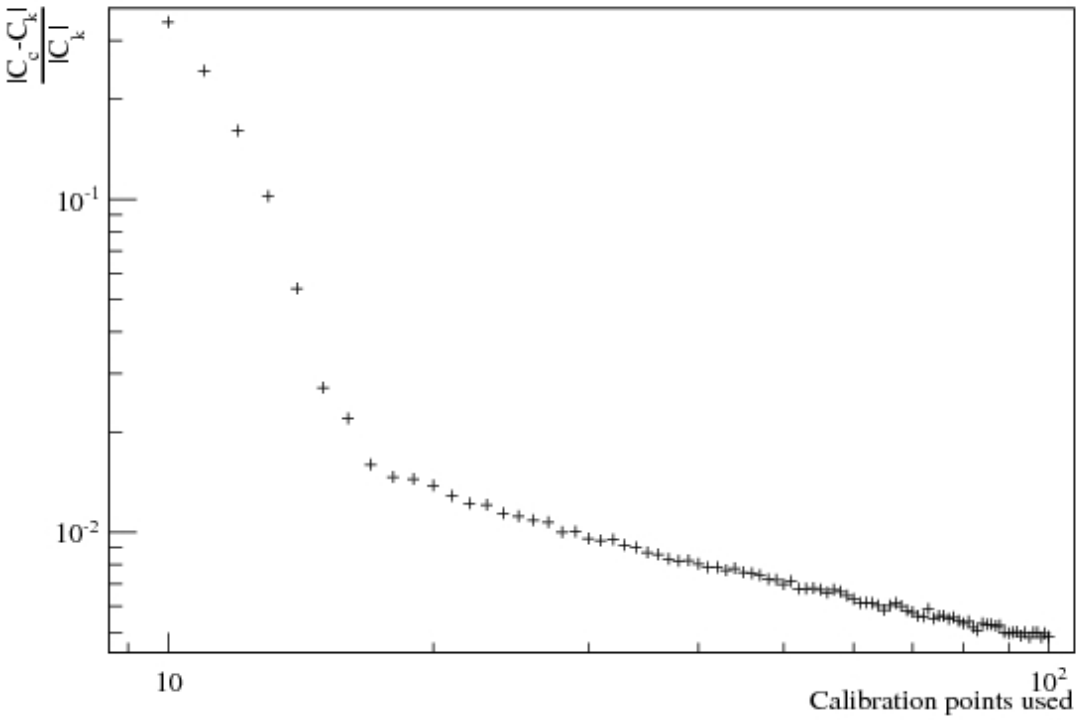

Figure 3.14: The results of determining the calibration in a simulation are shown, 10 calibration coefficients were to be determined. Plotted is the ratio of the moduli of the calibration error vector and the calibration vector as a function of the number of points used to calculate them. This value can be seen to drop significantly until nearly 20 points are used, and then decreases as $1/\sqrt{N}$ as the number of points used increases, as expected.

uncertainties are of the correct order. Figure 3.14 shows the ratio of the modulus of the error vector, the difference between observed and input calibrations, and the modulus of the calibration vector. Again, this falls quickly after a small number of additional points, and reduces slowly afterwards. This procedure was carried out for numbers of fitted parameters from 5 to 50. Empirically the point at which the fit becomes stable appears to occur after the degrees of freedom is larger than about half the number of coefficients to fit. Then to calibrate all 12 SQUIDs would require around 20 points, which is easily obtainable from a neutron resonance.

3.4.6 Proof of principle

A neutron resonance is yet to be observed in the experiment, and so a true test of this form of calibration was impossible. To give credibility to this method of calibration a small test set-up was created in an attempt to mimic the conditions found in cryoEDM.

As a proof of principle, data were taken using a set of five identical coaxial pick-up loops displaced from each other axially. The aim is to reconstruct the measurement of the end most SQUID from the measurements on the others, after calibration. This will be a measure of the surface average of the magnetic field not the volume average as expected, but this only changes the geometric factor, which is of no concern for this test. Figure 3.15 shows the arrangement used for this test. Five well defined pick-up loops formed from 20 mm diameter rings made of lead-tin coated copper tracks on a PCB, and connected to the SQUIDs using a niobium-titanium twisted wire pair through superconducting joints were used in the test. The pick-up loops are separated by 4 mm.

This is a factor of 15 smaller than the 300 mm maximum size of pick-up loop intended to be used in the full experiment, and scaled has a similar separation between the pick-up loops, which may vary but is typically between 50 mm and 100 mm. However, this arrangement does not, strictly, match the geometry of cryoEDM primarily due to the lack of a superconducting shield surrounding the region of interest, and so the assumption of linearity may be less valid. However the distance between the loops used here is short, and so it is hoped that this approximation will be close enough to true to conduct the test. The purpose of this measurement is simply a test of the algorithm to find calibration factors, so there is significant leeway in the permissible model.

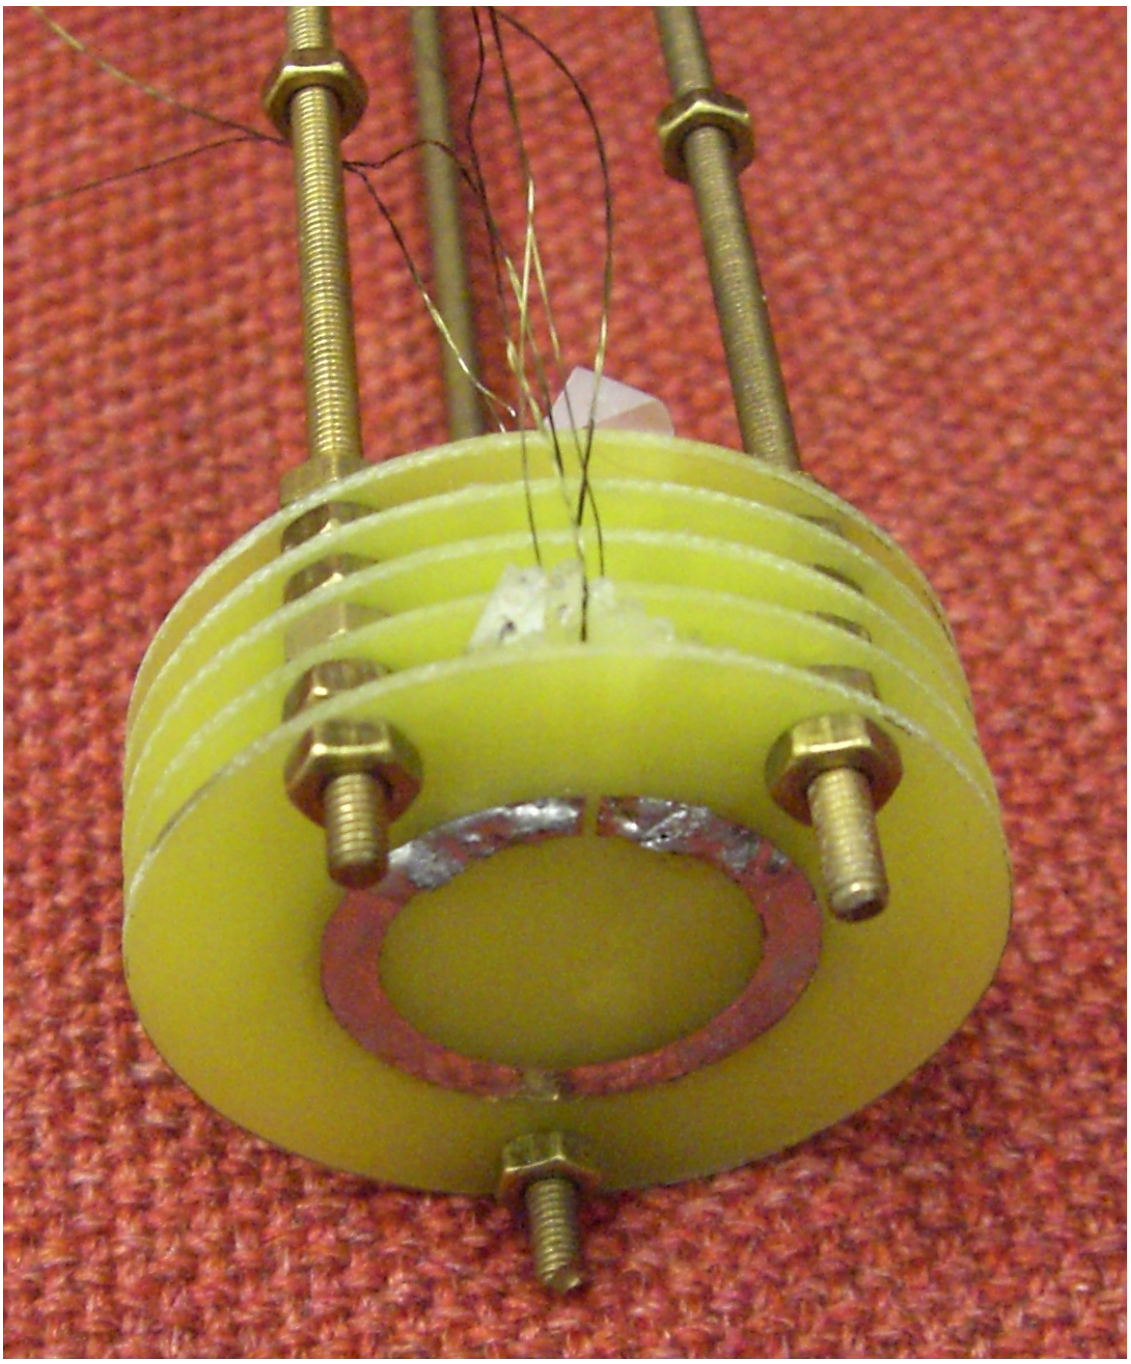

Figure 3.15: Five lead-tin coated copper tracks on PCBs forming the pick-up loops connected to five SQUIDs. The loops are arranged coaxially and have a separation much shorter than the size of the loops. This set-up was used to test the calibration algorithms.

Figure 3.16 shows data taken by the five SQUIDs for one hour at a low noise laboratory in France, the Laboratoire Souterrain à Bas Bruit (LSBB). All five SQUIDs show a similar signal and have RMS magnetically induced noise in the range 10 mV to 30 mV. The electronics

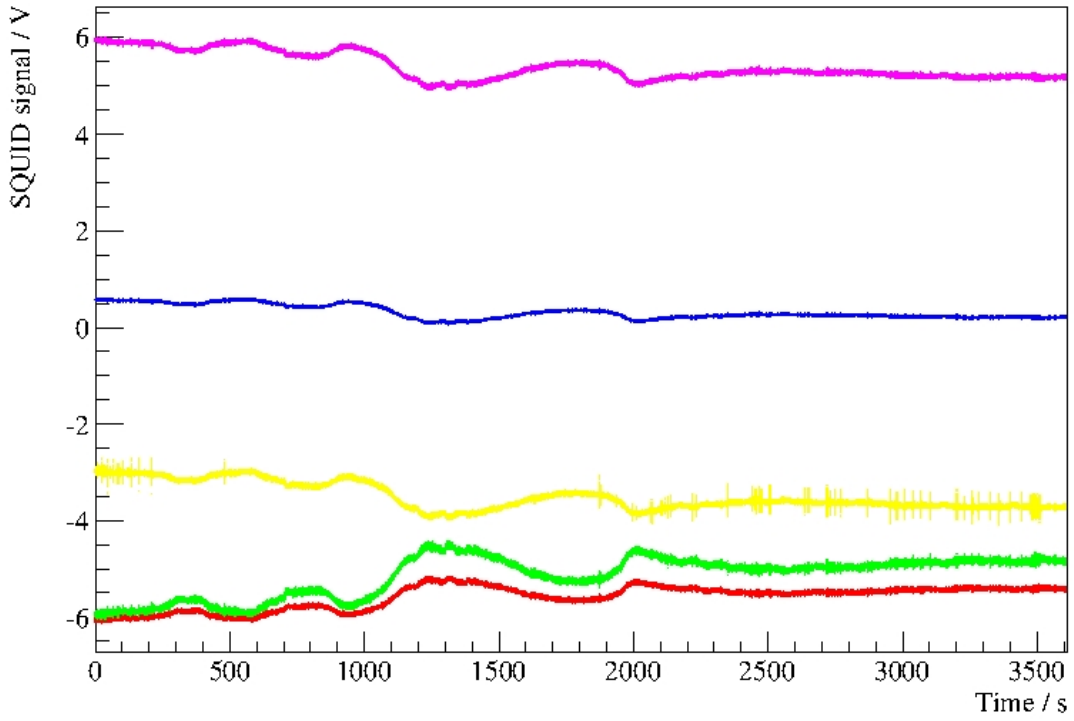

Figure 3.16: Shown is one hour of SQUID data taken at the LSBB taken at 1 kHz. The RMS noise of the electronics was 0.55 mV. The RMS noise of the channels is 17.6 mV (red), 29.1 mV (green), 10.1 mV (blue), 21.7 mV (yellow), and 19.6 mV (purple).

with no SQUID attached have a RMS noise of 0.55 mV. No calibration coil was present so the noise in terms of field values can not be quoted.

The algorithm was used to find the calibration factors of the first four SQUIDs using 600 seconds of the data. This calibration region was split into integral regions 30 seconds long, on which the calibration was performed. The results of applying the calibration on the next hour of data is shown in figure 3.17 with the residuals plotted in figure 3.18. The χ^2 between the observed and reconstructed signals of the last SQUID is 0.58, and the RMS of the residuals between the two is 14 mV. The residuals show a drift by 10 mV from the measured value after an hour, however the uncertainty in the reconstruction is approximately 50 mV in this region and the deviation is a similar size to the noise on the channels.

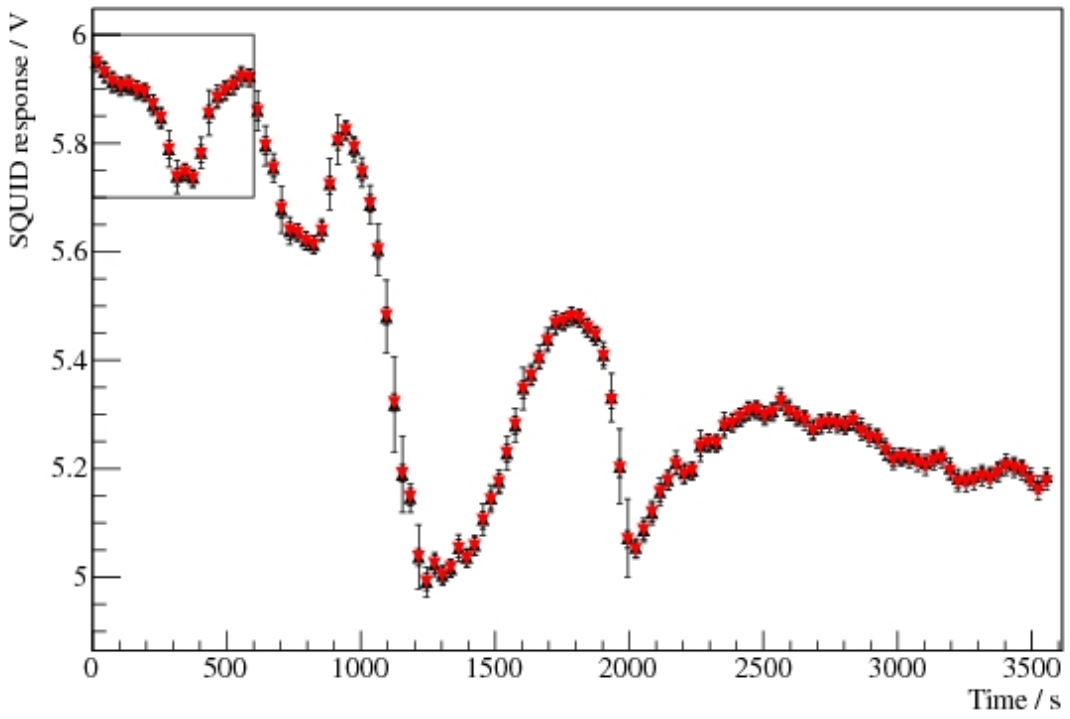

Figure 3.17: Shown is one hour of SQUID data taken at the LSBB. The box in the upper left indicates the region of data used for calibration. The χ^2 between the observed (red) and reconstructed (black) signals is 0.58 and the RMS of the residuals is 14 mV.

The stability of the calibration coefficients with time was computed by choosing different calibration regions from which to obtain the coefficients, and is plotted in figure 3.19. Regions beginning 500 s apart over 7000 s were were taken from the data, and a set of calibration coefficients computed from them. Each calibration coefficient is highly stable with time, showing that an arbitrary calibration region may be chosen, and also that regions with large field changes provide similar results to relatively small field changes.

All combinations of using one, two, three, and four SQUIDs to predict another were computed, the value for the average χ^2 and RMS residual is tabulated in table 3.3. The average RMS residual is also plotted in figure 3.20. There is a consistent improvement in the residuals as the number of SQUIDs used in the reconstruction is increased.

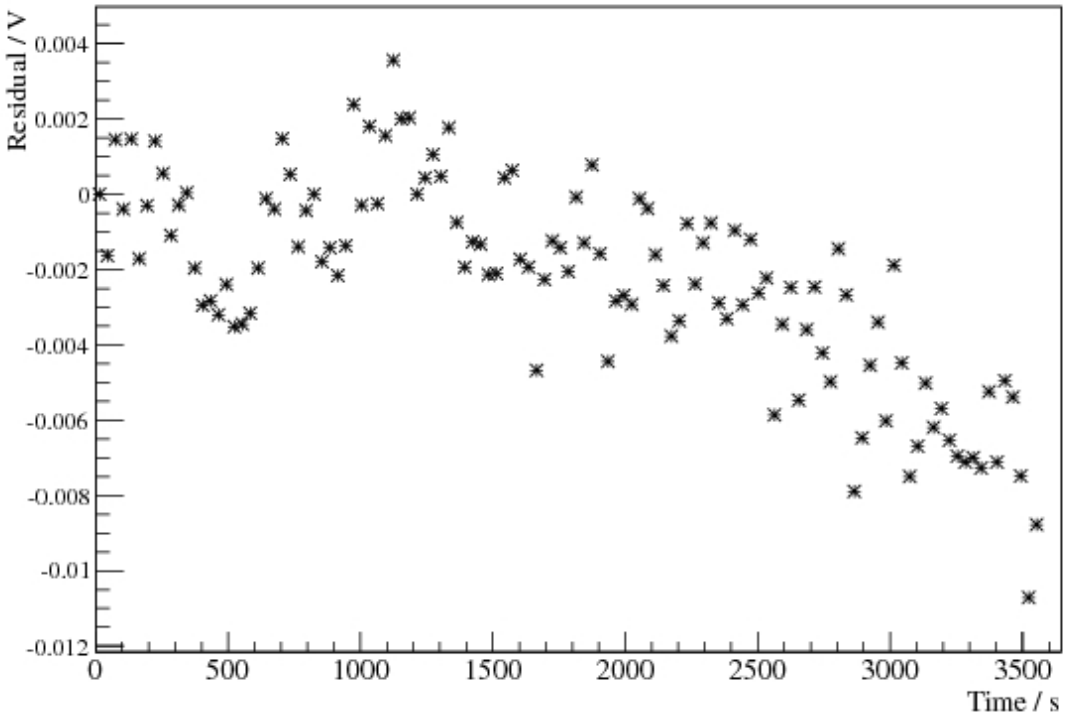

Figure 3.18: Shown are the residuals between the observed and reconstructed signals seen in figure 3.17. There appears to be a drift of about 10 mV, which is approximately 1 % of the total change observed during this period, but the uncertainty in the reconstruction at this point is approximately 50 mV.

Number of SQUIDs	Average χ^2	Average RMS of residuals / V
1	0.87	$2.09(13) \times 10^{-2}$
2	0.92	$1.54(8) \times 10^{-2}$
3	1.07	$1.43(7) \times 10^{-2}$
4	1.16	$1.38(8) \times 10^{-2}$

Table 3.3: The average χ^2 and RMS of residuals of the extrapolation for various numbers of SQUIDs used.

This improvement in RMS residual shows that, even in this non-ideal case, the reconstruction of a field can be obtained from a number of spatially separated measurements with better agreement to the actual field than any individual measurement.

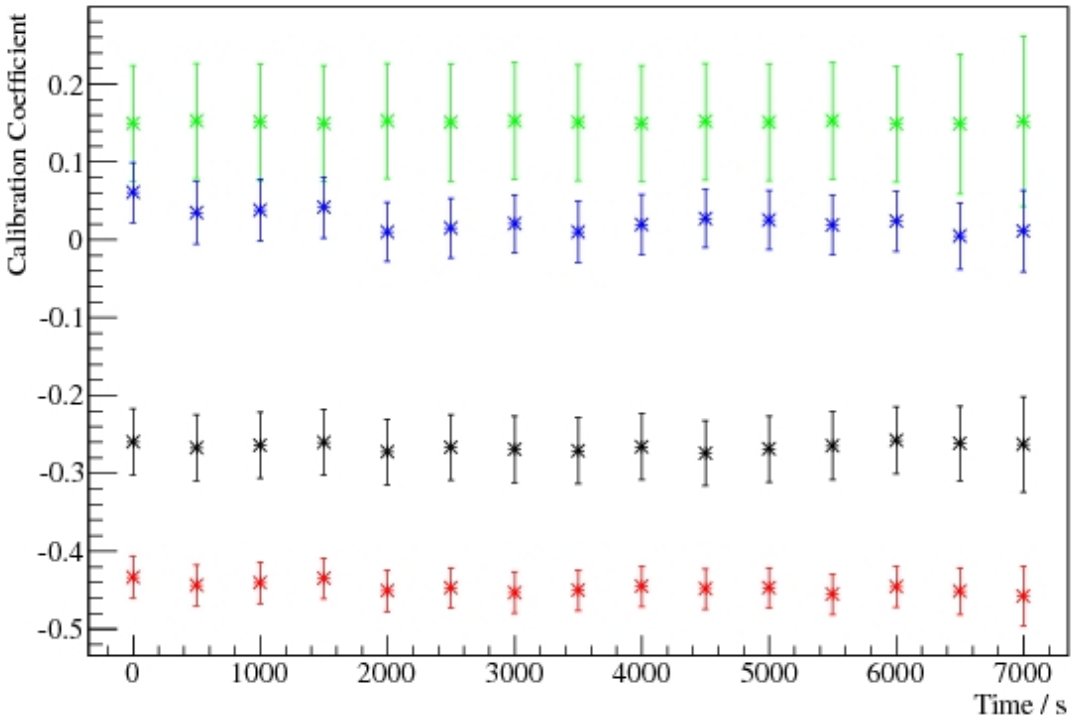

Figure 3.19: Shown are the calibration coefficients for the four SQUIDs as a function of time. The calibrations are very stable over the data set, a two hour period.

Finally, combinations that reconstructed signals from SQUIDs at the end of a set, and in between the extremes of a set were both computed. Figure 3.21 shows RMS residuals for both cases as a function of the number of SQUIDs used in the reconstructions. The "end" SQUID case shows the same reduction in residuals as the full data set, however the "between" SQUID case shows a fairly constant RMS residual. This indicates that the magnetometry should not suffer for having pick-up loops only on one side of the Ramsey cells. This result is presumably because a linear interpolation between two points can't be improved as much by additional points as extrapolation away from the points would.

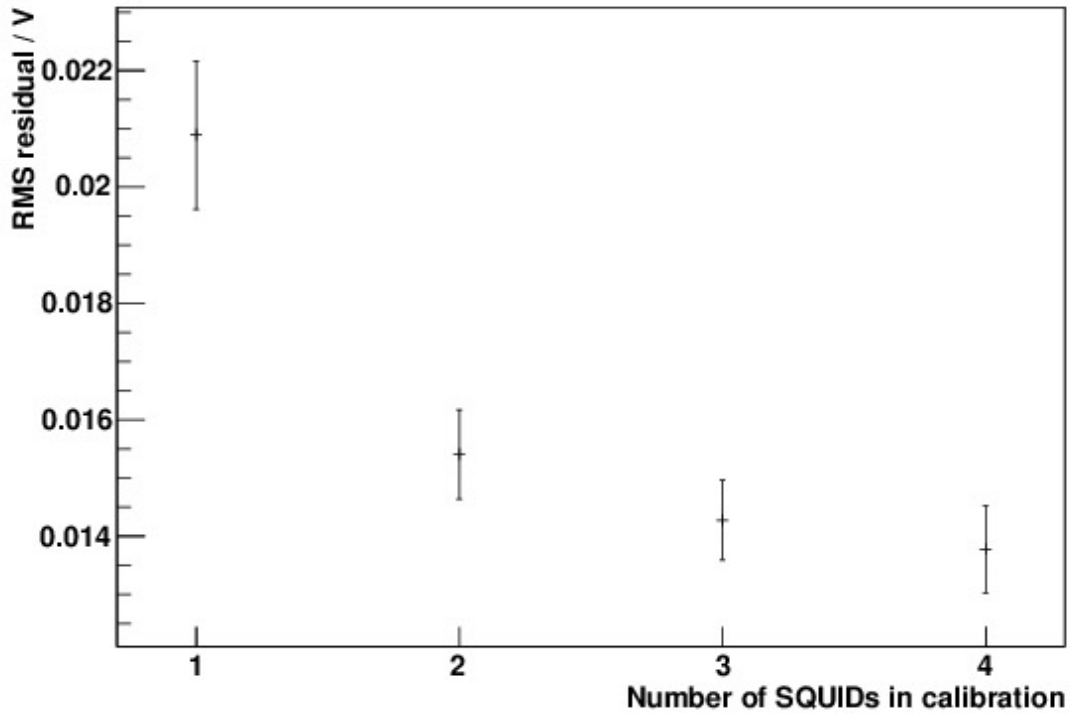

Figure 3.20: Shown is the average RMS of the residuals between the observed and the reconstructed signals for all combinations of the number of SQUIDs used in the reconstruction. It can be seen that the residuals decrease as the number of SQUIDs used increases. Meaning that the reconstruction is improved by adding more SQUIDs, as expected.

3.5 Summary

Two new systems of calibrating the SQUIDs in cryoEDM have been investigated, one using the existing calibration coils installed in the experiment, the other using a neutron resonance to obtain the calibration data.

The calibration using calibration coils has been show to work successfully in the running experiment, and to agree with calibrations obtained from comparison with the fluxgates. The limiting factor of this method is the knowledge of the position and orientation of the pick-up loops and calibration coils.

The second method of calibration has been show to work well on a small test set-up with similar geometry to the cryoEDM experiment, providing calibrations of the SQUIDs to fields in the region of interest without requiring knowledge of the geometry. Additionally, as this

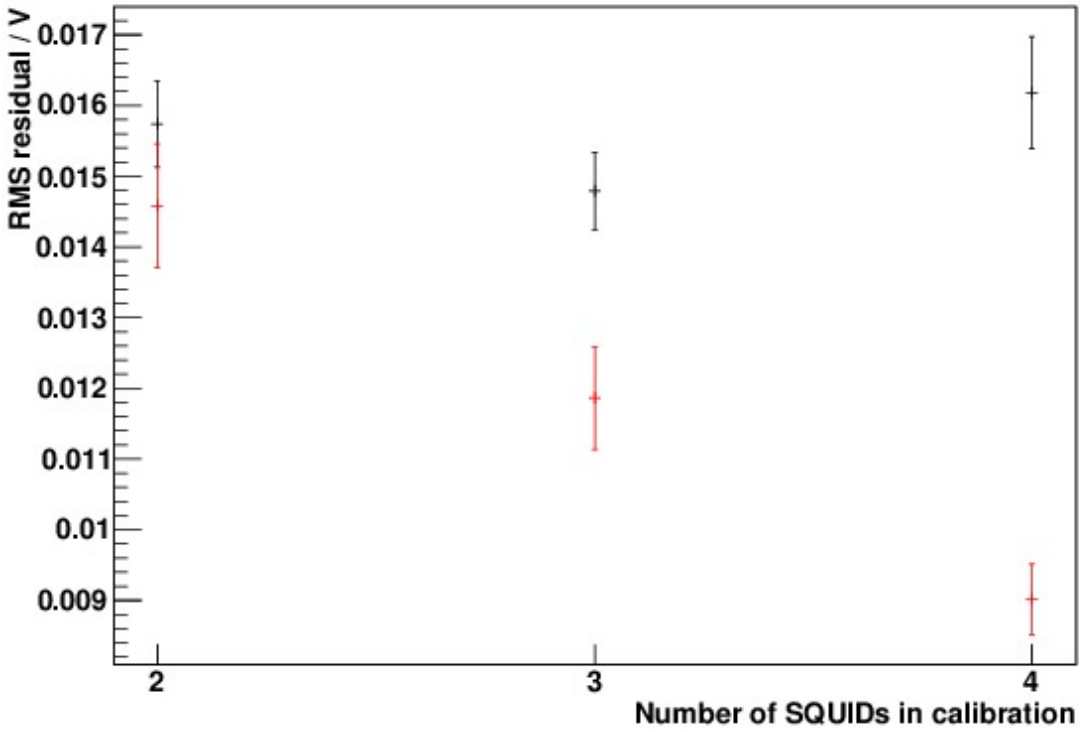

Figure 3.21: Shown is the average RMS of the residuals between the observed and the reconstructed signals for combinations reconstructing "between" SQUIDs (black) and "end" SQUIDs (red) of the number of SQUIDs used in the reconstruction. For the "end" SQUIDs it can be seen that the residuals decrease as the number of SQUIDs used increases as expected. However the "between" SQUIDs have a fairly constant RMS residual, this is presumably because once two points either side of the reconstructed SQUIDs are set only small deviations to this will occur. It is observed that the reconstructions of "end" SQUIDs have a lower RMS residual than the "between" SQUIDs which is unexpected.

method uses neutron measurements to determine the calibration, the effect of a non-uniform neutron distribution should be reduced.

To produce a fully functional magnetometry system all that remains is to correct the incoming data for hardware created artefacts.

Chapter 4

SQUID feedback artefact correction

To undertake the analysis and calibrations outlined in chapter 3, the outputs of the magnetometers must represent the magnetic flux observed through the magnetometer's sensing region in a linear fashion. While in general this is the case, the outputs of SQUIDs can have hardware introduced artefacts, which makes this relation not fully linear. These artefacts were discussed in section 2.7.2.1. A description of algorithms for their removal will be presented, and the performance of these algorithms on simulated and real data will be given.

4.1 Correction algorithms

4.1.1 Description

4.1.1.1 General

In general the algorithms were defined in terms of the magnetometer output $V(t)$, the magnetic field induced signal $V_B(t)$, and the offsets caused by artefacts $V_A(t)$ such that

$$V(t) = V_B(t) + V_A(t). \tag{4.1}$$

This can be differentiated to give

$$\frac{dV}{dt} = \frac{dV_B}{dt} + \frac{dV_A}{dt}. \tag{4.2}$$

We are only concerned about this signal during measurements, and so can define a region of interest between two times t_0 and t_1. We can integrate equation 4.2 over this region, giving

$$\int_{t_0}^{t_1} \frac{dV}{dt} dt = \int_{t_0}^{t_1} \frac{V_B}{dt} dt + \int_{t_0}^{t_1} \frac{dV_A}{dt} dt \tag{4.3}$$

$$V(t_1) - V(t_0) = V_B(t_1) - V_B(t_0) + \int_{t_0}^{t_1} \frac{dV_A(t)}{dt} dt, \tag{4.4}$$

If the artefacts are sufficiently quick they appear as a step function in the output, this leads to the gradient of the artefact induced signal being as follows

$$\frac{dV_A}{dt} = \sum_i V_i \delta(t - \tau_i), \tag{4.5}$$

where V_i are the step changes caused by artefacts occurring at times τ_i. Equation 4.4 can then be written as

$$V(t_1) - V(t_0) = V_B(t_1) - V_B(t_0) + \int_{t_0}^{t_1} \sum_i V_i \delta(t - \tau_i) dt \tag{4.6}$$

$$V(t_1) - V(t_0) = V_B(t_1) - V_B(t_0) + \sum_i \begin{cases} V_i, \text{ if } t_0 \leq \tau_i \leq t_1 \\ 0, \text{ otherwise.} \end{cases} \tag{4.7}$$

This can be rearranged to give a value for V_B

$$V_B(t_1) = V_B(t_0) + V(t_1) - V(t_0) - \sum_i \begin{cases} V_i, \text{ if } t_0 \leq \tau_i \leq t_1 \\ 0, \text{ otherwise} \end{cases} \tag{4.8}$$

and descretised

$$V_B(t_{j+1}) = V_B(t_j) + V(t_{j+1}) - V(t_j) - \begin{cases} V_i, \text{ if } t_j < \tau_i \leq t_{j+1} \\ 0, \text{ otherwise.} \end{cases} \tag{4.9}$$

As arbitrary initial offsets can be absorbed during the calibration process, it is possible to write

$$V_B(t_0) = V(t_0) \tag{4.10}$$

without affecting the results. Using equations 4.9 & 4.10 it is possible to determine the magnetic field response of the magnetometer from measurements of its previous outputs. It only remains to determine if between times t_j and t_{j+1} a jump has occurred and its size.

If the sampling rate is shortened significantly and the inductance of sensing circuits is increased such that the time constant of the coupling between the feedback and SQUID coils is very large, then it is possible for a reset artefact to extend across several samples (A reset itself takes only 2 μs, however to propagate the field changes through the system takes non-zero time). Much of the above argument is equally valid for these cases, the only modification is the number of points used to determine the presence of an artefact. This presents a problem if the outputs of the algorithm are used before a full artefact is detected, as outputs during an artefact could propagate to the analysis before it is corrected. To account for this, in general a delay buffer is introduced to ensure the full artefacts are detected before earlier outputs are used. Outputs during an artefact are then linearly interpolated between the known previous value before the artefact, and the newly reconstructed value after.

This process of detecting and reconstructing artefacts will be detailed in the next section.

4.1.1.2 Algorithms

There are two constraints for the algorithms that must detect and compensate for artefacts. The first is accuracy, a high detection rate and reconstruction accuracy is needed to give a true representation of the field. The second is speed, to act in the experimental control system the corrections must be done in real time, short delays may be added if compensated for but all channels must be corrected before the next sample is ready to be processed.

Many artefact correction algorithms were developed, a brief description of some will be given here.

1. A simple threshold, changes above a certain threshold value are ignored.

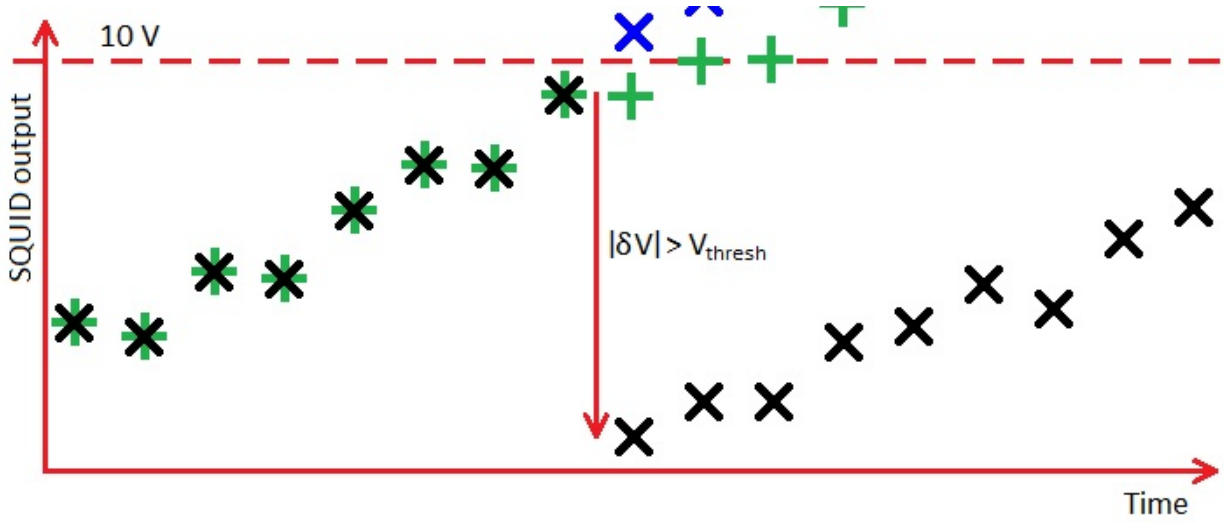

Figure 4.1: Shown is the operation of the threshold algorithm on a reset. A black 'X' represents the SQUID output, a blue 'X' represents where the output would be if not for an artefact. A green '+' represents the output of the algorithm.

   This is shown graphically in figure 4.1, clearly this approach always underestimates the value of a jump. However this is by far the quickest algorithm and useful in low noise environments for a fast field estimate.

2. A simple threshold with compensation, changes above a certain threshold have an integer number of flux quanta removed.

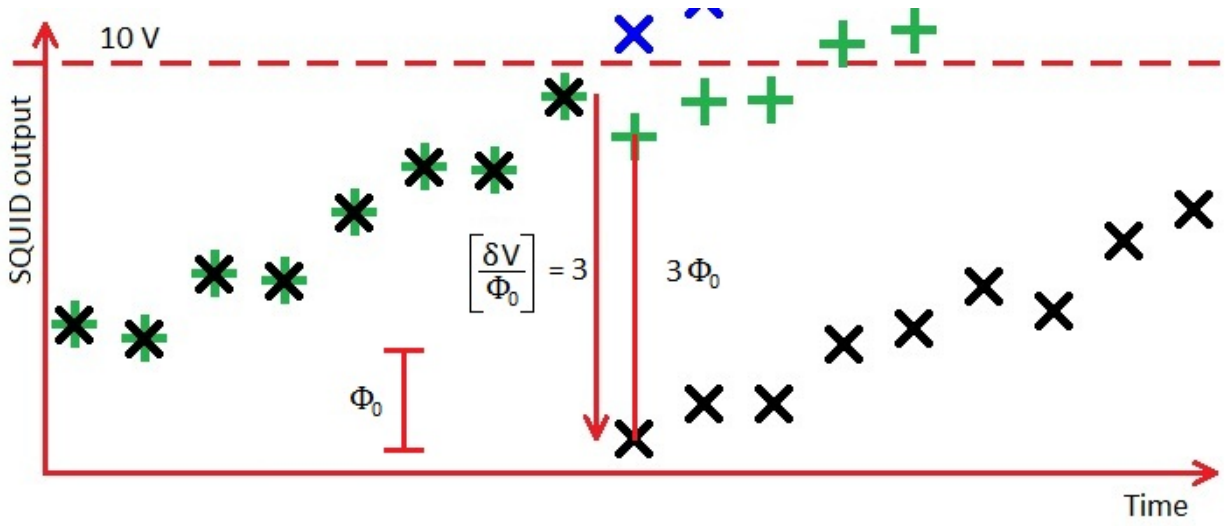

Figure 4.2: Shown is the operation of the threshold with compensation algorithm on a reset. A black 'X' represents the SQUID output, a blue 'X' represents where the output would be if not for an artefact. A green '+' represents the output of the algorithm.

This is shown graphically in figure 4.2, where the algorithm attempts to determine the number of flux quanta missing, but in the case shown the change in voltage between points is rounded incorrectly to 3 flux quanta rather than 4. Again this is useful in low noise environments where the change in voltage between points is small. In this environment the algorithm will be fast and correct the signals with superior results compared to the previous algorithm. For better results in noisy environments a better estimate of the voltage change is required.

3. A simple prediction, the average gradient of previous points is computed, and the current point is assumed to be different by this gradient. Any difference between the expected and observed value has an integer number of flux quanta removed.

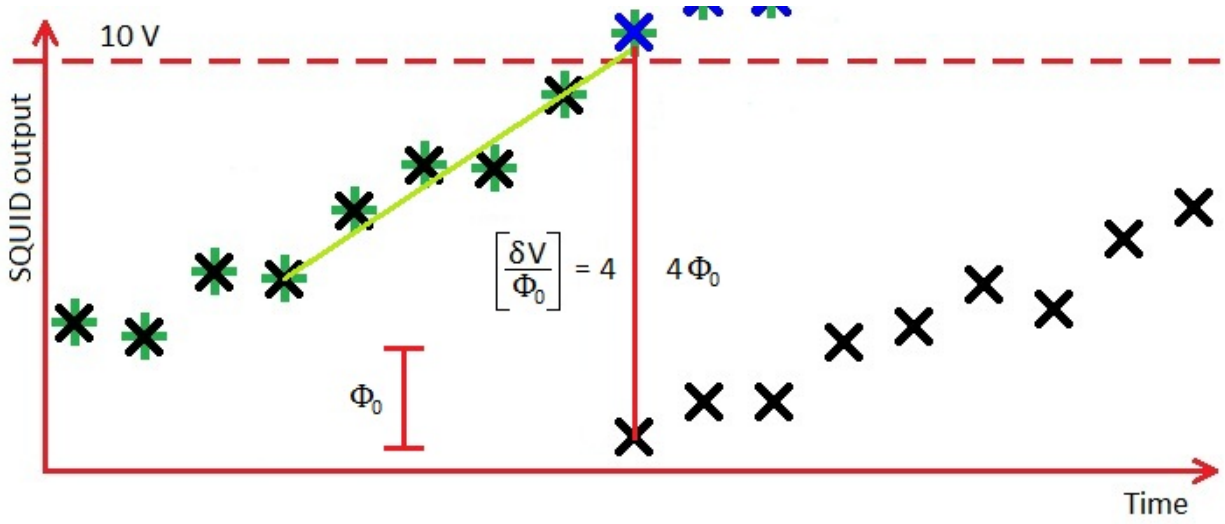

Figure 4.3: Shown is the operation of the average gradient algorithm on a reset. A black 'X' represents the SQUID output, a blue 'X' represents where the output would be if not for an artefact. A green '+' represents the output of the algorithm. The green line represents the algorithm's prediction using the average gradient of the previous points extrapolating to the new point.

This is shown graphically in figure 4.3, where the algorithm uses the previous five points to calculate an average gradient of the output, and adds this gradient on to the previous value to estimate where the current point should be. The difference between the predicted point and the actual point is used to determine the size of extraneous voltage changes. Due to the extra points processed this algorithm is significantly slower than the previous algorithms, but is more immune to noisy environments.

4. A simple prediction with spike detection, similar to algorithm 3, but delays the output by a point to check if any change is a noise spike rather than an artefact needing correction. This algorithm is useful when there are also spikes in the data stream, single points that are not recorded correctly by the ADCs and make an abrupt change that returns to the expected value after the next sample. Subsequent modification to the DAQ hardware has eliminated such spikes from the data.

5. A linear prediction, a linear fit is performed to a number of previous points, the current

point is extrapolated along this line, any difference between the expected and observed value has an integer number of flux quanta removed.

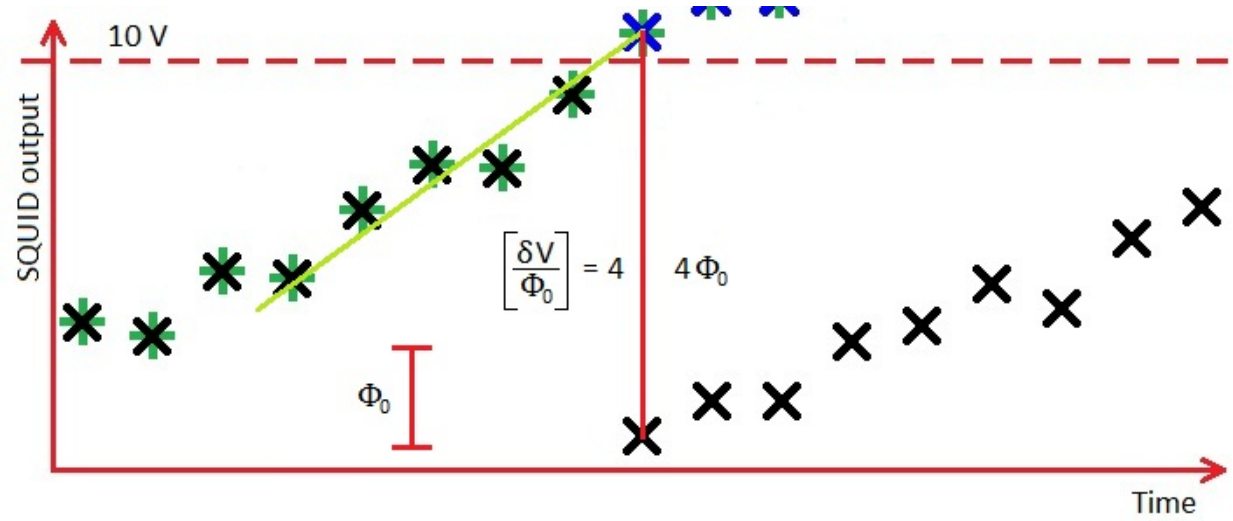

Figure 4.4: Shown is the operation of the linear fit algorithm on a reset. A black 'X' represents the SQUID output, a blue 'X' represents where the output would be if not for an artefact. A green '+' represents the output of the algorithm. The green line represents the algorithm's prediction using a linear fit to the previous points and extrapolating to the new point.

This is shown graphically in figure 4.4, where the algorithm collects the previous five points and fits a straight line through them. This line is used to extrapolate forward to where the current point should be. The difference between the predicted point and the actual point is used to determine the size of any artefact. This algorithm extrapolates to the next point very well in all but the noisiest conditions, again the price is increased correction time caused by the more complicated calculations.

6. A precognitive linear prediction, a linear fit is performed to a number of previous points and future points (taken using a delay line), if the difference between the forward and backward predictions is above a threshold the two lines are forced to intersect at the current point.

7. A precognitive linear prediction with compensation, a linear fit is performed to a

number of previous points and future points, an integer number of flux quanta are added to minimise the difference between the forward and backward predictions.

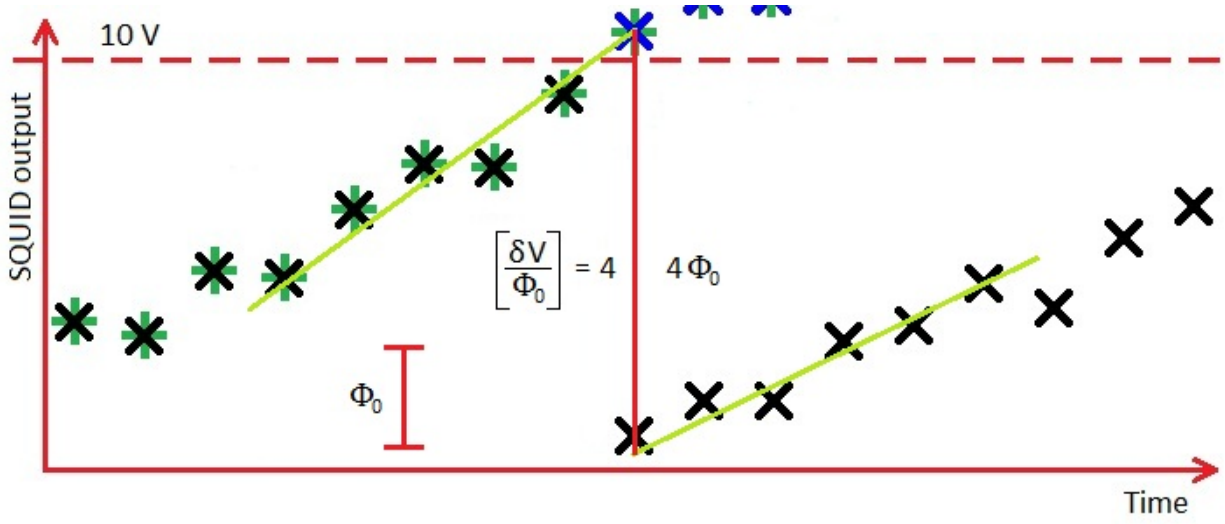

Figure 4.5: Shown is the operation of the precognitive linear fit algorithm on a reset. A black 'X' represents the SQUID output, a blue 'X' represents where the output would be if not for an artefact. A green '+' represents the output of the algorithm. The green line represents the algorithm's prediction using a linear fit to the previous and future points.

This is shown graphically in figure 4.5, where the algorithm fits a straight line to five points either side of the region of interest and calculates the size of any artefact from the difference of the predictions at the current point.

The precognitive algorithms are useful in situations where an artefact can take multiple points, as they use samples away from the artefact itself to identify it. Of course an artefact will appear in the 'future' samples before entering the region of interest, which will make any linear fit through the future points erroneous and cause the correction to be incorrect. To counteract this, a goodness-of-fit check is performed on the data, if this fails the algorithm falls back to using just the previously analysed values.

Each algorithm has a variant where multiple points can make up the region of interest, not just one. These algorithms each take different amounts of time to correct the data, but

generally increasing up the list. Each algorithm also has slightly different efficiencies in both detection and reconstruction, a fact that will be explored in section 4.2.

4.1.2 Finding the flux quantum

Many of the algorithms above require knowledge of the voltage representing a flux quantum, either to compensate for large changes or to set a threshold value. In general this value can be obtained from the SQUID tuning software. Small increments of magnetic field offset are applied to the SQUID through the feedback system, and after each increment a reset is forced. If the SQUID can settle into an operating point with lower feedback value then the output will change by one flux quantum. However, this relies on the feedback signal being the only applied field, with a pick-up loop connected the environmental field is also applied, and this is far less controlled. In a noisy environment the value of flux quantum returned by this method can be highly inaccurate.

The alternative is to take a significant amount of data in the noisy environment and histogram the point changes in the signal. This should produce peaks for the general noise, centred around zero, peaks for the resets near $\pm 10\,\mathrm{V}$, and peaks for the flux losses. Resets can be created in a controlled manner by applying large fields to the magnetometer via the calibration coils, and so large numbers can be generated easily. Flux jumps are essentially random and so a large amount of data is required to extract a good value for the flux quantum. Figure 4.6 shows these peaks, however, the flux jump peak is difficult to resolve from the central noise peak. Figure 4.7 shows Gaussian fits to flux jumps of up to three flux quanta, allowing a value for the flux quantum to be obtained even in a noisy environment.

The asymmetry of the peak is caused by the linearisation hardware searching for a new working point by increasing the feedback field in a single polarity, producing flux jumps in the same direction. To reduce the field and produce flux jumps in the opposite direction an intermediate reset would be required, therefore it is far easier for flux jumps to occur in one direction than the other.

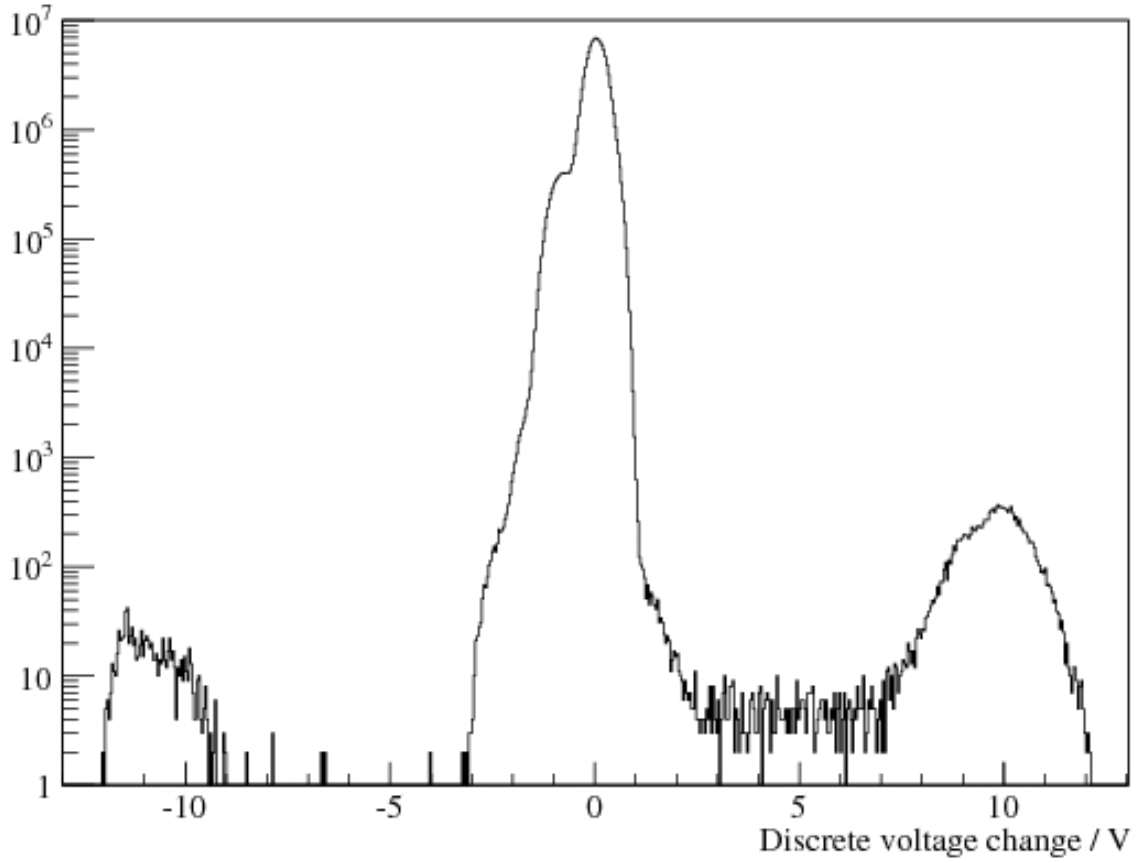

Figure 4.6: A Histogram showing the discrete voltage changes of a SQUID output taken over 30 hours. A sharp central peak can be seen, which is due to magnetic noise detected by the magnetometer. The isolated peaks at approximately ±10 V are due to feedback resetting, and the peak at -0.85 V is the collection of flux jumps and is difficult to resolve from the noise.

Additionally the value of the reset peaks can be used to add extra statistics, as the reset peaks should occur at an integer flux quanta.

In cases where the flux quantum is within the general magnetic noise of the system, a better tactic may be to set the flux quantum to be the size of a reset artefact and accept that flux jumps, which should be rare, cannot be identified. Similarly a threshold number of flux quanta can be set so that flux quantum sized differences in reset changes are handled, but within-the-noise changes are ignored. However, the magnetometry system is not intended to run in such noisy environments.

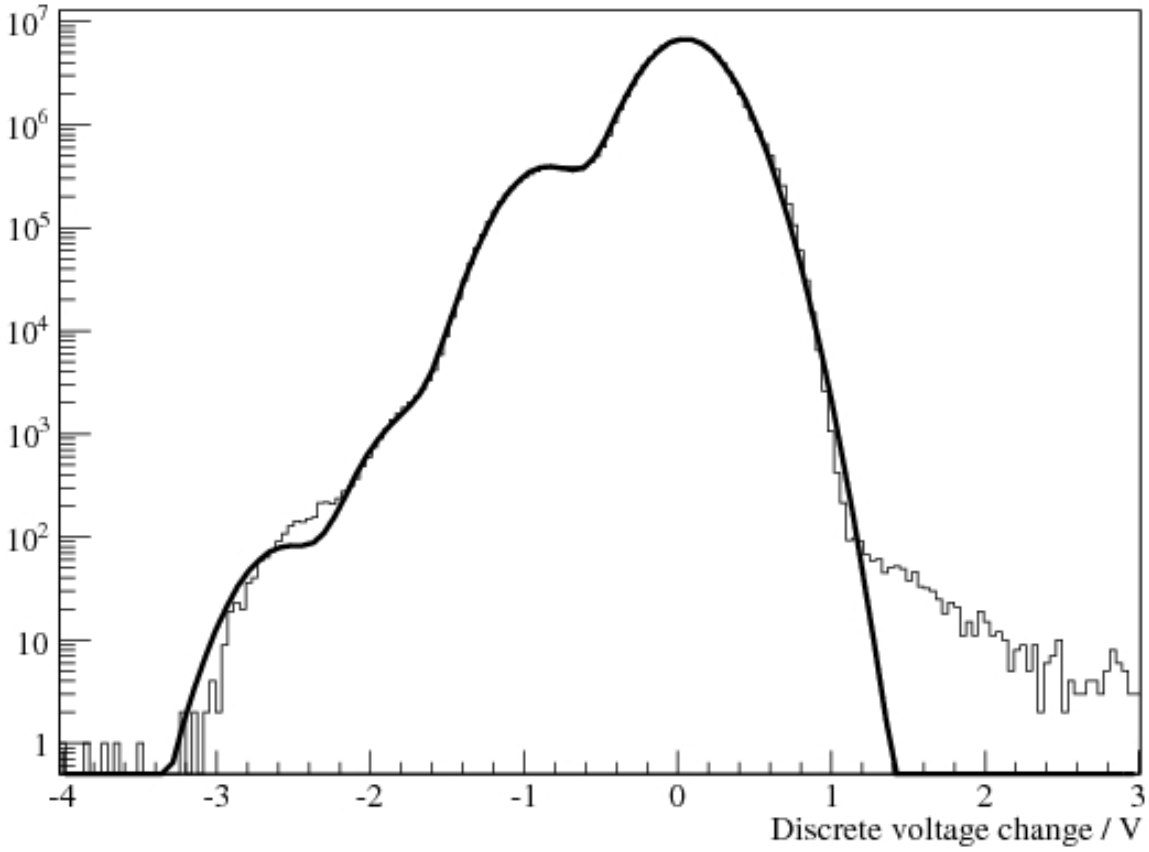

Figure 4.7: A Histogram showing the discrete voltage changes of a SQUID output as in figure 4.6. A Gaussian fit to the central noise peak and to $-\Phi_0$, $-2\Phi_0$, and $-3\Phi_0$ with the same widths is shown. This fit give a value of $\Phi_0 = 0.85\,\mathrm{V}$. So even in such a noisy environment a value of the flux quantum can be obtained.

4.1.3 Data quality

It is convenient to attribute a value for the data quality to the magnetometer output, so that certain outputs may be excluded from further analysis if they are noisy or unstable. It is during correction that many useful parameters are calculated, which might give a representative figure to the data quality. For example a number of resets in quick succession may indicate a loss of lock in the feedback system, especially if similar field changes are not seen on other magnetometers.

Additionally, multiple correction algorithms can be run on the input data, provided there is enough time available to run them. The data sets could be compared to determine any significant differences that could be used to attribute a quality to the data. For example if

the threshold and linear fitting correction algorithms do not detect an artefact in the same location, this would indicate the quality of the data used in the fitting region is low.

4.2 Performance

4.2.1 Simulated data

It is possible to construct a reasonable simulation of the SQUID response to different stimuli. The SQUID response will be the field applied to the magnetometer, with an offset and a measurement noise applied. The noise is applied simply using a Gaussian random number generator, centred on the value with the appropriate width to produce a white noise level. In this discussion the low frequency noise, which scales as $\frac{1}{f}$, shall be ignored as it can not interfere with the correction measurements due to its low frequency. The offset is calculated as an integer number of flux quanta: to bring the output within a specified range, simulating resets; and increments or decrements if the signal plus noise changes are greater than a predefined slew rate, simulating flux jumps. Both artefacts are treated as instantaneous changes in this simulation.

4.2.1.1 False positives and negatives

The first performance metric of the correction algorithms is how well they identify artefacts. Two important quantities are the number of data points misidentified as having artefacts (false positive), and the number of artefacts unidentified (false negative). The SQUID output simulation is aware of whether it has introduced an artefact or not so it is a simple matter to compare this to the artefacts found by the correction algorithms.

Figure 4.8 shows the number of false positives per 10000 simulated data points for each algorithm as a function of the SQUID noise expressed as a ratio to the voltage representing a flux quantum, Φ_0. Most algorithms show no false positive corrections when the FWHM noise is less than $0.1\Phi_0$. There is a sharp transition in the number of misidentifications when the

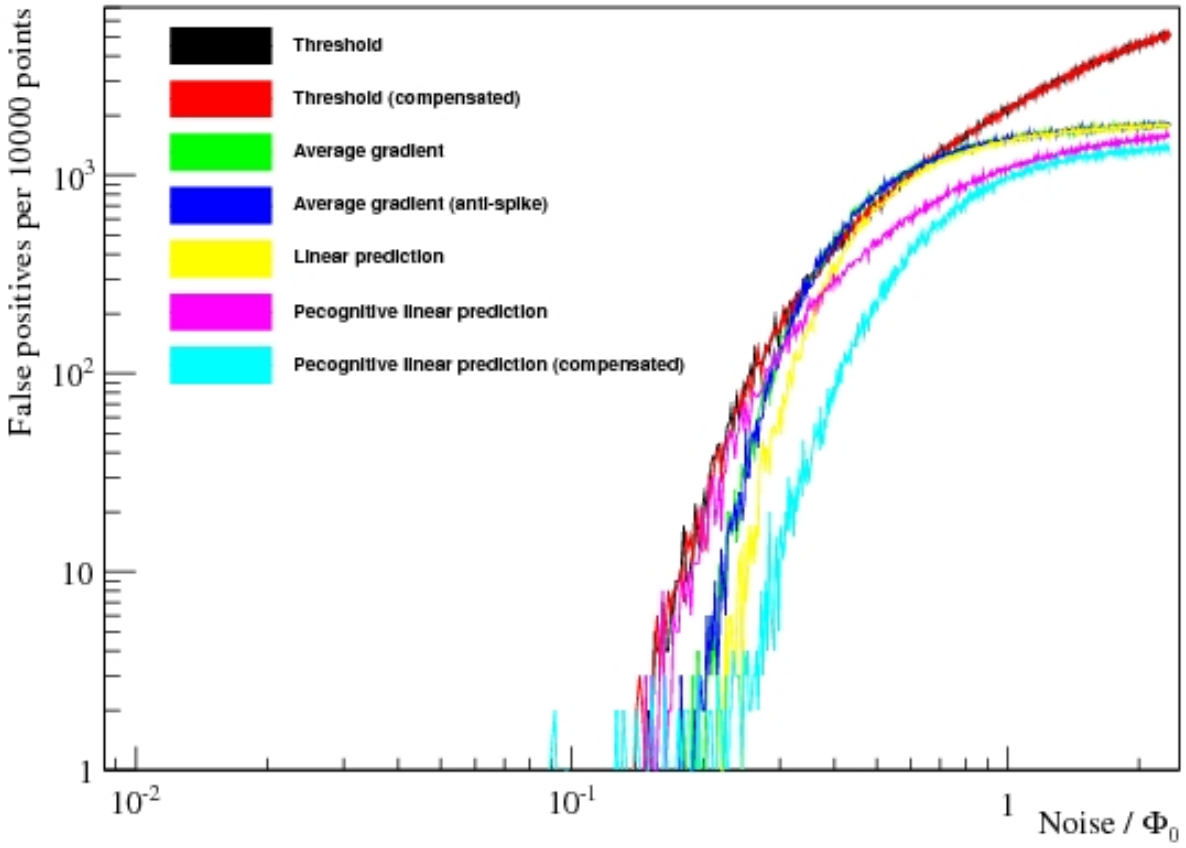

Figure 4.8: Shown is a plot of several correction algorithms' false positive rates as a function of the FWHM noise, expressed as a ratio with the size of a flux quantum. It can be seen that the false positive rate of most of the algorithms increases sharply around $0.2\Phi_0$. Below $0.1\Phi_0$ there are no false positives for most algorithms.

FWHM noise reaches $0.2\Phi_0$, after which the misidentification quickly saturates. The number of false positives is in line with the number of flux quantum size changes expected from the difference of two Gaussianly distributed data sets, this is further discussed in section 4.2.1.2.

Figure 4.9 shows the number of false negatives for each algorithm as a function of the SQUID noise expressed as a ratio to the flux quantum.

The number of false negative corrections shows similar behaviour to the number of false positives, in that there are no false negatives below $0.1\Phi_0$ and a sharp onset of false negatives at $0.2\Phi_0$. This similarity is expected, as the majority of the false negatives are actually caused by the number of false positives. This is because some algorithms limit the number of artefacts that can be identified in time, to avoid having to handle previously detected

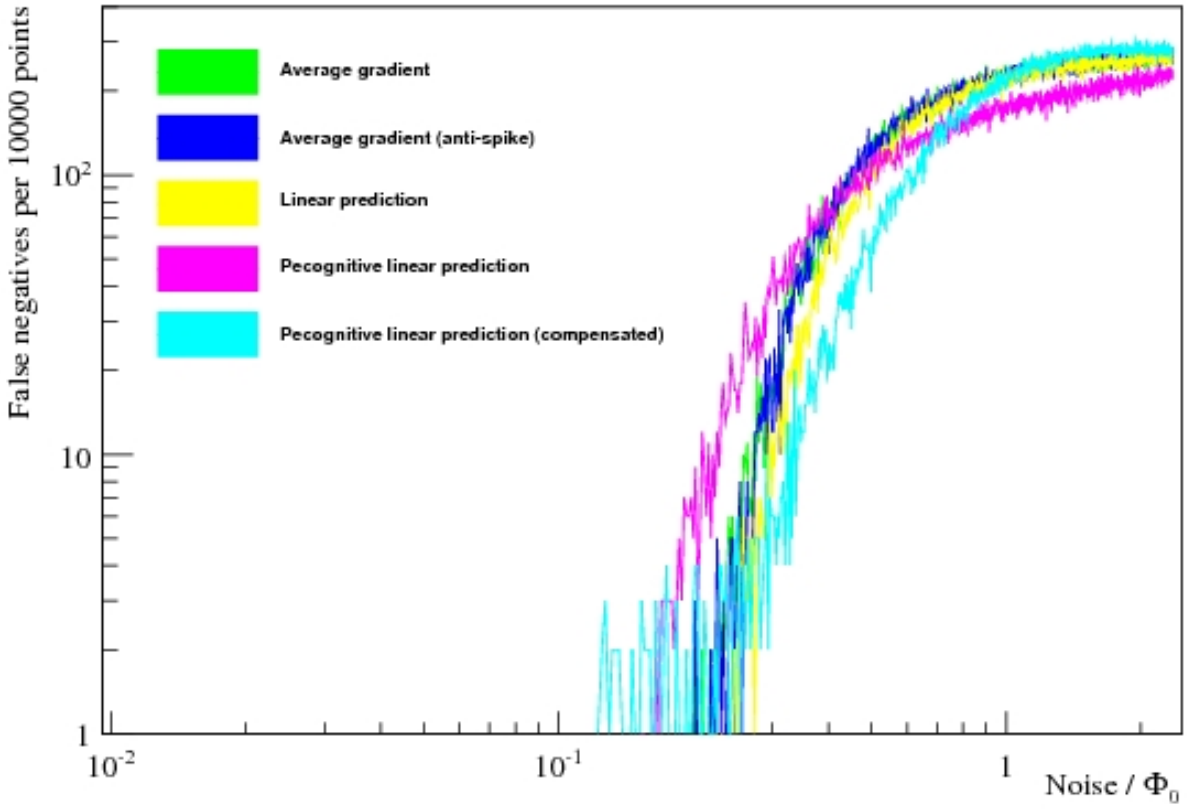

Figure 4.9: Shown is a plot of several correction algorithms' false negative rates as a function of the FWHM noise, expressed as a ratio with the size of a flux quantum. It can be seen that the false negative rate of most of the algorithms increases around $0.2\Phi_0$, below $0.1\Phi_0$ there are no false negatives for any algorithms. Neither the simple threshold nor simple threshold with compensation algorithms experienced false negatives and so are not shown.

artefacts in data buffers while checking for new artefacts. In normal operation this does not cause a problem as artefacts should not arrive in quick succession. However, if a false positive is identified slightly before a real artefact, the algorithm can ignore the real artefact leading to a false negative. For this reason the algorithms that do not require knowledge of previous data stability, the simple threshold based algorithms, exhibit no false negatives caused by previous false positives.

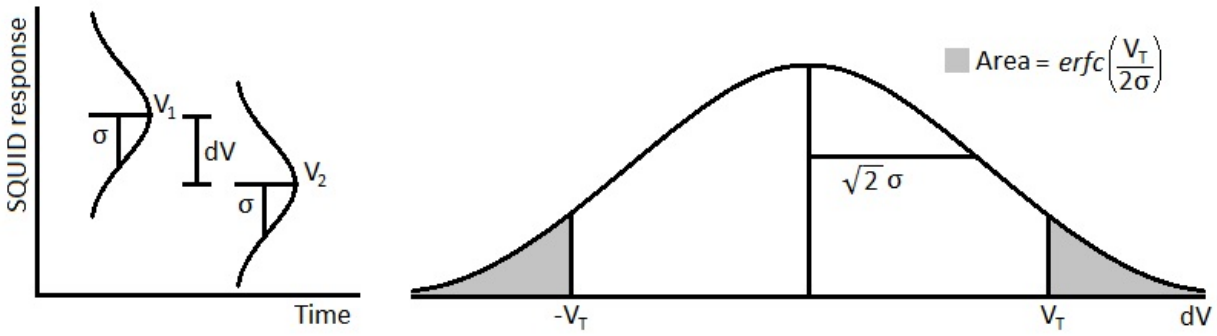

Figure 4.10: Shown is the distribution of the voltage changes between each sample, as related to the noise on the SQUID. The threshold algorithm will attempt to correct any change above the threshold value. A certain fraction of the noise distribution will leak into the threshold region, an amount erfc $\left(\frac{V_T}{2\sigma}\right)$, which will be misidentified as artefacts.

4.2.1.2 Comment on false positives

The false positive corrections exhibit a sharp loss in performance when the noise in the simulation reaches a level of about $0.2\Phi_0$. This is simply due to the expected rate at which the difference between two randomly generated Gaussian distributions will produce a value of a certain size.

Here the threshold algorithm will be taken as an example, the same principle holds for the other algorithms but is less clear due to the complications that come from the increased processing. Figure 4.10 illustrates the cause of the sudden onset of false positives. The algorithm determines the difference between two data points. In the absence of an input signal each data point will come from a distribution with the same mean and standard deviation, σ. As each point is an independent measurement of this distribution, the distribution of the differences will then have a mean of zero and a standard deviation of $\sqrt{2}\sigma$. The threshold algorithm assumes that any change in voltage above a certain level is an artefact. However, if the standard deviation of the SQUID noise is non-zero, as it will be in any physical system, there is a non-zero region of the distribution above the threshold value. Every entry in this region will be a false positive.

The region of the distribution above the threshold is given by the integral of a Gaussian

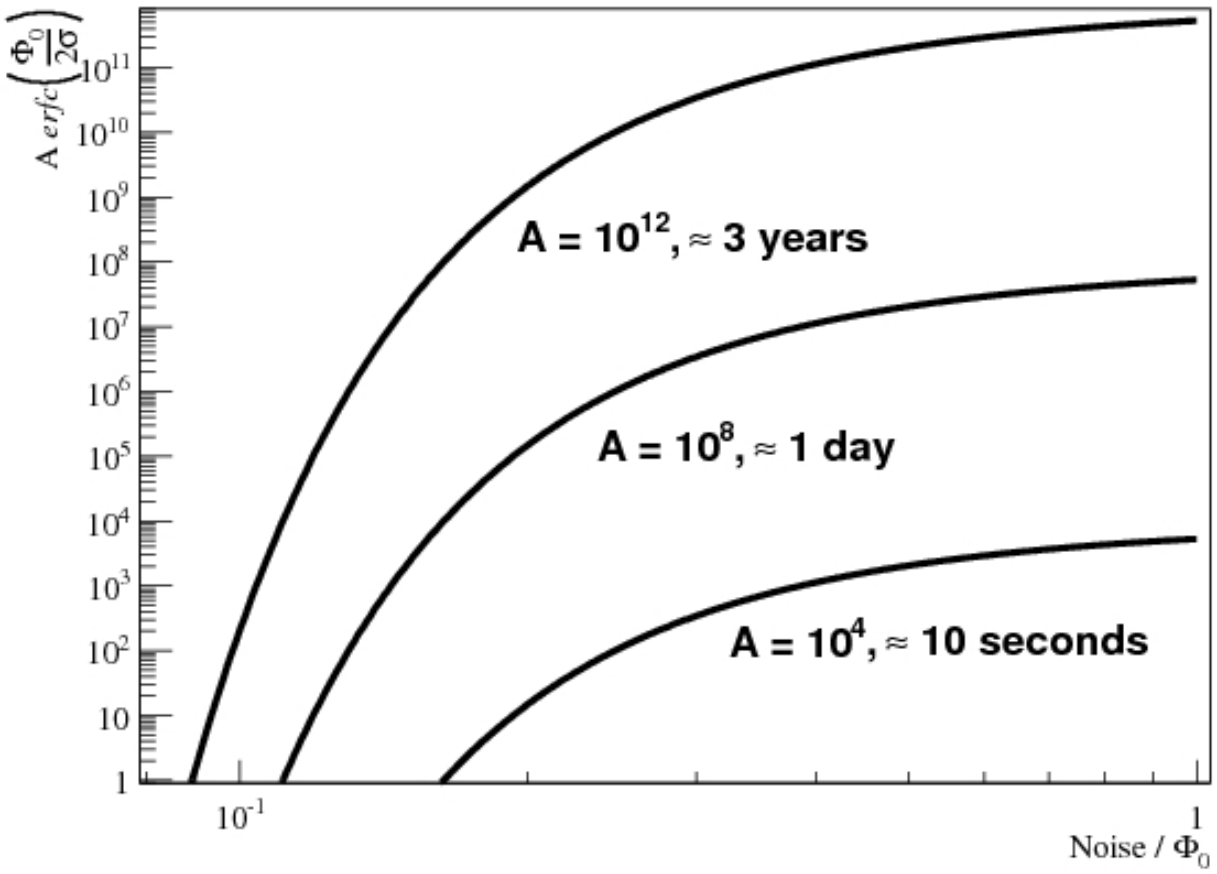

Figure 4.11: Shown is the value of $10^4 \operatorname{erfc}\left(\frac{V_T}{2\sigma}\right)$ in the noise range $10^{-1}\Phi_0$ to Φ_0. This is very similar to the values obtained for the false positives of the threshold algorithms, as expected. Also shown are the expected false positive numbers for amounts of data equivalent to 10 seconds, 1 day, and 3 years. With an eight orders of magnitude change in data there is less than a factor of two change in onset.

from $\pm V_T$ to $\pm\infty$, the complementary error function

$$N = \operatorname{erfc}\left(\frac{V_T}{2\sigma}\right), \tag{4.11}$$

where V_t is the threshold voltage.

Figure 4.11 shows the value of $10^4 \operatorname{erfc}\left(\frac{V_T}{2\sigma}\right)$, which gives a very similar curve to that of the false positives, again with an onset at approximately $0.2\Phi_0$. Also shown are the expectations in the numbers of false positives for 10 seconds, 1 day, and 3 years worth of of data. The onset slowly decreases, changing by a factor of less than two with an increase in the sample size of eight orders of magnitude.

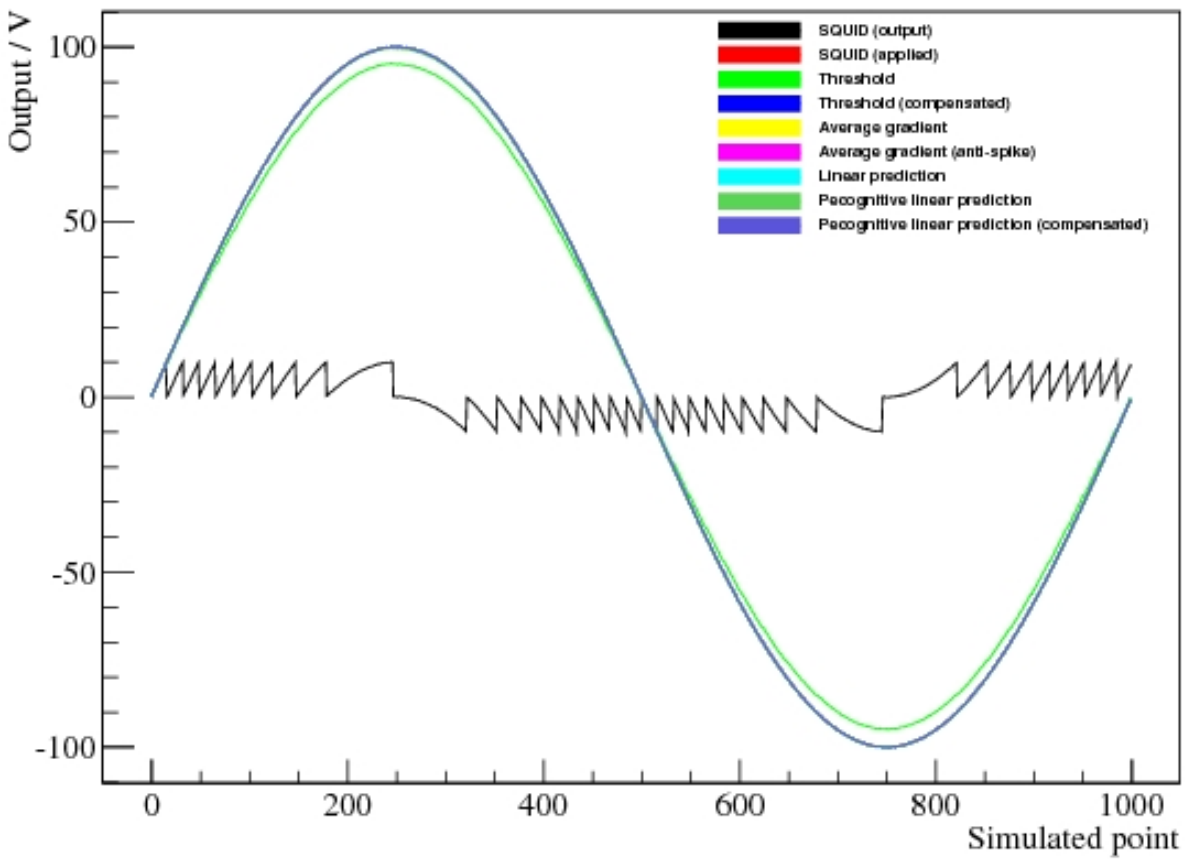

Figure 4.12: Shown are the outputs of the correction algorithms, along with the simulated SQUID output and simulated applied field for a simulation with a FWHM noise level of $0.02\Phi_0$, all algorithms except the simple threshold lie on top of the simulated data.

4.2.1.3 Accuracy

The second aspect that the correction algorithms must exhibit is the ability to recover the size of field changes in the SQUID output. The quantity used to measure the accuracy of the reconstruction will be the correlation of the algorithm's output with the applied signal. Figures 4.12 & 4.13 show a sinusoidal simulated signal, the SQUID's simulated response to the signal, and the output of the correction algorithms at two FWHM noise levels: $0.02\Phi_0$, and $0.5\Phi_0$. In the low noise simulation the output curves from all algorithms apart from the simple threshold model, which by definition underestimates the magnitude of changes, lie on top of the applied curve. In the high noise simulation, the outputs of the algorithms show a somewhat sinusoidal signal, interrupted by periods of instability caused by misidentification of artefacts and their size. These periods mean that the algorithms' outputs do not match

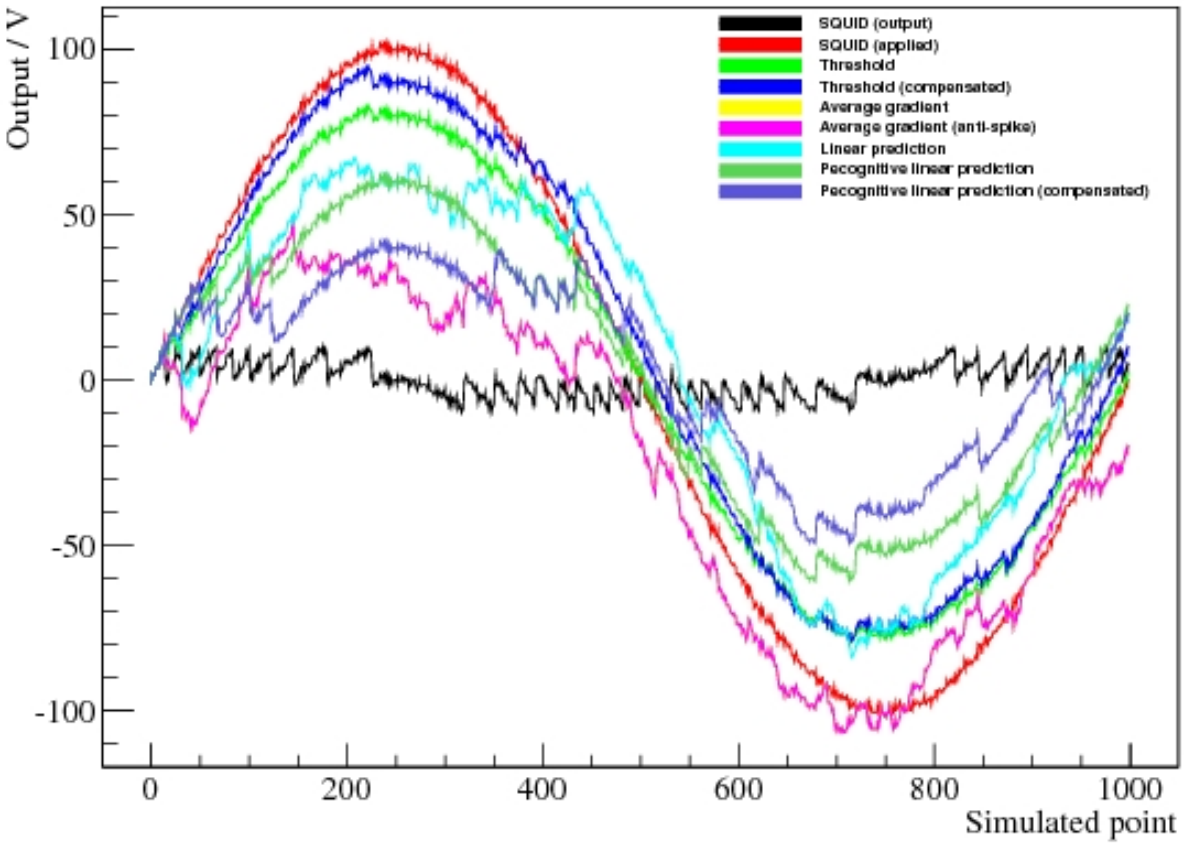

Figure 4.13: Shown are the outputs of the correction algorithms, along with the simulated SQUID output and simulated applied field for a simulation with a FWHM noise level of $0.5\Phi_0$, signals experiencing this much noise are not corrected well.

the input signal very well with this noise level.

Figure 4.14 shows the correlation, as defined in equation 3.52, between the input and reconstructed signals for each algorithm as a function of the FWHM of the SQUID's noise. The correlation is very good until the value of the FWHM noise reaches $0.2\Phi_0$, where there is a sharp decrease in the correlation.

4.2.1.4 Benchmarks

The other important aspect of these algorithms is their ability to run in real time. Typical sampling times of the DAQ system for the SQUIDs are between 1 kHz and 10 kHz, it is therefore a requirement that each algorithm is able to correct the data at a frequency much

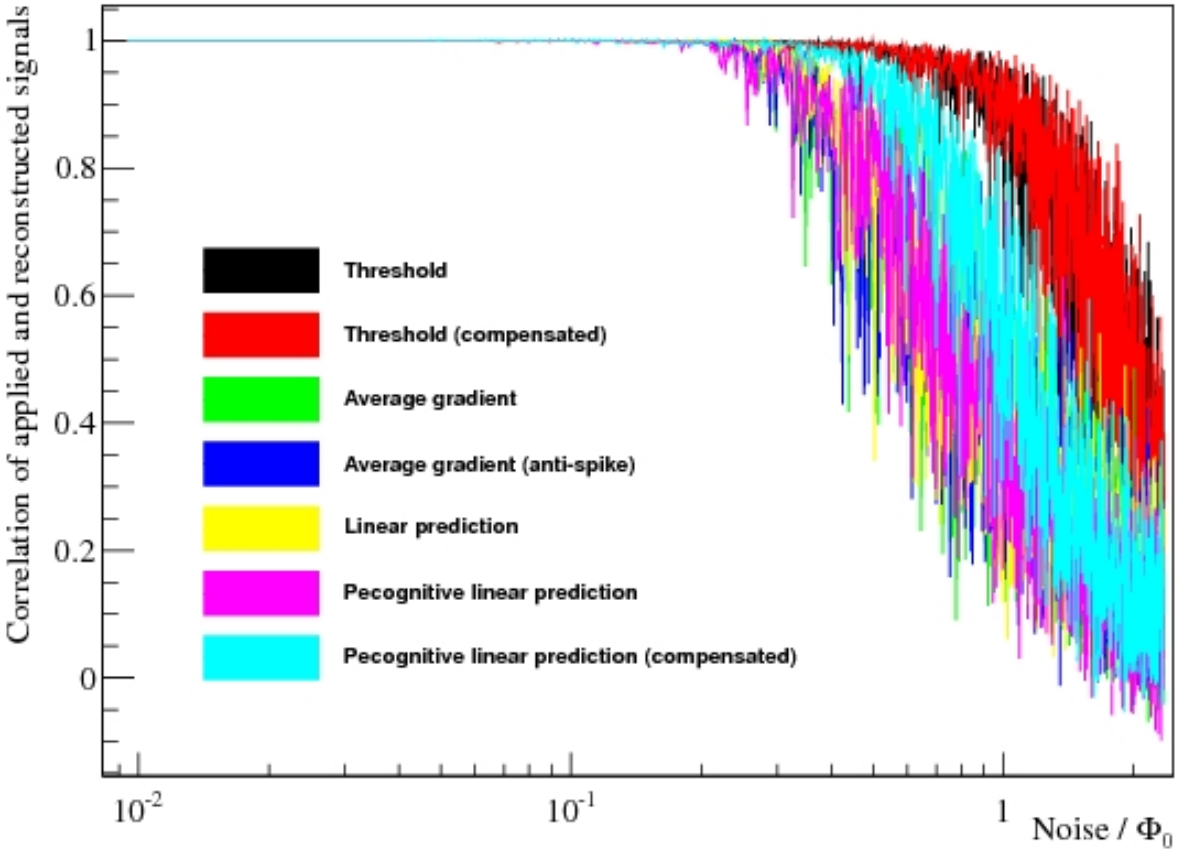

Figure 4.14: Shown is a plot of several correction algorithms' correlation rates with the simulated applied field as a function of the FWHM noise, expressed as a ratio with the size of a flux quantum. It can be seen that the correlation is very close to 1 until the FWHM of the noise is above $0.2\Phi_0$, at which point there is a sharp decrease in correlation.

greater than 10 kHz to allow time for performing the analysis, writing data files, etc. The average correction time of each of the above algorithms was measured on hardware configured identically to the DAQ machine with a clock speed of 2.5 GHz, and is tabulated in table 4.1. This also gives the maximum sampling frequency of a 12 channel SQUID system if that algorithm is used. All algorithms have this maximum sampling frequency above 50 kHz and so should be appropriate for use at 10 kHz data rates.

Algorithm	Correction time / s	Uncertainty / s	Max sampling frequency / kHz
1	1.65×10^{-7}	7.4×10^{-8}	505
2	1.71×10^{-7}	6.5×10^{-8}	487
3	8.23×10^{-7}	7.6×10^{-8}	101
4	9.87×10^{-7}	8.5×10^{-8}	84
5	8.57×10^{-7}	7.7×10^{-8}	97
6	1.38×10^{-6}	9.8×10^{-8}	60
7	1.42×10^{-6}	9.6×10^{-8}	58

Table 4.1: Table showing the average time taken to correct a single data point for each algorithm, and its associated uncertainty measured on the available hardware. The last column shows the maximum sampling rate for the 12 channel SQUID system if a particular algorithm is used. All of the algorithms comfortably allow a sampling rate of 10 kHz, the most likely largest sampling rate the system will be operated at, to be used.

4.2.2 Experimental data

During the 2010 data run the superconducting solenoid was seen to become superconducting and return to the normal state by both the fluxgates and the SQUIDs. This represents a large change in the magnetic field of approximately $4\,\mu$T, which is a substantial test of the correction algorithms. The field change is large and homogeneous enough that the fluxgates will experience essentially the same field, allowing a comparison with the corrected SQUID signals to be made. This $4\,\mu$T field change corresponds to 4240 V and 124000 V changes of the magnetometer output of two axial SQUIDs. The SQUIDs therefore experience approximately 420 and 12000 resets, which need correction, in this time. Figure 4.15 shows the raw signals of the two SQUIDs over a 30 hour period, during which the solenoid becomes normal. Figures 4.16 & 4.17 show the output of one of the correction algorithms on the SQUID data overlaid with the fluxgate data taken during the period.

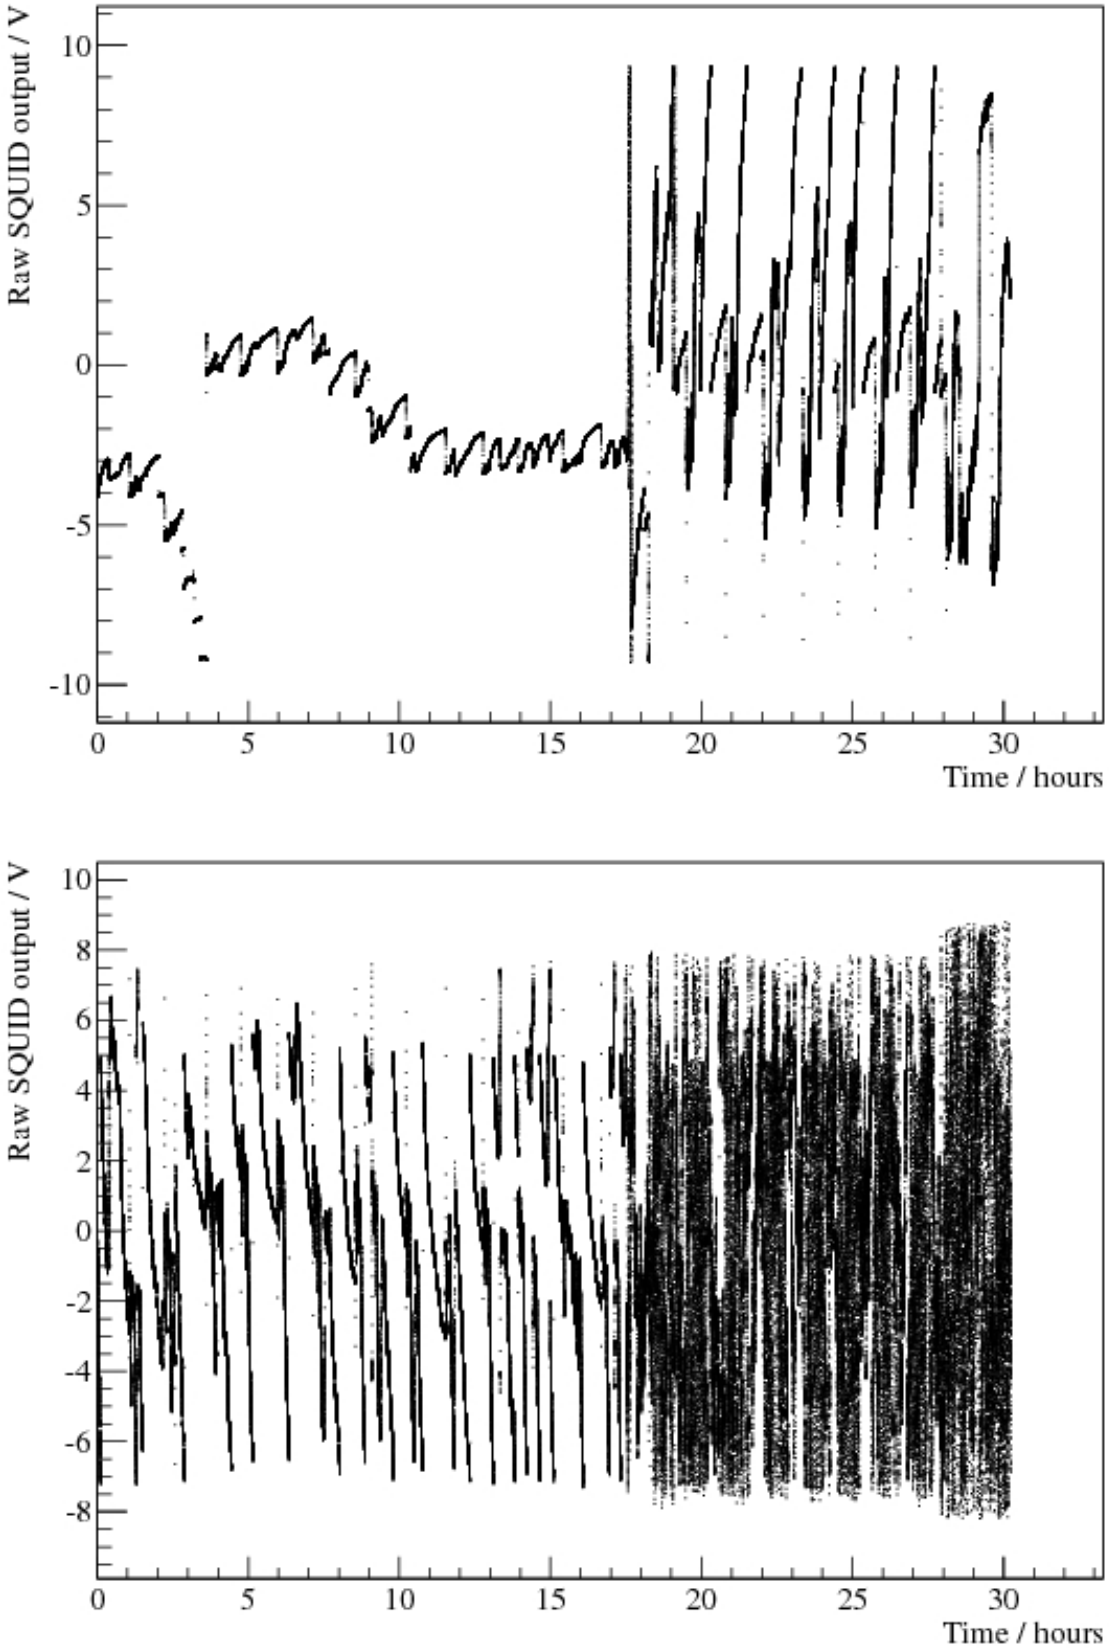

Figure 4.15: Shown is 30 hours of data around a transition of the solenoid as observed by SQUIDs 5 *(top)* & 11 *(bottom)*. A change of $4\,\mu$T to the field through the SQUIDs occurs. The number of resets during this period is 514 and 16559, for SQUID 5 & 11 respectively.

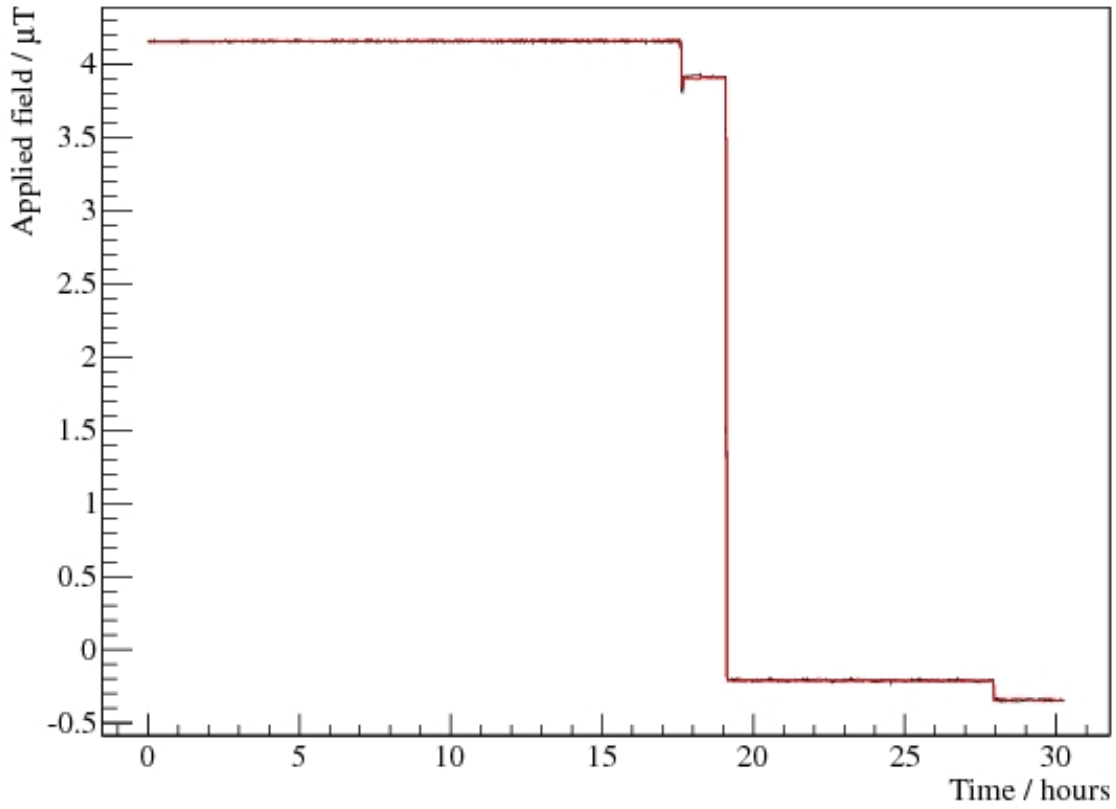

Figure 4.16: Shown is the transition of the solenoid from figure 4.15 seen on SQUID 5 with its output corrected for artefacts (red) plotted with a fluxgate observing the same period (black). The χ^2 between the reconstructed field and the fluxgate signal is 4.22. Comparison of the two signals agrees well with the calibration measured in section 3.3, giving $-9.40(2)\ \mathrm{nT\,V^{-1}}$.

The χ^2 between the reconstructed field and the fluxgate signal is 4.22 and 3.47 for the two SQUIDs. The value of the calibration obtained in section 3.3 agrees very well with that obtained by scaling the SQUID and fluxgate signals during this period: $-9.40(2)\ \mathrm{nT\,V^{-1}}$ and $0.03137(8)\ \mathrm{nT\,V^{-1}}$ for SQUID 5 and 11 respectively.

Table 4.2 gives the χ^2 for the raw data and that using correction algorithms 1 to 5 for both SQUIDs. The χ^2 of the raw data is 8 and 11 orders of magnitude higher than any of the correction algorithms for the two SQUIDs respectively, indicating how little it resembles the fluxgate data. The χ^2 for the corrected data is generally close to 4, this is likely larger than it would be if compared to the field entering the SQUID pick-up loop, as there are small non-homogeneous magnetic fields within the experimental volume discussed in section 6.4.

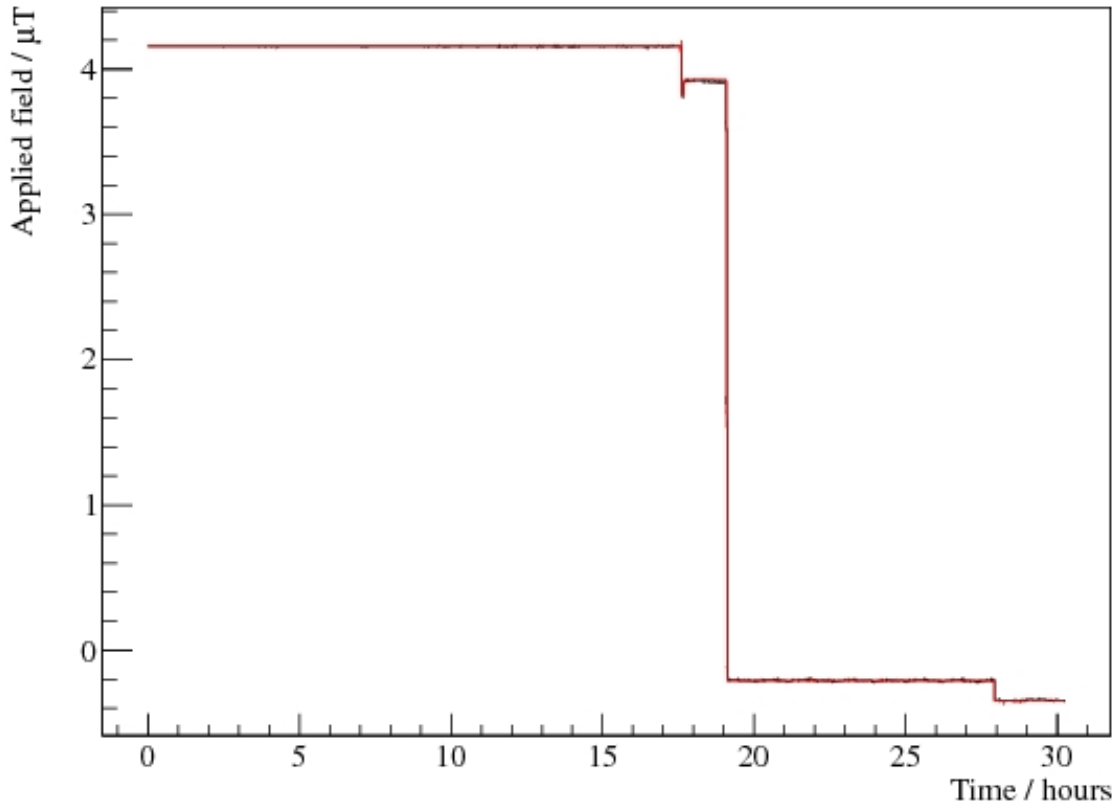

Figure 4.17: Shown is the transition of the solenoid from figure 4.15 seen on SQUID 11 with its output corrected for artefacts (red) plotted with a fluxgate observing the same period (black). The χ^2 between the reconstructed field and the fluxgate signal is 3.47. Comparison of the two signals agrees well with the calibration measured in section 3.3, giving 0.03137(8) $\mathrm{nT\,V^{-1}}$.

Algorithm	SQUID 5 reduced χ^2	SQUID 11 reduced χ^2
None	1.92×10^8	1.98×10^{12}
1	3.85	23.8
2	3.81	16.7
3	4.06	3.45
4	4.05	3.44
5	4.22	3.47

Table 4.2: Table showing the reduced χ^2 comparing the reconstructed signal of the SQUIDs to the fluxgate signal for SQUIDs 5 & 11.

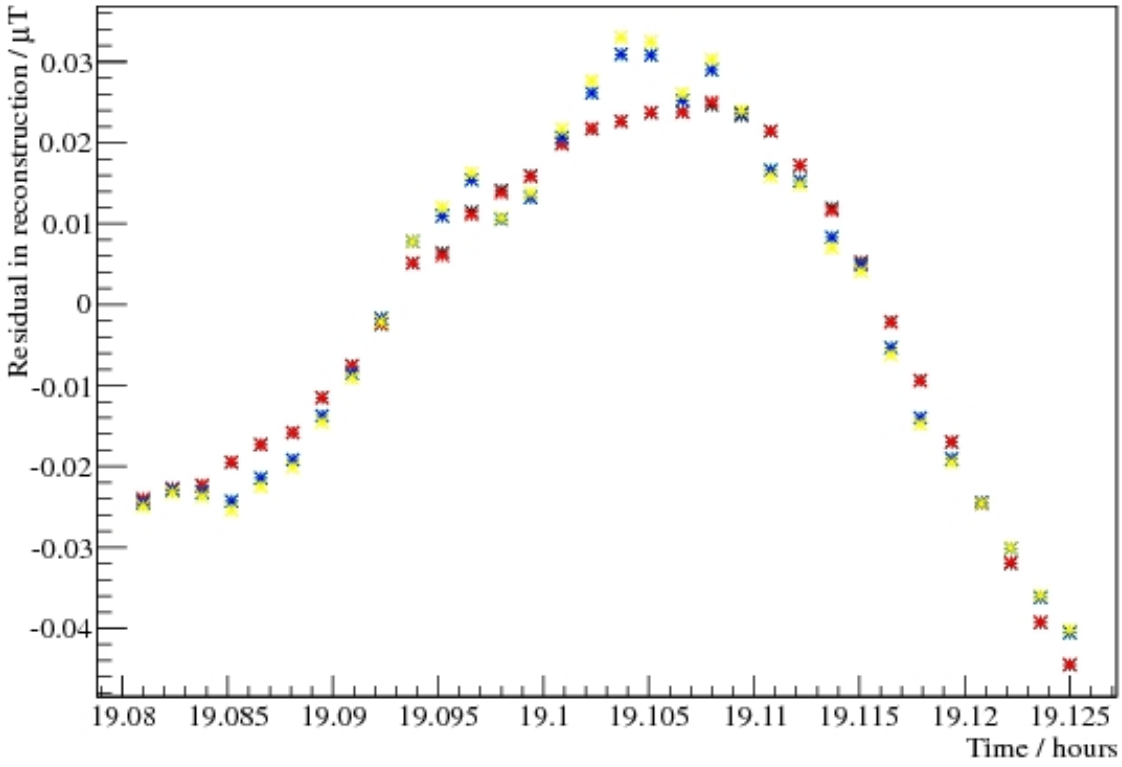

Figure 4.18: Shown is the residual of the fluxgate signal and the reconstruction in the region of the transition for algorithms 1 - 5 (black, red, green, blue, yellow) for SQUID 5. The average RMS change in the fluxgate signals over periods of this length before the transition was 15 nT, and for the SQUIDs 0.1 nT which can go some way to explaining the deviations of order 10 nT, the RMS of the residuals is 20 nT.

The threshold algorithms have a χ^2 significantly larger than the other algorithms when correcting SQUID 11's data, which highlights that the faster algorithms are unsuitable when the signal is moving very rapidly: on SQUID 11 the signal is changing at nearly 800 $\mathrm{V\,s^{-1}}$ compared to 25 $\mathrm{V\,s^{-1}}$ seen through SQUID 5.

Figures 4.18 & 4.19 show the residuals between the fluxgate and reconstructed signals for the SQUID 5 and 11 respectively during the transition region. The RMS of the residuals are all approximately 20 nT, this is comparable to the average RMS of the fluxgate signal over periods of this length, which is 15 nT. The maximum deviation from the fluxgate signal is 40 nT, which is 1 % of the total change seen. If this error is shared equally between each artefact corrected then it would mean an error of approximately 90 pT and 3 pT would be introduced with each correction for SQUID 5 & 11 respectively. Of course such large changes

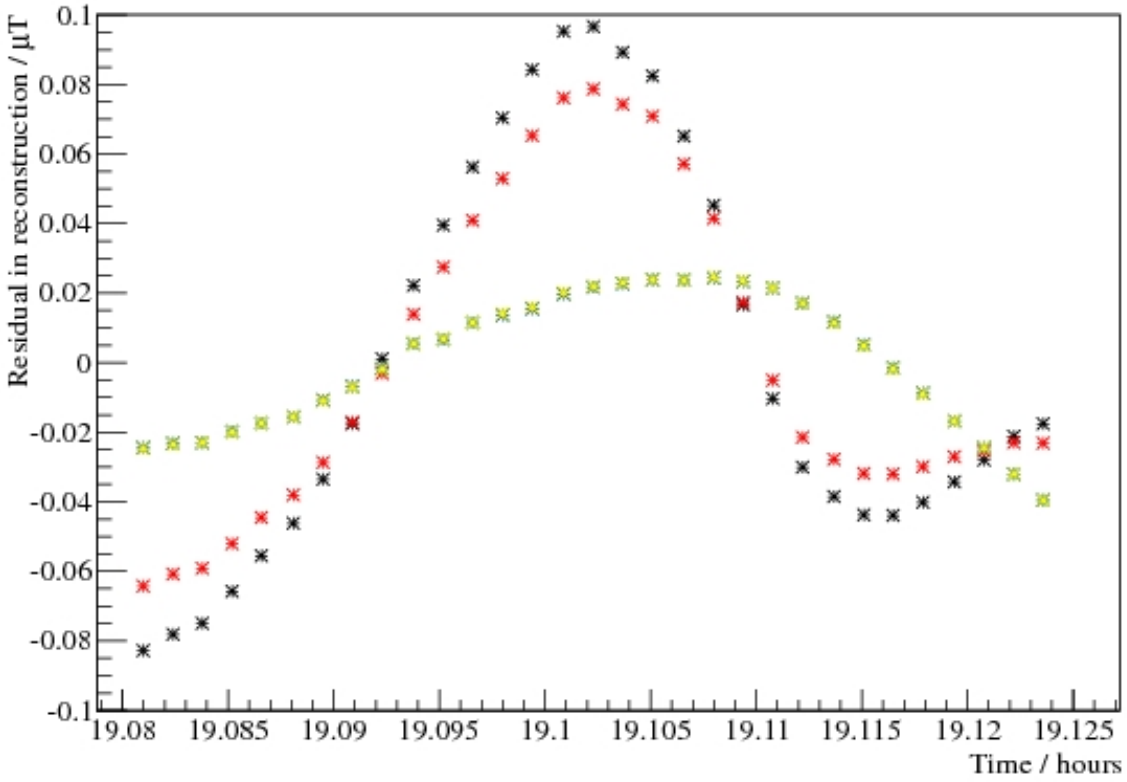

Figure 4.19: Shown is the residual of the fluxgate signal and the reconstruction in the region of the transition for algorithms 1 - 5 (black, red, green, blue, yellow) for SQUID 11. The residuals are much larger for the threshold algorithms, an RMS of 56 nT compared with 20 nT for the other algorithms.

are not to be expected within data taking, and the SQUIDs will be operating in a sensitivity range 100 to 1000 times better, which should scale the error accordingly. Additionally, it is unclear how much of this difference in signal is caused by variations of the fluxgates from the true SQUID signal as the RMS of the fluxgates and the residuals are of similar size.

4.3 Conclusion

It has been demonstrated that artefacts introduced by the hardware linearisation system of the SQUIDs can be removed, and the magnetic signal reconstructed to a very good level if the FWHM environmental noise is below $0.1\Phi_0$. The FWHM noise on the SQUIDs in the most recent data run was typically 0.09 V and 0.4 V for the SQUIDs shown above. As discontinuities caused by resets are typically around 10 V, and the noise values are signifi-

cantly less than 1 V, correcting resets can be performed easily, as demonstrated by the good agreement of the reconstruction of large field changes above. The value equivalent to Φ_0 was 5.1 V and 0.85 V for the same SQUIDs, in the former case this is again in the good quality operating window, but the latter case the noise is $0.47\Phi_0$, which is 5 times too large to reliably correct for flux jumps. Correction for resets is possible in the current noise levels within the experiment to a very good level. However, the ability to correct flux jumps varies across SQUIDs, as the noise levels on some SQUIDs is far too high. The high noise is a problem for more than just artefact correction and will be revisited in section 6.3.

The good quality of reconstruction allows the magnetometer output to be related linearly to the magnetic field experienced by the sensing region, as required by the analysis and calibration schemes. The correction can also be performed before a new sample is ready from the data acquisition system, which allows the magnetometry to be used in real time.

Chapter 5

Material investigation

Samples of the grade 5 titanium alloy obtained from various sources were measured with an AC susceptometer to determine the superconducting transition temperature of the alloy. Transitions were found to occur within the range 1.3 K to 5.6 K for the samples. A 1 K difference in transition temperature was seen between samples from the same source, one unmodified, and one containing a example weld, the welded sample produced the higher transition temperature. The aluminium alloy Dural was also measured with this technique and found to have a transition of (1.08 ± 0.03) K.

In addition to high levels of field stability, cryoEDM requires high levels of field homogeneity. This is due to the fact that magnetic field gradients can depolarise the neutrons, obscuring the resonance fringes. The requirement for a highly uniform magnetic field leads to the exclusion of all magnetic materials from the region in the vicinity of the neutrons. Of particular concern are superconductors, due to their ability to distort a magnetic field significantly through external fields induced by the Meissner effect. This is problematic as the temperature at which the experiment is run is below the transition temperature of many known superconductors.

Additionally, the materials used to build the containment volumes must meet certain physical criteria: the ability to machine to the correct dimensions; the ability to withstand certain pressures without deformation; to be sealable and leak-tight against a superfluid; and of course must come at an acceptable cost. Many materials, e.g. steel alloys and aluminium, which may traditionally be used for such purposes fail the magnetic requirements. The

design therefore requires consideration of more exotic materials.

For a number of candidate materials there are no data concerning their magnetic behaviour in cryoEDM's temperature range, and so there is a need to measure their magnetic properties before use in the experiment. The Oxford group is in a unique position to run rapid tests for superconductivity in samples of these materials, and therefore remove them from consideration if found. This work is carried out before a more detailed, and expensive, magnetic analysis is performed by specialised external labs.

A proposed SCV for the experiment was to be made of titanium. Pure titanium has a superconducting transition below 0.4 K and would be suitable from a magnetic standpoint, but is prohibitively expensive. An industrial standard alloy was suggested instead, grade 5 titanium alloy [102]. The properties of this material were investigated along with the aluminium alloy Dural [103]. This chapter outlines work done to measure any superconducting transitions of these materials.

5.1 Measurement

To measure the transition temperature of the materials the magnetic susceptibility of samples was measured as a function of temperature. In the superconducting state a material will be perfectly diamagnetic with a susceptibility of -1, which is several orders of magnitude larger than the susceptibility of a non-magnetic material, and so a transition between the two states is easily observed.

Susceptibility is defined as the ratio of the magnetisation of a material in an applied field to the applied field:

$$\boldsymbol{M} = \chi \boldsymbol{H}, \tag{5.1}$$

where χ is the susceptibility, $\boldsymbol{M}$ is the induced magnetisation, and $\boldsymbol{H}$ is the applied field

To determine the susceptibility of the sample, it is sufficient to measure the magnetic response to a known applied field. This was achieved using an AC susceptometer constructed

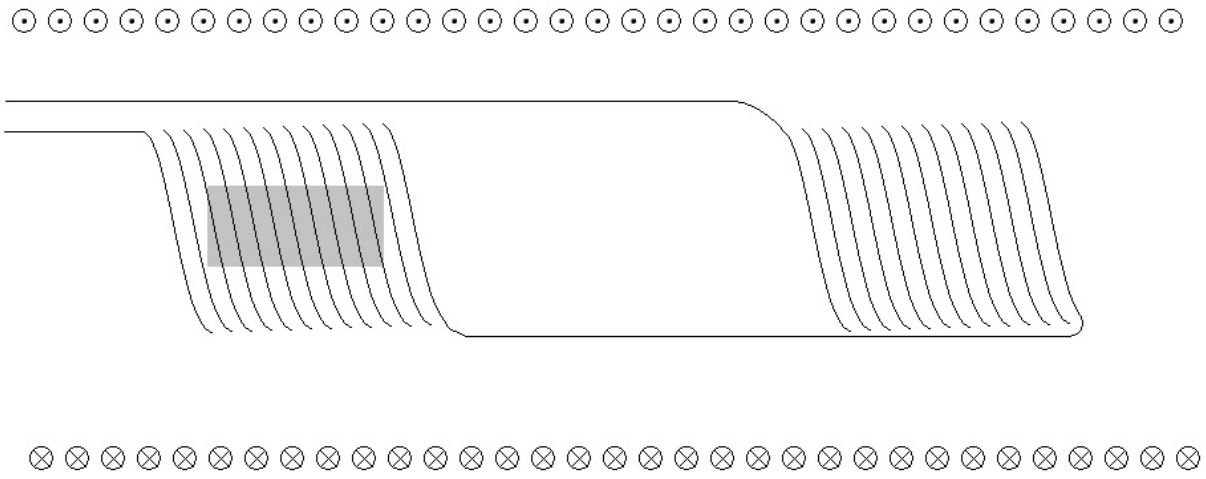

Figure 5.1: Shown is a diagram of the coil arrangement used for susceptibility measurements. The outer (primary) and inner (secondary) coils are coaxial, the secondary coil is used for read out, it is composed of two counter-wound coils one containing the sample. When no sample is present, the coupling between the primary and secondary coil produces zero output. Any change in magnetisation of the sample induces a current in the secondary coil not compensated by the counter-wound portion, and so a response is measured.

from two coaxial coils (See figure 5.1). An alternating magnetic field was applied by the outer primary coil to two spatially-separated, but connected, coaxial counter-wound secondary coils. The coils were tuned such that the induced current in one of the coils is cancelled by the induced current in the counter-wound coil, meaning there is zero total current induced in the secondary coils due to the applied field when no sample is present. When one of the secondary coils contains the sample, any current induced in the coil is caused by a change in magnetisation of the sample in response to a change in applied field. The ratio of the induced signal and the applied field is, with correct calibration, the value of the susceptibility of the sample. The applied and induced signals were read using a lock-in amplifier to produce the in-phase and out-of-phase components of the susceptibility, the former giving the magnetisation signal and the latter giving information on dissipation processes in the material. These can be seen as the real and imaginary components of a complex susceptibility. Only the modulus of this value is required for these purposes.

The applied field was created by passing a current through the primary coils. A constant current supply was used to ensure any changes in the resistance of the coils or connecting

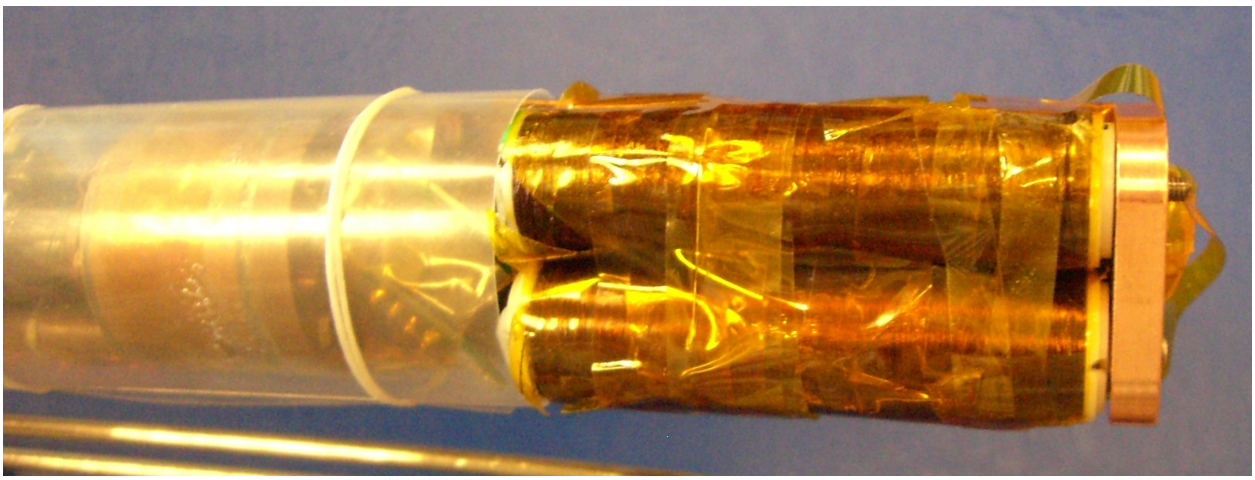

Figure 5.2: Shown is a image of two sets of primary coils side-by-side, mounted in a copper holder to the ^{3}He pot of a ^{3}He cryostat. A thermometer is embedded into the copper mount at the far right of the picture.

wires as the system was cooled did not change the applied magnetic field, which would appear as a false susceptibility signal if assumed constant. The current was then sinusoidally modulated to produce an alternating field. The applied current was recorded during measurements to compensate for any change in frequency or amplitude of the field.

The resolution of the susceptometer was estimated to be of order 1×10^{-4} for SI susceptibility, which is significantly better than the order 1 change expected for a superconducting transition. However, when compared to typical values of non-magnetic materials, which are of the order 1×10^{-6} , it is clear this susceptometer is suitable to find superconducting transitions, but not to probe the magnetic properties of the material in detail. As only transitions in the susceptibility are required to determine a superconducting state and, due to the poor resolution, no useful information could be gathered from the absolute values of the susceptibility an accurate calibration of the susceptometer was not required.

Figure 5.2 shows two AC susceptometers mounted on the experimental plate of a ^{3}He cryostat. The samples were screwed into the copper mount to ensure a good thermal contact and the coils mounted over them. A thermometer is embedded into the copper mount to ensure it measures the sample temperature as close to the samples as possible, without interfering with the measurement.

5.2 Temperature calibration

To accurately determine the temperature of the sample, a Cernox thermometer was inserted into the sample holder. The resistance of the thermometer varies repeatably with temperature, and so measurement could be performed via a four point resistance measurement using a resistance bridge. The only obstacle was to relate the resistance of the thermometer to its temperature.

The resistance calibration was performed by using measurements of the transition temperature of samples of known superconductors, and fitting a curve with the functional form [104]

$$\ln R = a(\ln T)^2 + b \ln T + c, \tag{5.2}$$

where R is the thermometer resistance, T is the temperature of the device, and a, b and c are constants to be determined. This form is valid for this type of thermometer at low temperatures, below approximately 10 K.

Six reference samples were measured with transitions between 0.39 K and 7.2 K. These were combined with measurements of the resistance of the thermometer at the boiling point of liquid helium, the boiling point of liquid nitrogen, and room temperature. The curve defined by equation 5.2 was fit to the reference values in the region of interest, but this curve is only valid for low temperatures. From 10 K to room temperature, which is outside the region of interest, the calibration is interpolated between the reference values using an Akima spline [105]. Figure 5.3 shows the thermometer calibration curve obtained in the region of interest in good agreement with the measured transitions.

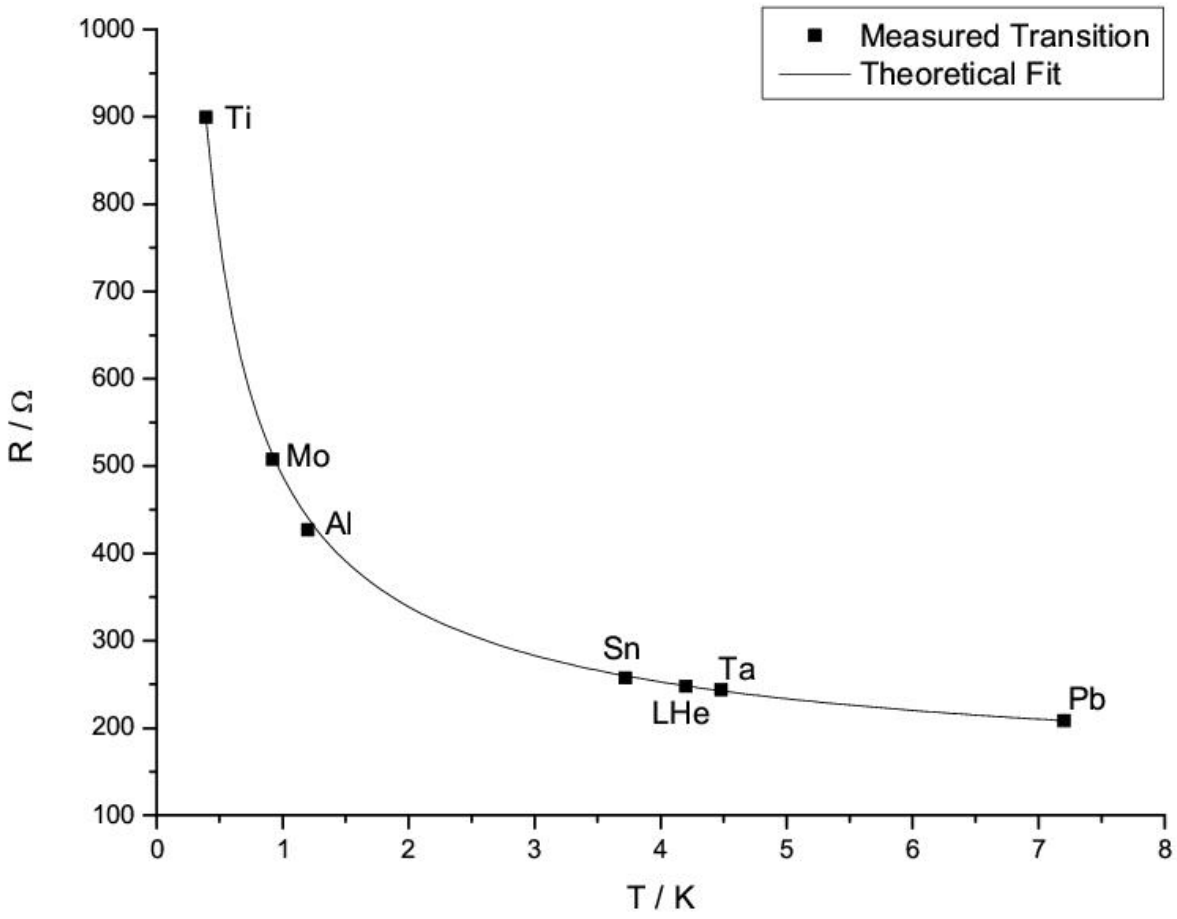

Figure 5.3: Calibration curve of the Cernox thermometer showing the fit of equation 5.2 to the known transition temperatures at measured resistances of reference samples.

5.3 Determination of T_c

The standard definition of the transition temperature is the point where the susceptibility increases from zero [106]. However, all signals are subject to noise, and these measurements are no exception, so deciding whether a change in the data was the onset of superconducting behaviour, or an artefact in the noise was open to interpretation. To remove any subjectivity in determining the departure from zero, a method using computable quantities was developed, as illustrated in figure 5.4. A region far above the transition was determined, this region was averaged to obtain a value and standard deviation of the susceptibility signal in the normal state. From this temperature down a moving average was taken, when this average was 1σ away from the value for the normal-state region, this temperature was taken

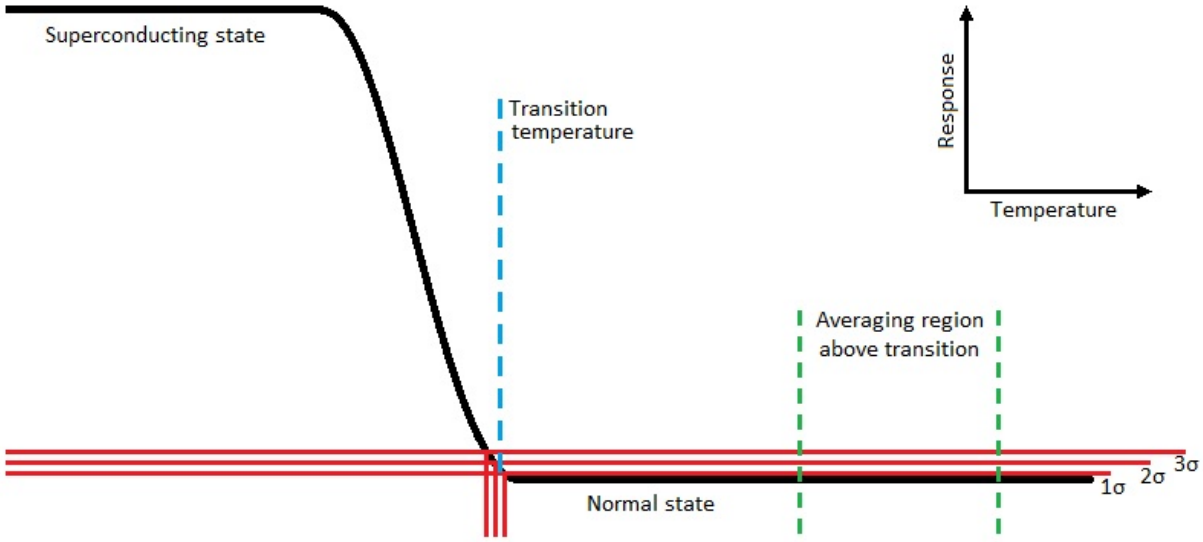

Figure 5.4: Shown is an illustration of a transition curve, and the definition of transition temperature used. A region well above the transition is used to define an average and standard deviation in the value of the normal state response. The region of the curve that lies greater than 1σ, and less 2σ away from the normal state is found. The transition temperature is defined as the centre of this region. Clearly due to the gradient of the transition the definition of the lower temperature edge does not change the result significantly.

to be the upper bound of the transition region. To determine the lower bound of the region a similar procedure was carried out, but starting well below the transition and moving towards higher temperatures. Once the moving average was within 2σ of the normal-state region this was taken as the lower bound. The transition temperature was taken to be the centre of the transition region, so the average of the warmest point where the signal was more then 1σ from normal, and the coldest point where the signal was within 2σ of normal. An uncertainty in the transition temperature is attributed from the uncertainties in finding the correct transition region. Changing the length of the moving average by a factor of two in each direction, or changing the lower bound threshold from 2σ to 3σ did not change the results significantly, from this is was concluded that this method is stable with respect to the arbitrarily chosen parameters.

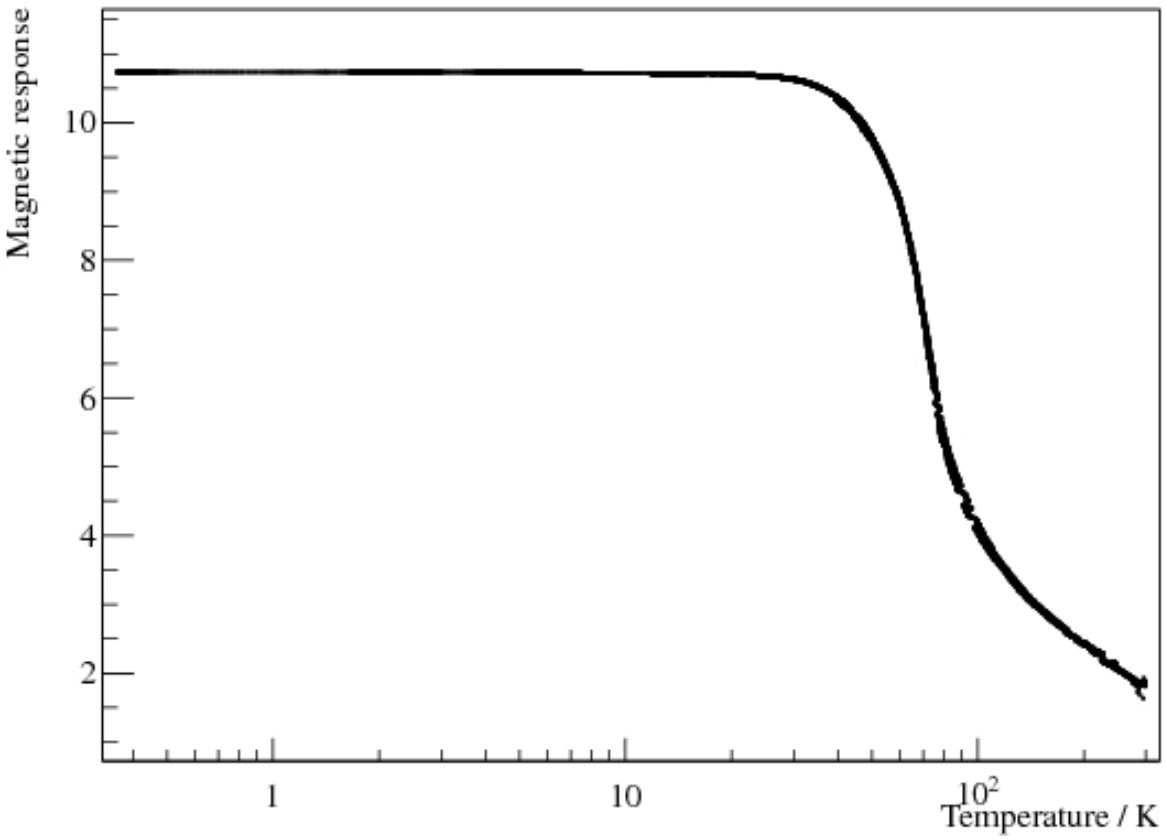

Figure 5.5: Susceptibility curve of a copper sample measured for testing purposes. Visible is a change in magnetic signal from room temperature to approximately 40 K due to changes in the conductivity of the metal with temperature, a feature shared with all samples.

5.4 Results

As a test of the apparatus a copper sample was used. Figure 5.5 shows the curve produced for the copper sample, there is a change in susceptibility between 40 K and 300 K, this is a well known phenomenon caused by the change of conductivity of the sample with temperature creating eddy currents and described in [106]. This change was visible in all samples but occurred significantly above the region of interest, it therefore did not interfere with determination of the transition temperature of samples.

A number of samples of grade 5 titanium were tested from various sources:

1. A sample provided through the Rutherford Appleton Lab (RAL) engineering department.

2. Another sample provided through RAL from the same batch as 1, but with a sample weld. As the vessel would require welding the nature of the welds would need to be determined also.

3. A sample of the titanium welding rod.

4. A sample from the manufacturer Goodfellow.

5. A sample from the manufacturer Dynamet.

6. A sample obtained through contacts in Oxford.

7. A sample reported to have no transitions in the region of interest [107].

In addition to the titanium samples, a sample of the aluminium alloy Dural was tested.

All samples were measured while both cooling and warming through the transitions, due to the design of the cryostat it was not possible to stabilise the temperature in the vicinity of the transitions and map them precisely. Some samples showed a small but noticeable hysteresis in the transition temperature, likely due to the intrinsic thermal properties of the sample itself. The hysteresis of such samples was typically about 2 % of the transition temperature, an example of the hysteresis is shown in figure 5.6, which shows the transition curve of the welded grade 5 titanium sample and the two transition temperatures attributable. Due to the hysteresis, all curves and transition temperatures will be quoted for the cooling of the sample, this ensures that all data are comparable and, as the cooling transition temperature is lower, that the superconducting behaviour in the material is below the quoted transition temperature.

All samples tested showed a superconducting transition in this experiment. Table 5.1 shows the values of the transition temperature and its associated uncertainty obtained for each of the samples listed above.

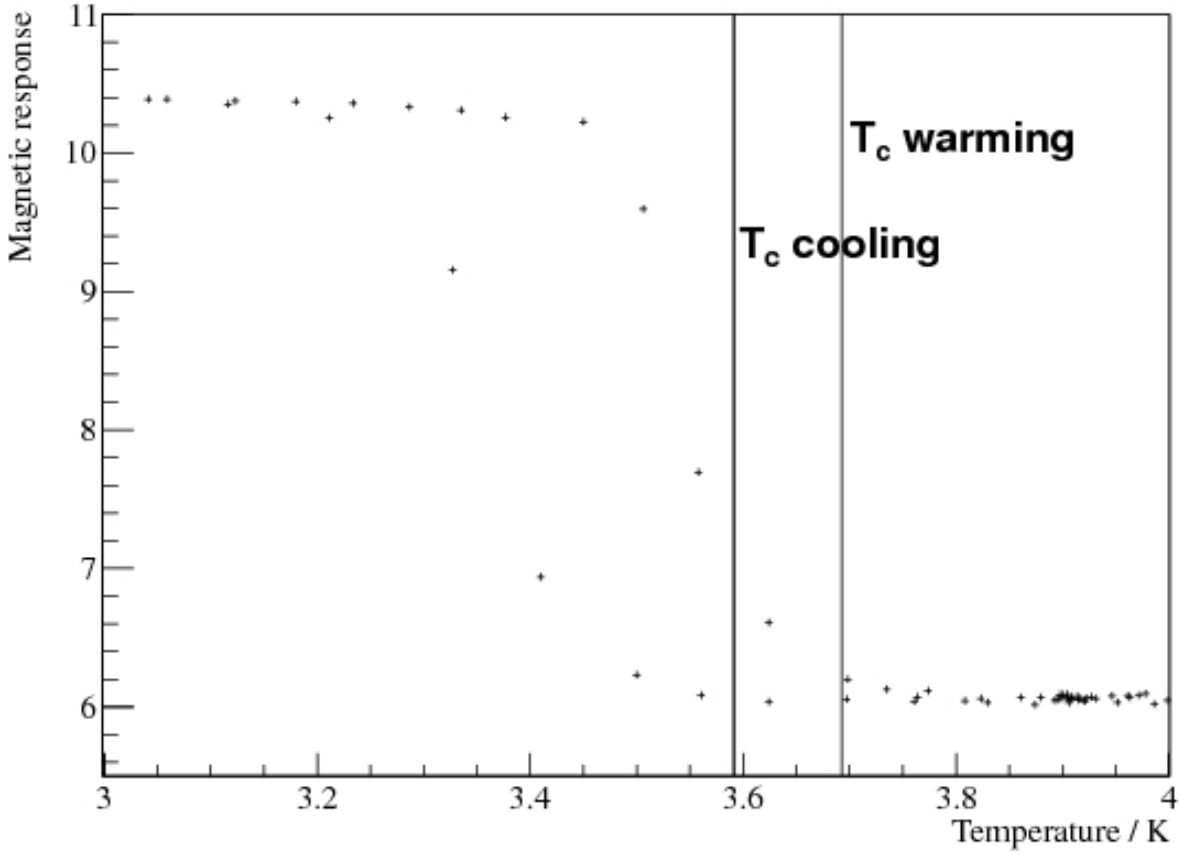

Figure 5.6: Shown is an example of the hysteresis exhibited by some samples. The welded grade 5 titanium sample is shown here with a transition temperature of $(3.59 \pm 0.04)\,\mathrm{K}$ and $(3.69 \pm 0.04)\,\mathrm{K}$ for the cooling and warming phase respectively.

Sample	Label	Transition / K	Uncertainty / K
Dural		1.08	0.03
Ti G5 RAL (welded)	T1	4.42	0.03
Ti G5 RAL (unwelded)	T2	3.59	0.04
Ti welding rod	T3	1.38	0.03
Ti G5 Oxford FD	T4	5.57	0.02
Ti G5 Goodfellow	T5	1.34	0.03
Ti G5 Dynamet	T6	1.43	0.03
Ti UD	T7	3.21	0.02

Table 5.1: Table showing all measured samples' transition temperatures, with associated uncertainty. Labels on the titanium samples refer to the labels on the curves displayed in figure 5.7.

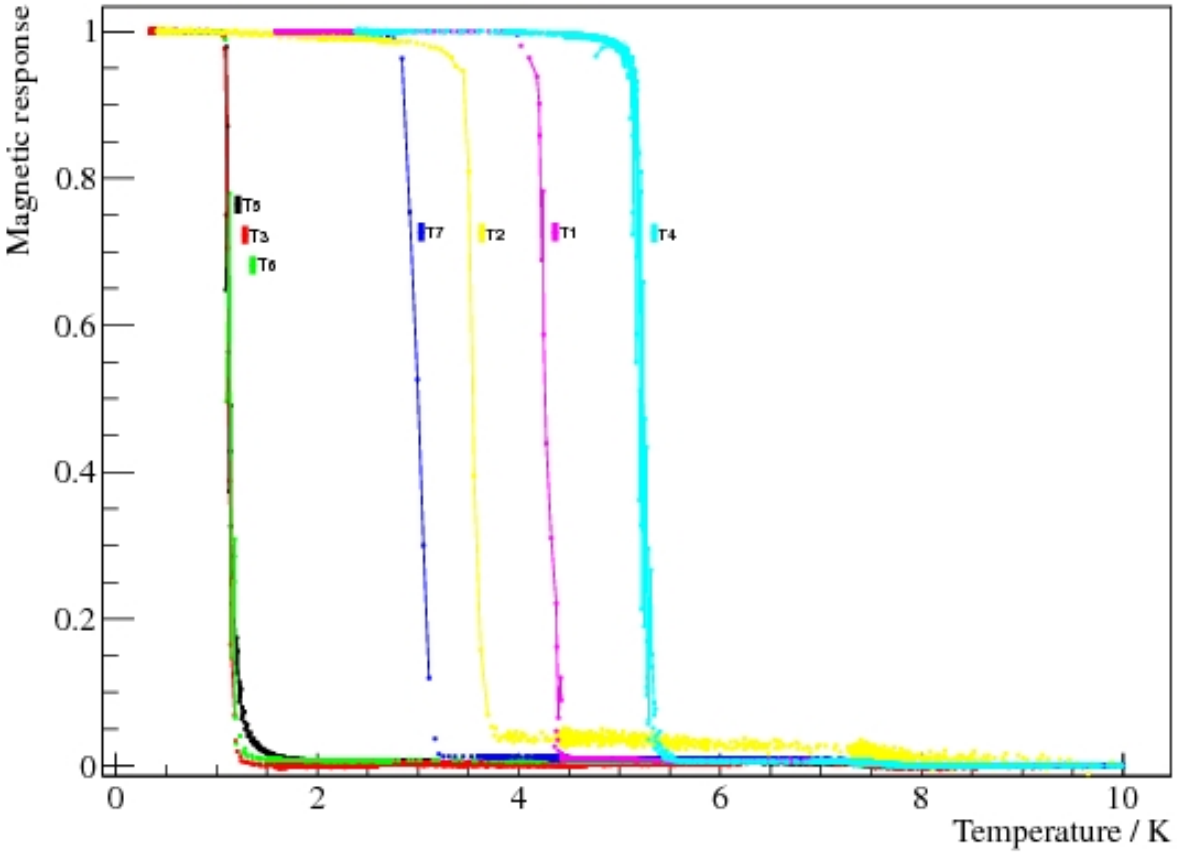

Figure 5.7: Measured susceptibility curves for the titanium samples, offset and normalised. Each sample shows an transition in the range 1.3 K to 5.6 K.

The associated curves for the transitions of the titanium samples may be seen in figure 5.7, here the curves have been normalised and vertically offset for ease of comparison. The grade 5 titanium samples each have different and incompatible transition temperatures, occurring in the range 1.3 K to 5.6 K. The transition curve for Dural can be seen in figure 5.8, and occurred with a transition temperature of (1.08 ± 0.03) K.

5.5 Conclusion

There is a significant spread in the values of the transition temperatures of the chemically standard grade 5 titanium samples, a range between 1.3 K and 5.6 K. Of particular note is the approximately 1 K increase in transition temperature of the RAL sample after welding. As there is no visible transition at the same temperature as the not-welded sample then

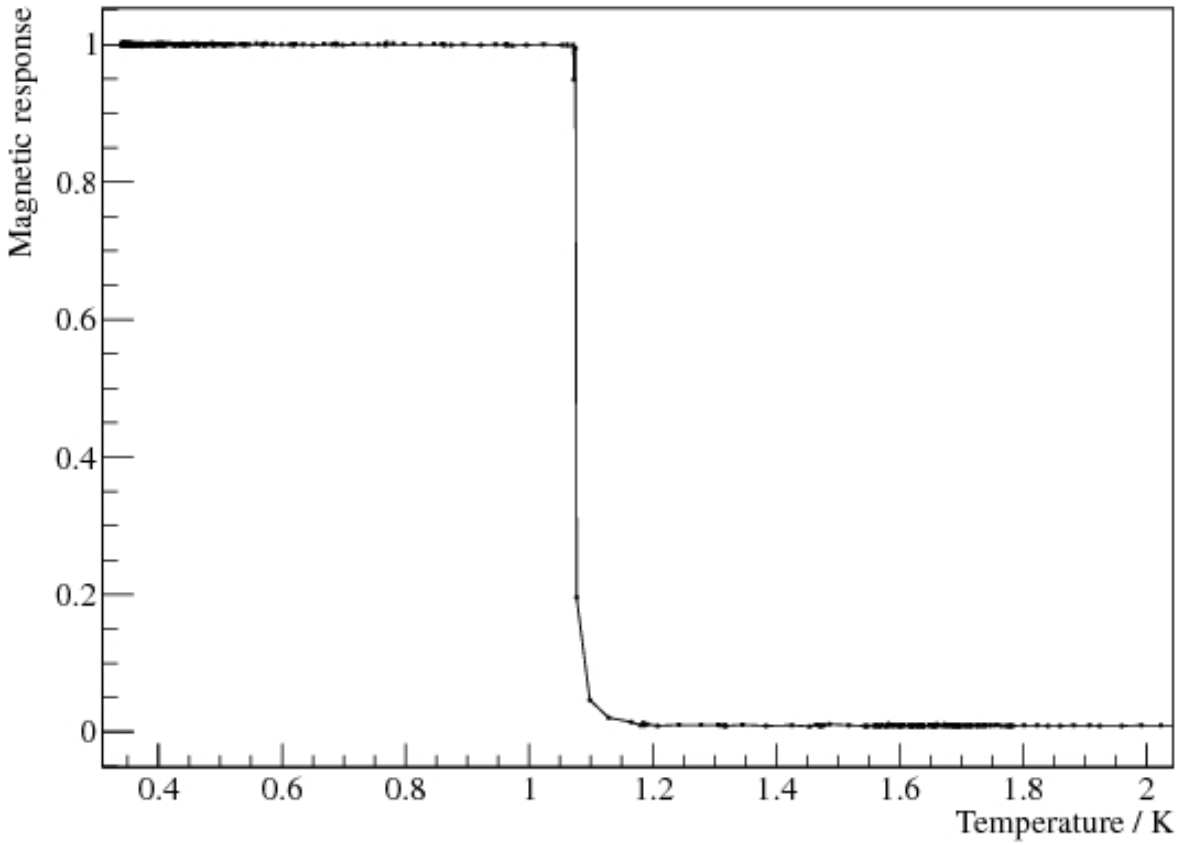

Figure 5.8: Susceptibility curve of the dural sample with a transition of (1.08 ± 0.03) K.

the welding must have changed the bulk properties of the material, rather than simply introducing contaminants with a higher transition temperature. These results highlight the importance of obtaining test samples from the same batch of material as that intended to be used. Especially as this range of temperatures straddle the often used cryogenic temperature of liquid helium. Indeed it is easy to imagine a situation where a material, such as the RAL sample, would be used in an application assuming its normal state at liquid helium temperature, but after welding into position would be superconducting at liquid helium temperature.

The above results clearly indicate that grade 5 titanium from any of the manufacturers tested would not make a suitable material for construction of the superfluid containment vessel in cryoEDM, which requires materials that remain in a normal state down to a temperature of 0.5 K. Additionally with a transition of (1.08 ± 0.03) K, Dural would be unsuitable

for this purpose as well.

Also worthy of note is the sample claimed to have no transition above 0.2 K is clearly seen to have a transition at a temperature of (3.21 ± 0.02) K. This corrects the previous result.

Chapter 6

The magnetometry system and the magnetic environment

A significant amount of data were taken by the cryoEDM magnetometry system in October-December of 2010 during a prototype data run. This was the first successful run of the magnetometry system, as it provided functional SQUIDs that could be calibrated, and which were generally stable and able to measure the magnetic field in the apparatus. A summary of the performance of the magnetometers, and of the magnetic environment in the experiment will be presented.

6.1 Software

The magnetometry system is used both to remove false EDM signals from the final EDM analysis, and as a component in a system keeping the magnetic field at the experimental operating point. While the former use can be done offline and only requires data files to be written while running, the latter requires a full real-time analysis to be performed on the data as it is taken. Additionally, the magnetometry data must be made available to the other components in the field stabilisation system. This is done through communications over a network, with a master DAQ server issuing commands to the magnetometry system,

and requesting data from it.

In order to have a working real-time analysis, a significant software framework had to be developed, this included data taking, data recording, signal processing, configuration, user display, etc. The full software system is described in appendices D and E. Comments on its use and performance in real data taking situations are given here.

6.1.1 Monitoring software

The user accessible side of the data analysis software is a graphical environment in which to monitor the magnetic environment remotely, which allows a user to view and manipulate the raw and corrected/calibrated data in real time. Data sets may be stored for future reference, or combined to produce new data sets via a scripting system (detailed in section D.5.1). An example of the monitoring software in use is shown in figure 6.1. If provided with an administrative password the client program can also mimic commands sent by the rest of the cryoEDM DAQ system to the magnetometry for testing purposes.

It is useful to have the ability to observe certain parameters of the magnetic environment in real time for data quality checking. Many of these parameters are not directly measured, but are simple manipulations of the measurements, for example the stability of the dynamic shielding factors. These parameters are produced by the scripting system.

The script system was used to great effect during the susceptibility measurements described in chapter 5. The script system was used to plot the calibrated superconducting transition curves of the samples as the data were being taken, and to monitor various parameters. This allowed a rapid check of the data quality, meaning experimental issues could be spotted and corrected before the sample was removed. The system will play a similar role in the main experiment.

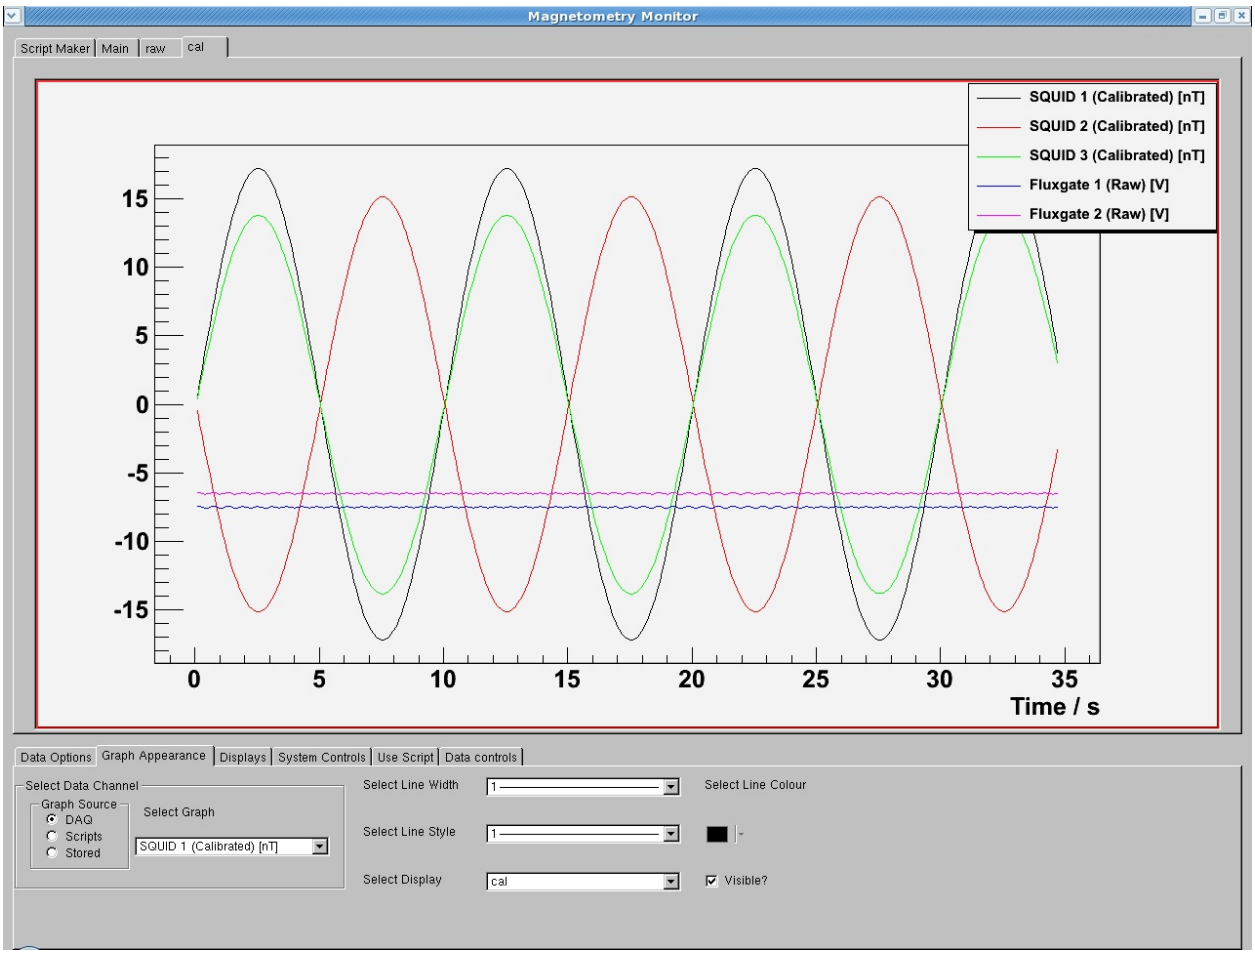

Figure 6.1: Shown is a screen shot of the magnetometry monitoring software displaying the corrected and calibrated output of three example SQUIDs, and the raw output of two example fluxgates at the same time. The display properties tab at the bottom of the figure is open to allow display changes. Other tabs under the drawing region allow access to other functionality, e.g. changing incoming data rates or sending commands to the server.

6.1.2 Performance

The software detailed above was used extensively during the 2010 data run. It was tasked with recording data from the SQUIDs through a newly designed data acquisition system, and from the fluxgates through an off the shelf ADC card. A further task was controlling and recording DAC values used to supply currents to compensation coils. During this period it recorded approximately 65 days worth of data with minimal user intervention, and successfully processed and responded to instructions provided from other parts of the experiment. Data taken using this software can be used to deduce the performance of the magnetometry and to assess the magnetic environment of the experiment.

6.2 Magnetometry progress

Much work has been done by the Oxford group to achieve a functioning magnetometry system in the environment of cryoEDM. Original designs of pick-up loops were to be circular loops of approximately 20 cm diameter, a sensing area of 3×10^{-2} m^2. However, this size of pick-up loop collected significant amounts of RF interference. The transmission of RF interference on to a SQUID chip can diminish the SQUID's performance, or prevent the SQUID's operation entirely [108]. With the levels present in the experiment, the SQUIDs could not be operated with this size of loop attached.

The loops were redesigned as rings, still with a 20 cm diameter, but with an area of approximately 3×10^{-3} m^2. The RF interference picked up was reduced, and the SQUIDs could be operated intermittently, but not reliably. The sensing area of the pick-up loops was reduced further, to approximately 9×10^{-4} m^2 using small, off axis circular loops. This allows the SQUIDs to be operated in most situations, but introduced a new problem. With such a small sensing area, contributions to this area from the twisted wire pairs connecting the loops to the SQUIDs became significant. Ageing shielding used to screen these wires from external sources of magnetic field had breaks, and allowed significant pick-up of field in the regions not covered by the experiments magnetic shielding. The shielding of these wires had to be remade, and this has now eliminated the extraneous pick-up problem.

Investigations into the optimal pick-up loop designs for this environment are ongoing by the group. The pick-up loops attached to the magnetometers used in this data run have the maximum sensing areas that can tolerate the RF environment and remain functional. SQUID 5 has a diameter of 16 mm and a sensing area of 2×10^{-4} m^2, SQUIDs 1, 4, & 11 each have a diameter of 35mm and a sensing area of 9.6×10^{-4} m^2. 6 other pick-up loops were installed with sensing areas ranging from 5×10^{-4} m^2 to 3×10^{-2} m^2, none of which could be brought into consistent operation, presumed to be due to pick-up of RF interference in the experiment. The performance of the functional pick-up loops is investigated below.

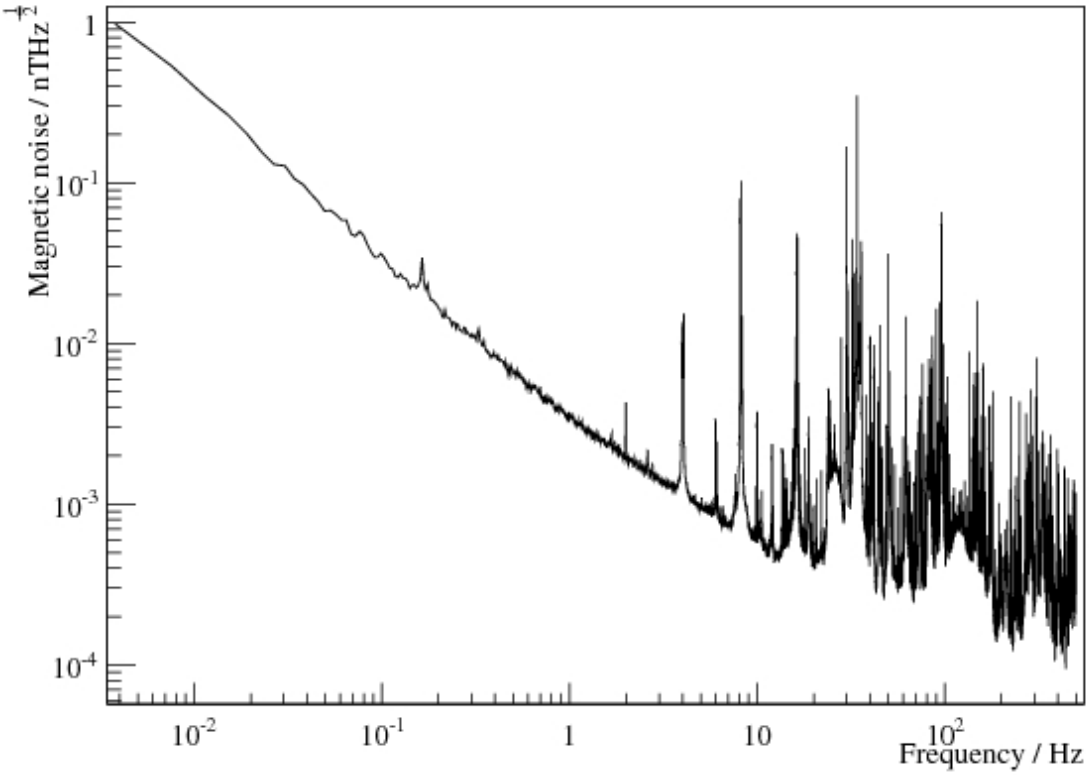

Figure 6.2: Shown is the noise spectra of the axial SQUID 5 in the frequency range important to a Ramsey resonance measurement 4×10^{-3} Hz to 500 Hz. SQUID 5 has a resolution in this range of 95 pT. The noise power is dominated by a $\frac{1}{f^2}$ contribution, meaning that it is the environment limiting the resolution and not the true resolution of the SQUIDs. At higher frequencies large peaks in the noise can be seen, the source of this noise is not yet determined, but will be investigated further in section 6.5.2.2.

6.3 Magnetic noise

Oscillating magnetic fields with frequencies corresponding to times longer than the measurement cycle, or frequencies much greater than the neutron precession frequency do not significantly affect the measurement and can be ignored initially. The important range is then between approximately 4×10^{-3} Hz and 500 Hz. A measurement of the magnetometer noise in this region is essential.

The noise spectra, again defined as the square root of the power spectra, for two axial SQUIDs are shown in figures 6.2 & 6.3. The data were taken during the normal operating mode of the experiment with all magnetic shielding available.

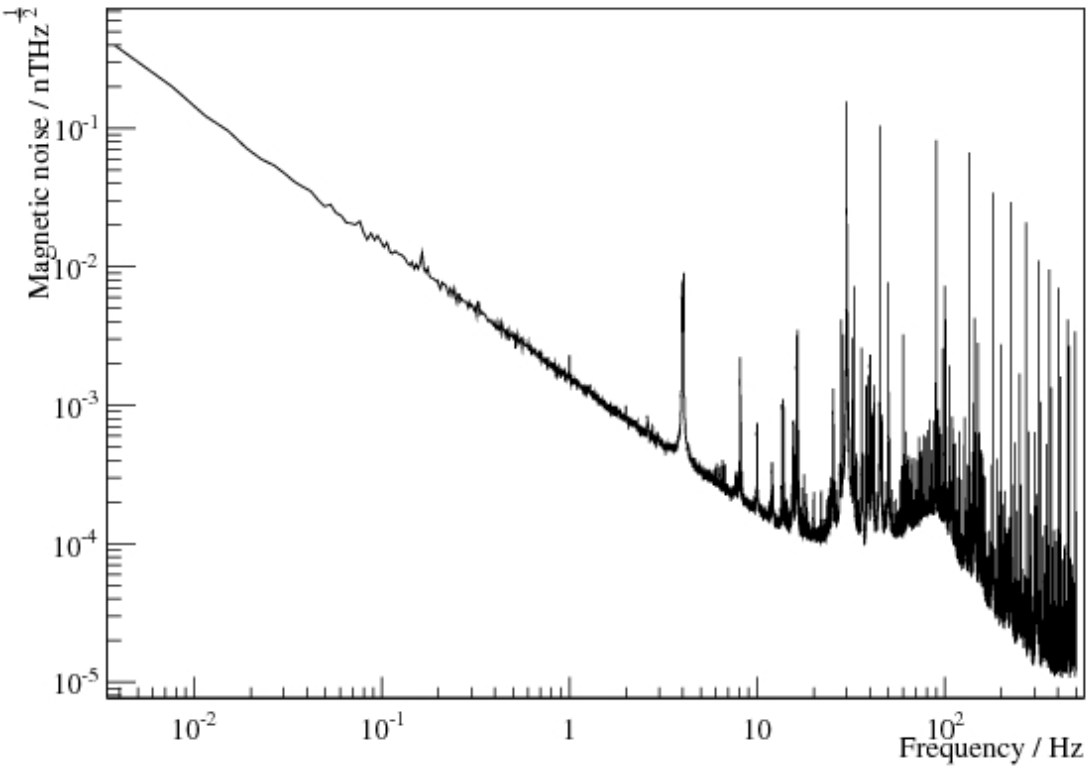

Figure 6.3: Shown is the noise spectra of the axial SQUID 11 in the same frequency range as figure 6.2. SQUID 11 has a resolution of 28 pT. Again the noise power is dominated by a $\frac{1}{f^2}$ contribution, and large peaks in the noise can be seen.

The noise on the SQUIDs in this region is dominated by a $\frac{1}{f^2}$ contribution, which is not uncommon in background magnetic noise [109]. So the SQUIDs are being operated in a mode where the environment is limiting their resolution and not the thermal noise of the SQUID. The higher frequency range above 4 Hz contains many peaks from unknown sources, critically these are in the region of 30 Hz to 200 Hz and could very well interfere with the precession of the neutrons at 150 Hz.

Figure 6.4 shows the noise spectra of SQUID 1, which has a pick-up loop in the transverse direction. This shows a similar behaviour to the axial SQUIDs, meaning there is also a lot of transverse noise. Oscillating fields in the transverse direction will cause a shift in the resonance frequency that scales as $\frac{B_{osc}^2}{\omega_0 - \omega_{osc}}$, given that there are components of the noise very close to the resonance frequency, the shift could be very large.

Figure 6.5 shows the noise spectra for an internal fluxgate in the frequency range 4×10^{-4} Hz

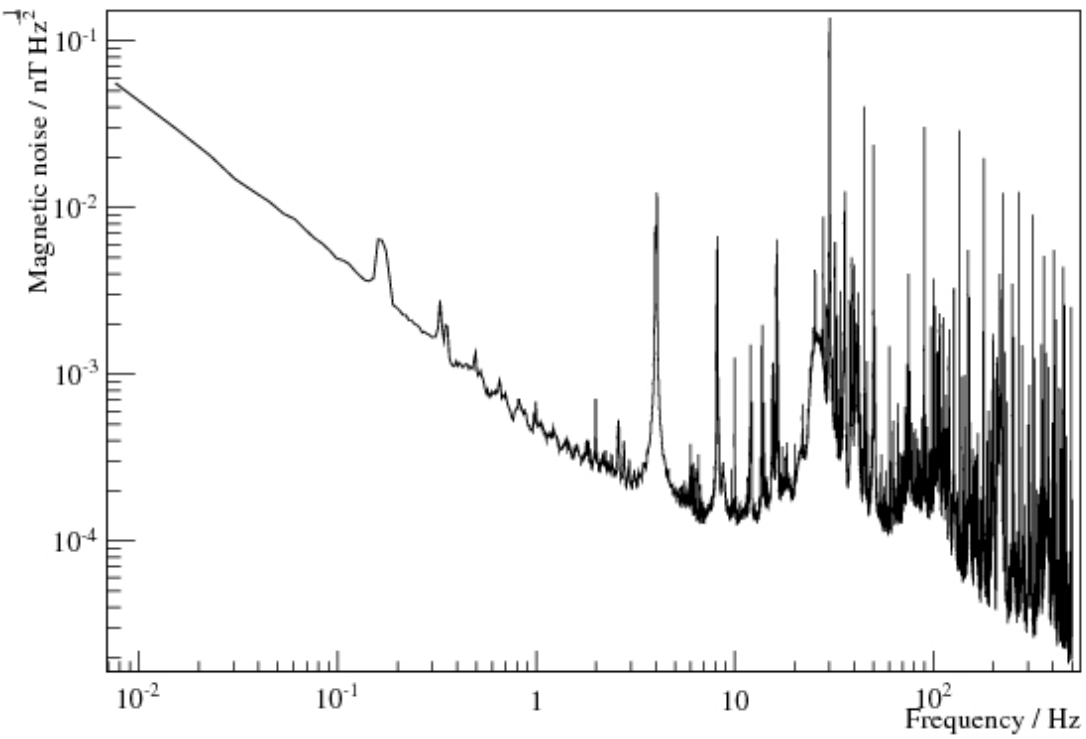

Figure 6.4: Shown is the noise spectra of the transverse SQUID 1 in the same frequency range as figure 6.2. SQUID 1 has a resolution of 25 pT. Again the noise power is dominated by a $\frac{1}{f^2}$ contribution, and large peaks in the noise can be seen. Oscillating fields in the transverse direction will cause a shift in the resonance frequency, and so are dangerous.

to 50 Hz. The noise on the internal fluxgates below 1 Hz is dominated by the $\frac{1}{f}$ contribution. This is what would be expected of a magnetometer being limited by internal noise. However above 1 Hz the white noise level of the fluxgate is reached. This white noise takes a value of 0.5 nT Hz$^{-\frac{1}{2}}$. The fluxgates should have a bandwidth of 10 Hz, so to see the white noise dominate above a frequency of only 1 Hz is an indication of a problem explored further in the next section.

Figure 6.6 shows the noise spectra for an external fluxgate in the same frequency range as figure 6.5. Below 1 Hz the noise on the external fluxgates is dominated by the $\frac{1}{f^2}$ contribution, which would indicate that the noise is coming from the field. The difference in low noise behaviour of the two fluxgates would indicate that the shielding is enough to lower the background noise below that of the instrument noise on the internal fluxgate. Similarly to the internal fluxgate, above 1 Hz the white noise level of 0.5 nT Hz$^{-\frac{1}{2}}$ is reached. The

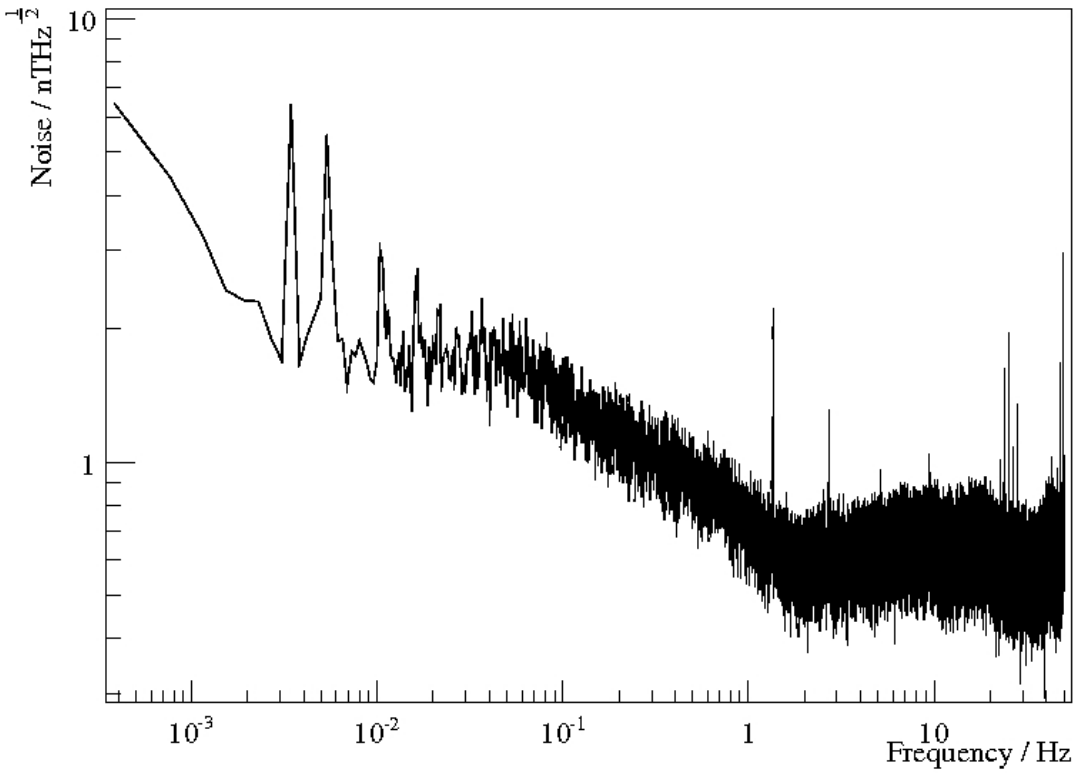

Figure 6.5: Shown is the noise spectra of an internal axial fluxgates in the frequency range 4×10^{-4} Hz to 50 Hz. Above 1 Hz the fundamental noise of the fluxgate is reached at a level of 0.5 nT Hz$^{-\frac{1}{2}}$. At low frequencies the internal fluxgate's noise power has a $\frac{1}{f}$ dependence, which indicates the internal fluxgate is shielded enough to reach its intrinsic noise limit.

fact that both fluxgates have the same white noise level is consistent with that being the fundamental limit of the magnetometers. The low frequency behaviour of the SQUIDs and external fluxgate being the same would also indicate that the $\frac{1}{f^2}$ contribution is very likely part of the experimental field.

6.3.1 Resolution

By using Parseval's theorem [110] it is possible to relate the noise spectra to the resolution of the magnetometer via

$$\sigma_B = \sqrt{\int |\tilde{B}(f)|^2 df}, \tag{6.1}$$

where σ_B is the RMS of the variations in the signal, taken as the resolution of the magnetometers, and $|\tilde{B}(f)|$ is the noise spectrum.

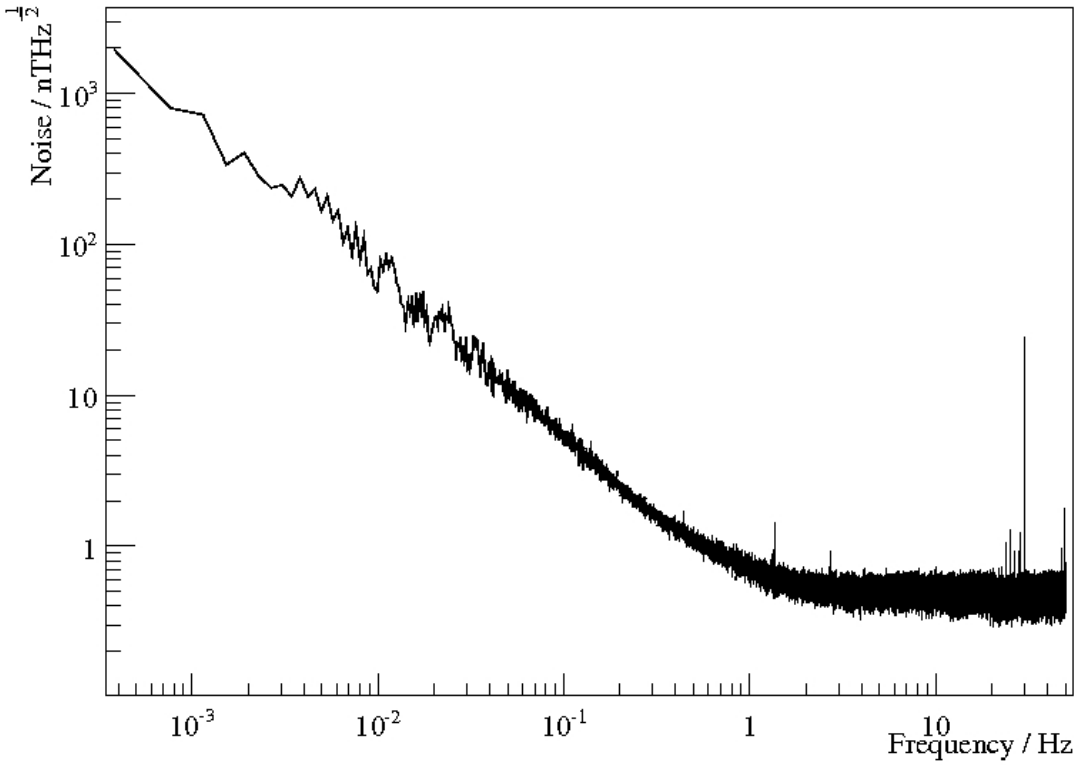

Figure 6.6: Shown is the noise spectra of an external axial fluxgate in the same frequency range as figure 6.5. Again above 1 Hz the fundamental noise of the fluxgate is reached at a level of 0.5 nT Hz$^{-\frac{1}{2}}$. At low frequencies the external fluxgate's noise power has a $\frac{1}{f^2}$ dependence, which indicates that the $\frac{1}{f^2}$ is a real property of the magnetic field.

Using this with the spectra above, without any signal processing, the deviations in the signal of the SQUIDs in the 4×10^{-3} Hz to 500 Hz band have an RMS of 95 pT for SQUID 5, 28 pT for SQUID 11, and 25 pT for SQUID 1. It is worth noting that the two SQUIDs with pick-up loops of the same dimensions have very similar resolutions. For the fluxgates in the 4×10^{-4} Hz to 50 Hz band, the RMS is 4.4 nT and 31 nT for the internal and external fluxgates respectively. The shape of the noise spectra and the similarity of these fluxgate resolutions despite the magnetic shielding is an indication that the resolution is not being limited fully by the magnetic noise, and so the fluxgates are being operated at close to the instruments' resolutions, especially in the case of the internal fluxgate. However, the quoted resolution of the fluxgates is 0.1 nT, which is significantly below the observed value, and so this is likely to be electrical noise rather the fundamental magnetometer resolution. The wiring of the fluxgates is such that once the sensor cables leave the cryostat they travel

approximately 3 m to the processing electronics, and then through a further 3 m of common ground cabling with numerous connector changers to the ADCs. This great distance between the magnetometer and the DAQ systems has great potential for introducing electrical noise, this situation is not helped by the fact that the full range of the fluxgate processing units is $\pm 2\,\mathrm{V}$, while the ADC range is $\pm 10\,\mathrm{V}$ reducing the effective ADC resolution to 13 bits rather than 16. Improvements in the cabling are planned to occur before the next data run, as well as the development of new DAQ hardware that contains on-board amplifiers to better make use of the full ADC resolution. This should improve the resolution of the fluxgates to reach a level closer to the 0.1 nT expected value.

The noise entering the SQUID pick-up loop scales very roughly with the inverse of the area of the pick-up loop [111], and so larger loops would have less noise giving a better resolution. However, larger loops increase the transmission of RF interference on to the SQUID chip and, as noted above, that can diminish the SQUID's performance.

As perturbations below these levels cannot be resolved, and so cannot be corrected, the resolution of the magnetometers will produce noise in the determination of an EDM signal. Using equation 2.5 with an electric field of 40 $\mathrm{kV\,cm^{-1}}$ and the RMS SQUID resolution in the displayed band, the RMS EDM noise induced on a typical batch of neutrons if each were the sole magnetometer would be: $7.1 \times 10^{-23}\,\mathrm{e\,cm}$ for SQUID 5; $2.1 \times 10^{-23}\,\mathrm{e\,cm}$ for SQUID 11; $1.9 \times 10^{-23}\,\mathrm{e\,cm}$ for SQUID 1; and $3.3 \times 10^{-21}\,\mathrm{e\,cm}$ for the internal fluxgates. Contributions to the uncertainty in the EDM will be 1000 times larger than that the of neutron counting statistics for SQUID 5, about 300 times larger than that of the neutron counting statistics for SQUID 11 and SQUID 1, and 45000 times larger than the neutron counts for the internal fluxgates. These are significantly above the experimental aim of resolution of EDMs of approximately $1 \times 10^{-25}\,\mathrm{e\,cm}$ per batch originating from the magnetometry's resolution.

To allow the SQUIDs to be operated in the sub-picotesla resolution range the pick-up loops must be made larger, and so elimination of the RF interference is essential. This alone will not be sufficient, as with the current environmental noise the pick-up loops would need

to have a radius on the order of 1 m, which is far from feasible, so it is necessary for the noise to be reduced by at least an order of magnitude for the largest pick-up loop that can be accommodated to have an acceptable resolution.

6.4 Internal magnetic fields

Magnetic perturbations in the neutron volume can come from internal sources of magnetic fields, or from external fields penetrating the magnetic shielding. The former are completely under the control of the experiment and can in principle be eliminated. A summary of the internal magnetic perturbations observed by the magnetometry system will be presented, and their sources and possible removal discussed.

6.4.1 Cryogenic artefacts

The experiment is not a thermally static system, the cryostat is continually using cryogens and often replenishing them. This can cause various thermal and physical stresses. An example of a large field change that correlates with cryogenic processes is shown in figure 6.7. This shows that there is a strong magnetic event, an approximately 15 nT change lasting approximately 20 seconds, after a spike in pressure on the nitrogen tanks surrounding the shield, solenoid and SCV assembly. The field then slowly returns to its previous value over the course of an hour.

The experimental volume is not fixed in place, to allow for relative thermal contraction of the components during cooling the solenoid and shields are suspended on flexible hangers connected to the cryostat wall. The shield is loosely coupled to the SCV, which is itself attached to the rest of the experiment by a set of steel bellows. The magnetic signal is consistent with the spike in nitrogen pressure, which leads to activation of a relief valve, causing a displacement of the experimental volume, producing the rapid change in magnetic field due to movement through an inhomogeneous field. Over time the set of bellows allow

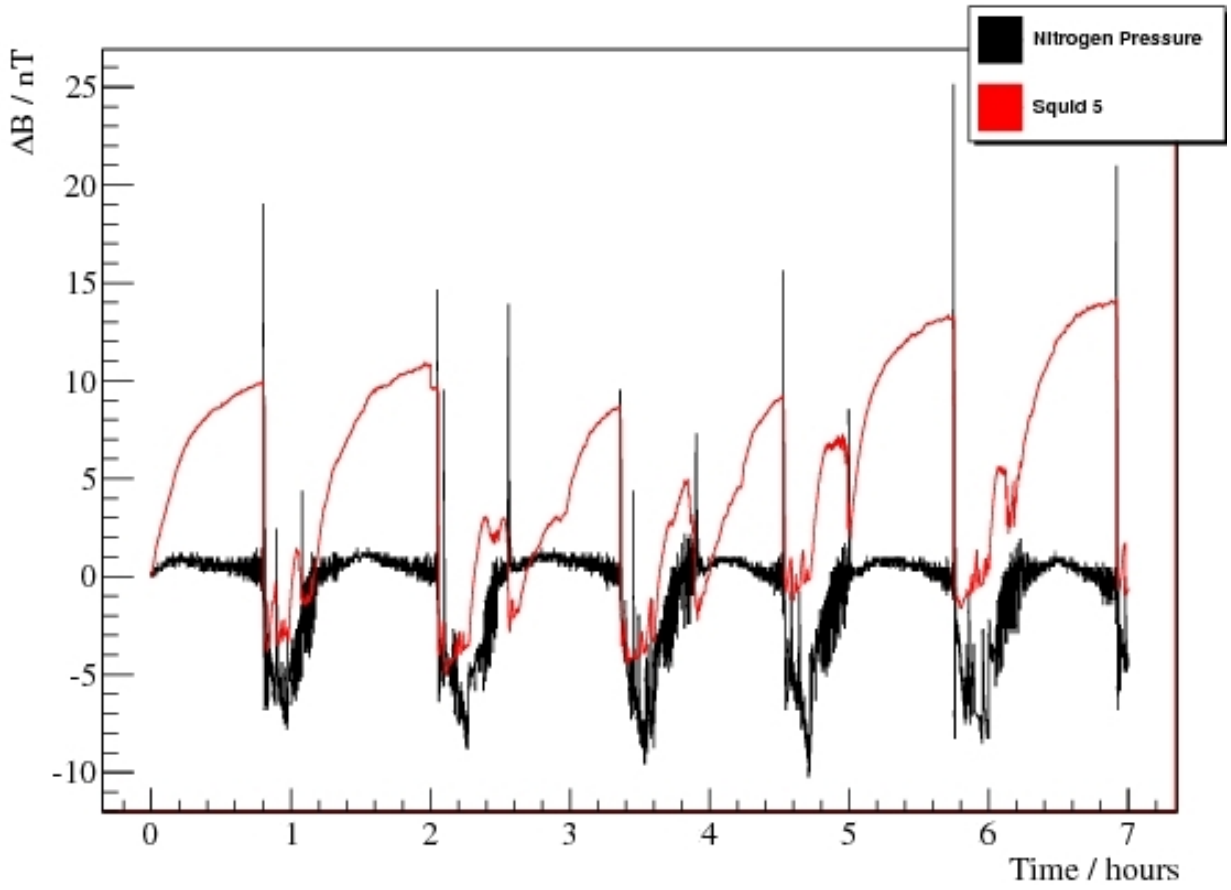

Figure 6.7: Shown is a typical signal from the SQUIDs seeing a periodic signal of approximately 15 nT magnitude (red), clearly visible is the association of this signal with large changes in the pressure of Nitrogen tanks in the cryostat (black).

the neutron volume to move back to its original position, causing the slow return in field seen. However, this alone does not narrow down the source of the inhomogeneous magnetic field.

This fluctuation appears on all the magnetometers, but varies in size and polarity. Figure 6.8 shows the response of the axial SQUIDs and internal fluxgates over the same period as shown in figure 6.7. The variations across the magnetometers show that the source is local rather than external to the cryostat. It is approximately equal and opposite on SQUID 5 and the nearest fluxgate, which indicates it is equidistant to the two magnetometers. It is also attenuated by the superconducting shields, which indicates it is outside the superconducting shielding. Assuming the source is dipole like, it is possible to get a rough idea of the position of the source from the relative size of the magnetometers' signals. Figure 6.9 shows the

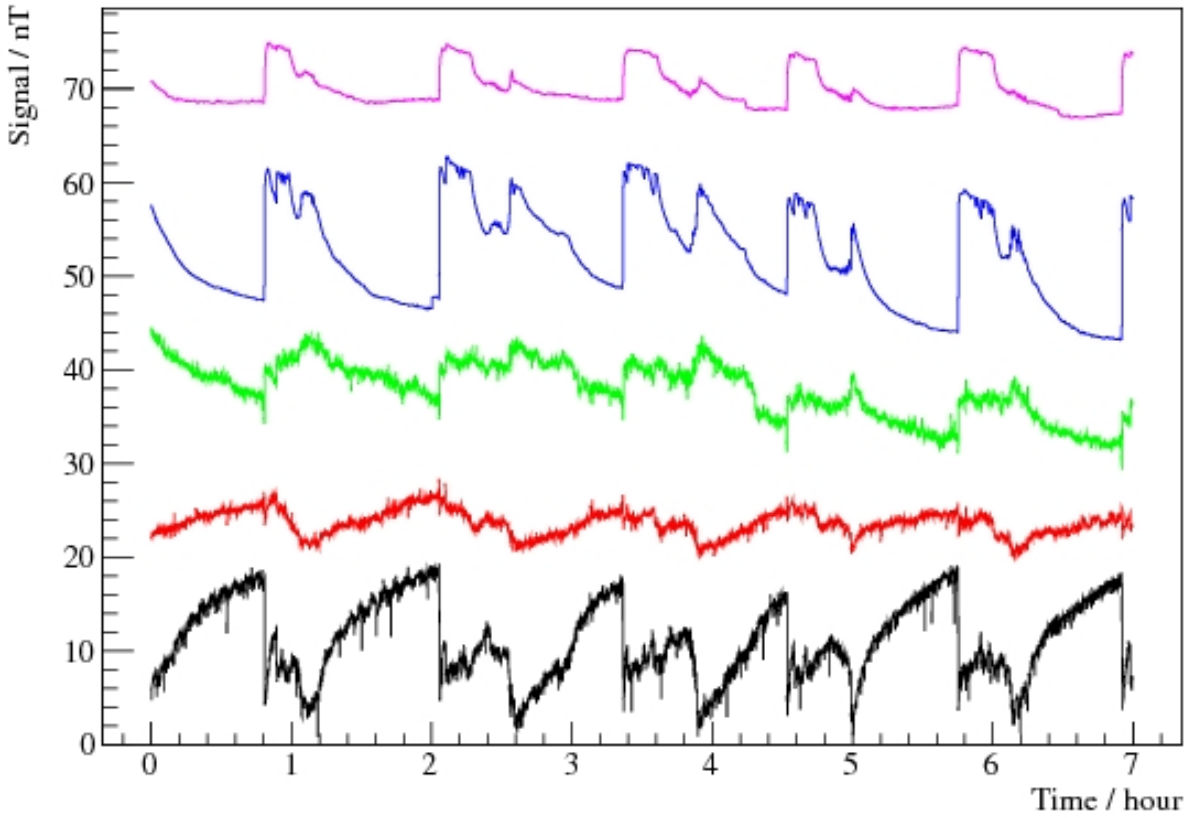

Figure 6.8: Shown is a typical magnetic signal arising from the cryogenic movement of the SCV. Seen is a periodic signal on the internal fluxgates (black, red & green) and axial SQUIDs (blue & magenta). The polarity and size variations hint to a local source, and can be used to narrow down its location.

points in the experiment that would give the observed SQUID and fluxgate responses, if occupied by a dipole source, to produce a guide to the location for the perturbation. These positions are located in the lower east quadrant of the experiment roughly in line with the neutron cells.

6.4.2 Valves and motors

The experiment has several moving parts, mostly valves and actuator rods that control the flow of the neutrons through the neutron volume. The valves are powered by motors, which are housed outside the main cryostat, and operate the valves via the actuator rods. Shown in figure 6.10 is the SQUID response when the valves controlling neutrons' entries to the

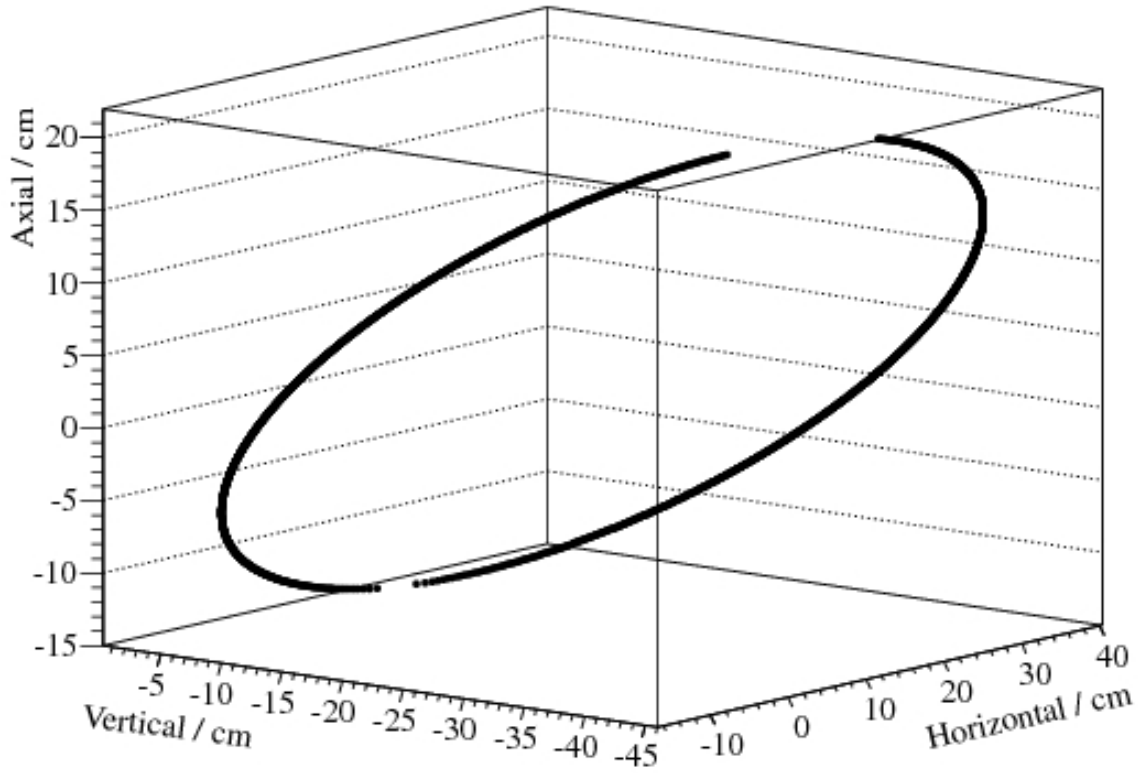

Figure 6.9: Shown are the approximate possible locations of the magnetic perturbation, responsible for the signal dependent on the nitrogen pressure, extrapolated from the magnetometer measurements assuming a dipole source. The majority of these locations are in the lower east quadrant of the experiment.

cells are moved. The first disturbance, starting at 0.29 s and lasting about 0.08 seconds, is the opening of the first cell. This causes a 100 Hz sine wave (aliased in the figure) on SQUID 5, and is likely due to the electromagnetic field of the motor. This sends the other two SQUIDs shown out of lock, as the signal is exceeding the slew rate of those SQUIDs. The second disturbance, starting at 0.37 s and lasting about 0.08 seconds, is the opening of the second cell. This causes a DC change in field as the actuator moves, and is likely due to a magnetic item attached to the actuator rod, e.g. a steel screw. Offline tests where the motors and rods are decoupled confirm that movement of one of the motors and movement of one of the actuator rods causes these signals. Other valves in the system also cause magnetic effects when used. This shows the magnetometry system in its current state can be used as a powerful diagnostic tool to determine, and hence improve, the state of the magnetic

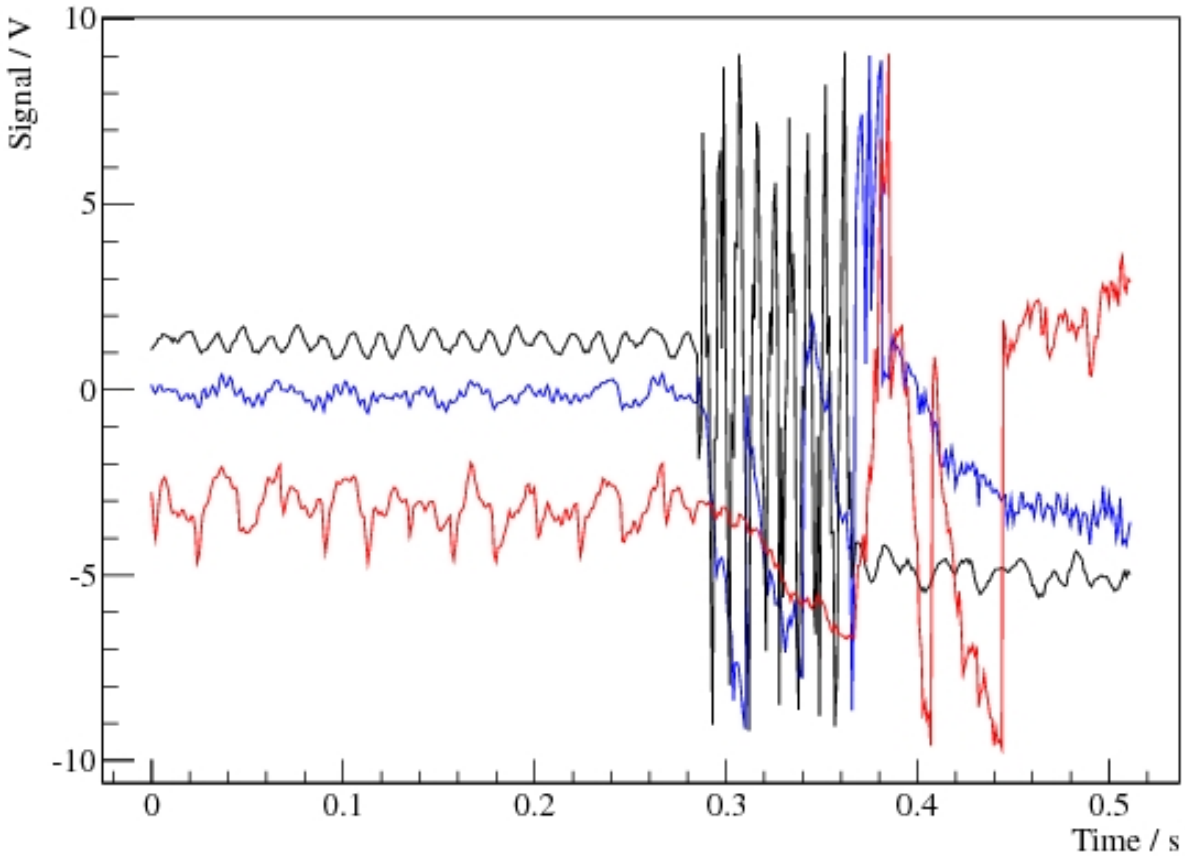

Figure 6.10: Shown is a typical signal from the SQUIDs caused by the operation of the neutral and high voltage cell valves. SQUID 5 (black), 11 (red), and 1 (blue) are stable until the first cell valve is operated at 0.29 s. A 100 Hz sinusoid is visible on SQUID 5 during this time, and SQUIDs 1 and 11 lose lock. The second cell is operated at 0.37 s and causes a change in the DC field.

environment within the experiment.

6.5 Magnetic shielding

As the experiment is conducted in a neutron scattering facility, there are large sources of external magnetic fields, especially from other experiments. Figure 6.11 shows the field produced by a neighbouring experiment measured using the external fluxgates. The axial component of the field shows large variations of around 300 nT over periods of 40 seconds and a repetition rate of about an hour. Uncorrected such fluctuations could cause a systematic uncertainty at the level of 2×10^{-19} e cm on a single batch of neutrons. With the magnitude

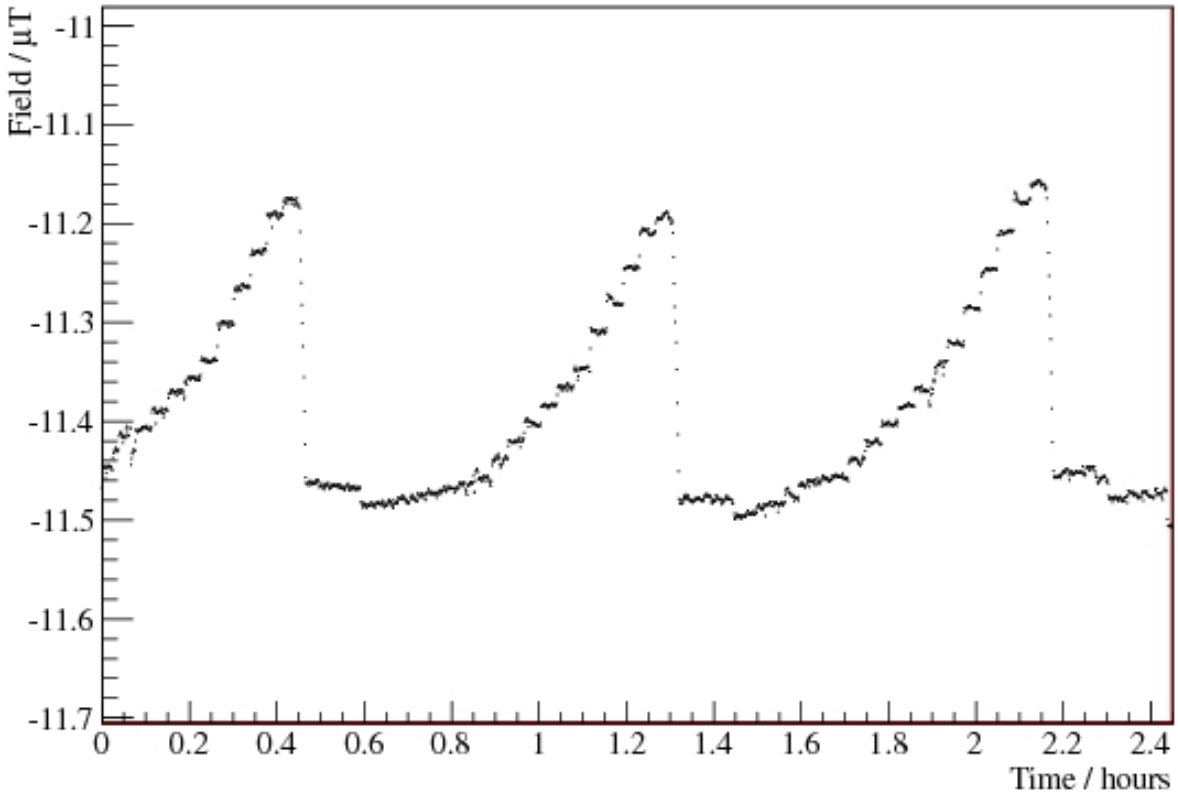

Figure 6.11: Shown is a typical example of the field measured in the vicinity of cryoEDM produced by a neighbouring spin-echo spectrometer (IN15). The signal shown is measured with the external fluxgate orientated along the experimental z-axis. Clearly visible is the ramping-up of IN15's experimental field, with a sharp decrease in field of 300 nT over 40 seconds. This signal has a period of approximately an hour.

of deviations of this scale, it is important to know how effective the magnetic shielding described in section 2.4.1 is at removing these perturbations. Shielding decomposes into two parts: the static shielding, the reduction of a fixed magnetic field; and dynamic shielding, the reduction of time dependent perturbations. As we wish to see the response to external perturbations, only the dynamic shielding will be investigated here. The shielding factor is determined as the ratio of the fields at a point if no shielding is in place and when shielding is in place. Measurements of the dynamic shielding factors of the cryoEDM shields will be performed in the next two sections, which will determine the dynamic shielding factors against rapid changes in the external field, and the frequency response of the shields respectively.

6.5.1 Shielding factor for discrete changes

In this section the shielding against large changes to the external field like that in figure 6.11 shall be determined. As the shielding cannot be moved to make this measurement, the mu-metal will always be present, so the field at the magnetometers without shielding is estimated from an external magnetometer. The source is approximately 20 m away from the experiment and so the field should be fairly homogeneous over the apparatus. This shielding factor is simply a measurement of ratio between signals produced outside the shields, and the response within the shields. Defining the shielding factor as

$$S_{f_i} = \frac{\partial B_{ext_i}}{\partial B_{int_i}} \ , \tag{6.2}$$

where B_{ext} is the external magnetic field, B_{int} is the internal magnetic field, and i the component of interest: axial or transverse. The quantity should be easy to extract from measurements of field changes taken both sides of the shields. However, as there are sources of magnetic signals within the shields that are considerably larger than the shielded external perturbations, the situation is complicated.

To extract shielding factors it is sufficient to locate periods of large external fluctuations that correspond to periods of stability in the internally produced fluctuations. To do this the measurement of the internal field during large external changes was fit linearly to the measurement of the external field over the same period of time. If the correlation between the signals was high enough, and the reduced χ^2 was sufficiently small, the period was deemed to be stable. The shielding factor could then be extracted from the fitting's scale parameter $\frac{\partial B_{ext}}{\partial B_{int}}$ directly.

The three forms of shielding each contribute to the total shielding factor. It is possible to run the three shields in the following configurations;

1. Mu-metal only

2. Mu-metal + superconducting shield

3. Mu-metal + compensation coils

4. Mu-metal + superconducting shield + compensation coils

During the 2010 data run, the experiment was operated in all four of these configurations while the magnetometers were functional. The shielding factors were determined for them as follows.

6.5.1.1 Mu-metal shields only

With both the compensation coils turned off and the superconducting shields normal, data can be taken in the mu-metal shields only mode.

In this mode external z-axis aligned fluctuations of approximately 250 nT were recorded along with internal fluctuations of nearly 7 nT. An example of such a fluctuation is plotted in figure 6.12 showing the external fluctuation on the external fluxgate and the internal response on SQUIDs 5 & 11. This results in an axial shielding factor for the first layer of shields of approximately 36.

The radially aligned SQUIDs show no resolvable changes that correspond to external field changes in this mode. To produce an estimate of the shielding factor in the transverse direction the FWHM noise of the SQUID signal must be used as an upper bound, which has a value of 0.24 nT over the period of the external changes. There is also uncertainty in the orientation of the pick-up loop within the cryostat, which had rotated during the cool down from its installation position. So to produce a conservative lower limit to the shielding factor the lowest valued transverse perturbation will be used, which is 40 nT. This gives a value of the transverse shielding factor as $S_{f\perp} > 160$. The transverse shielding of the mu-metal is estimated to be of order 10^2, so this limit is consistent with the estimate. As the perturbations are too small to be resolved through the mu-metal, this method will not provide shielding factors for the transverse direction for other shield combinations.

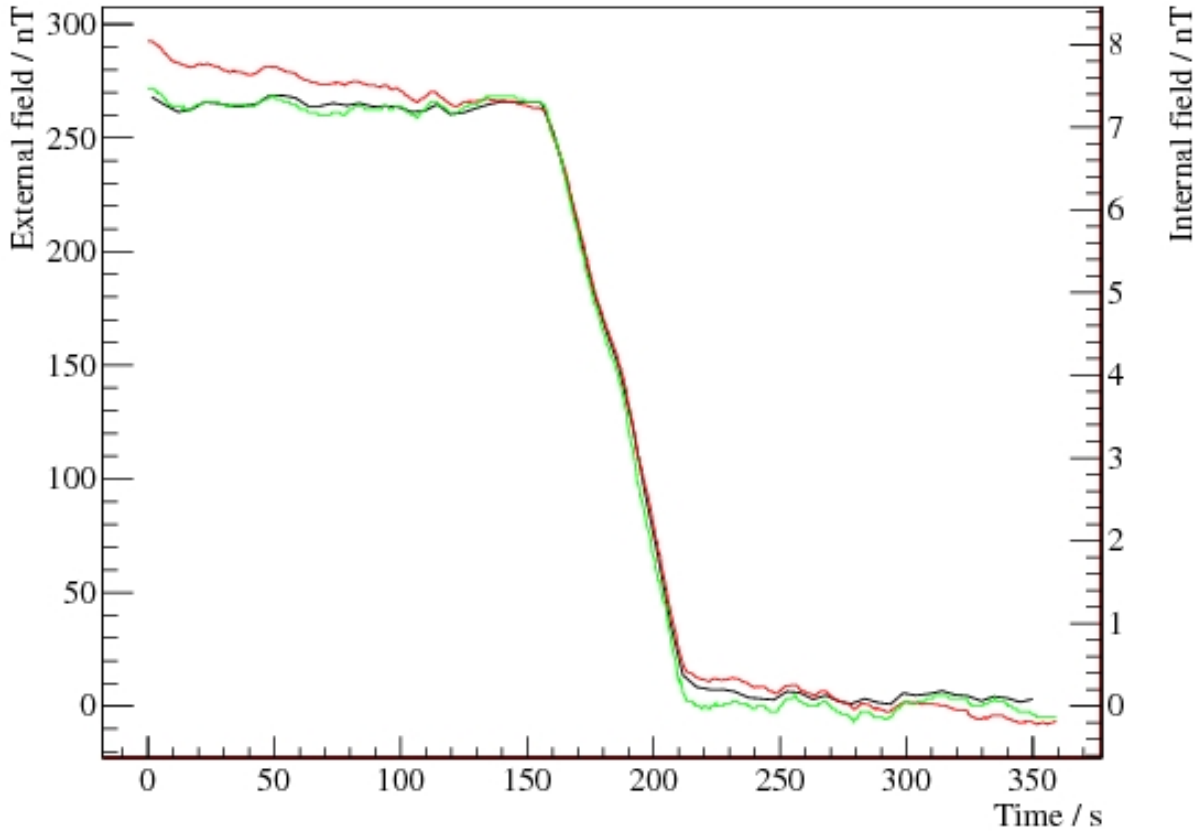

Figure 6.12: Shown is an example of an external field change observed outside the cryostat with a z-orientated fluxgate (black), and inside the mu-metal shields with the axial SQUIDs 5 & 11 (red and green), giving an axial shielding factor of 36.

6.5.1.2 Mu-metal & superconducting shields

With the compensation coils turned off, data can be taken in the mu-metal and superconducting shields mode. In this mode external z-aligned fluctuations of approximately 250 nT were recorded, along with internal fluctuations of nearly -0.5 nT. An example of such a fluctuation is plotted in figure 6.13. This results in an axial shielding factor for this combination of shields of -500.

It is interesting to note that the shielding provided by the superconducting shields is negative, as would be expected from flux leakage through the end of a finite superconducting tube [112]. The shielding effect of this combination is much lower than the 10^8 expected. This will be discussed further in the individual shields section below.

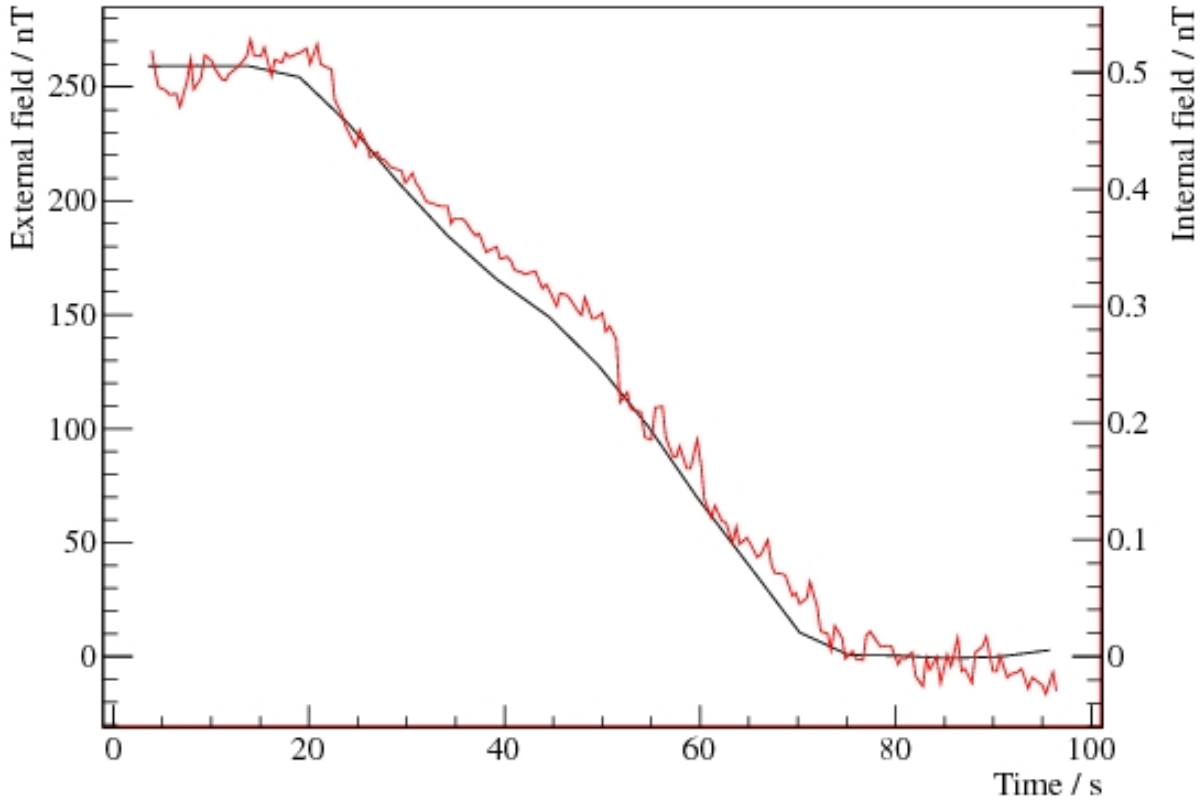

Figure 6.13: Shown is an example of an external field change observed outside the cryostat with a fluxgate (black), and inside the mu-metal and superconducting shields with SQUIDs (red), giving an axial shielding factor of about −500. The SQUID response has been inverted to make the comparison more obvious.

6.5.1.3 Mu-metal shields & compensation coils

With the superconducting shields normal, data can be taken in the mu-metal shields and compensation coils mode. In this mode external z-aligned fluctuations of approximately 300 nT were recorded along with internal fluctuations of nearly 0.9 nT. An example of such a fluctuation is plotted in figure 6.14. This results in an axial shielding factor for this combination of shields of 300.

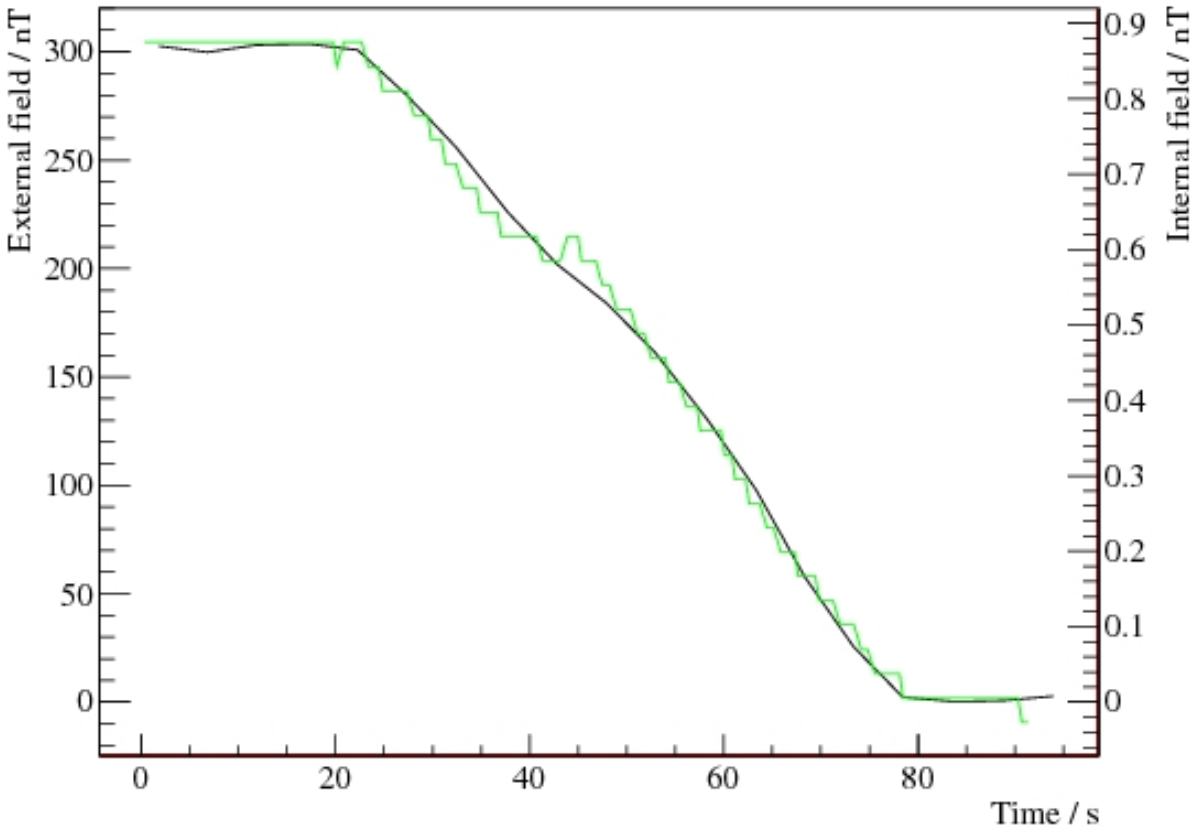

Figure 6.14: Shown is an example of an external field change observed outside the cryostat with a fluxgate (black), and inside the mu-metal shields and compensation coils with SQUIDs (green), giving an axial shielding factor of about 330.

6.5.1.4 All magnetic shields

With all shields operating, data can be taken in the all shields mode. In this mode external z-aligned fluctuations of approximately 300 nT were recorded along with internal fluctuations of nearly 60 pT. An example of such a fluctuation is plotted in figure 6.15. This results in a axial shielding factor for this combination of shields of 5000.

6.5.1.5 Individual shields

From the above measurements the shielding factor for each shield can be extracted. Given the geometry of the shielding in cryoEDM the shielding factors are expected to be multiplicative [83]. As there are four measurements and three unknowns it is possible to construct a better

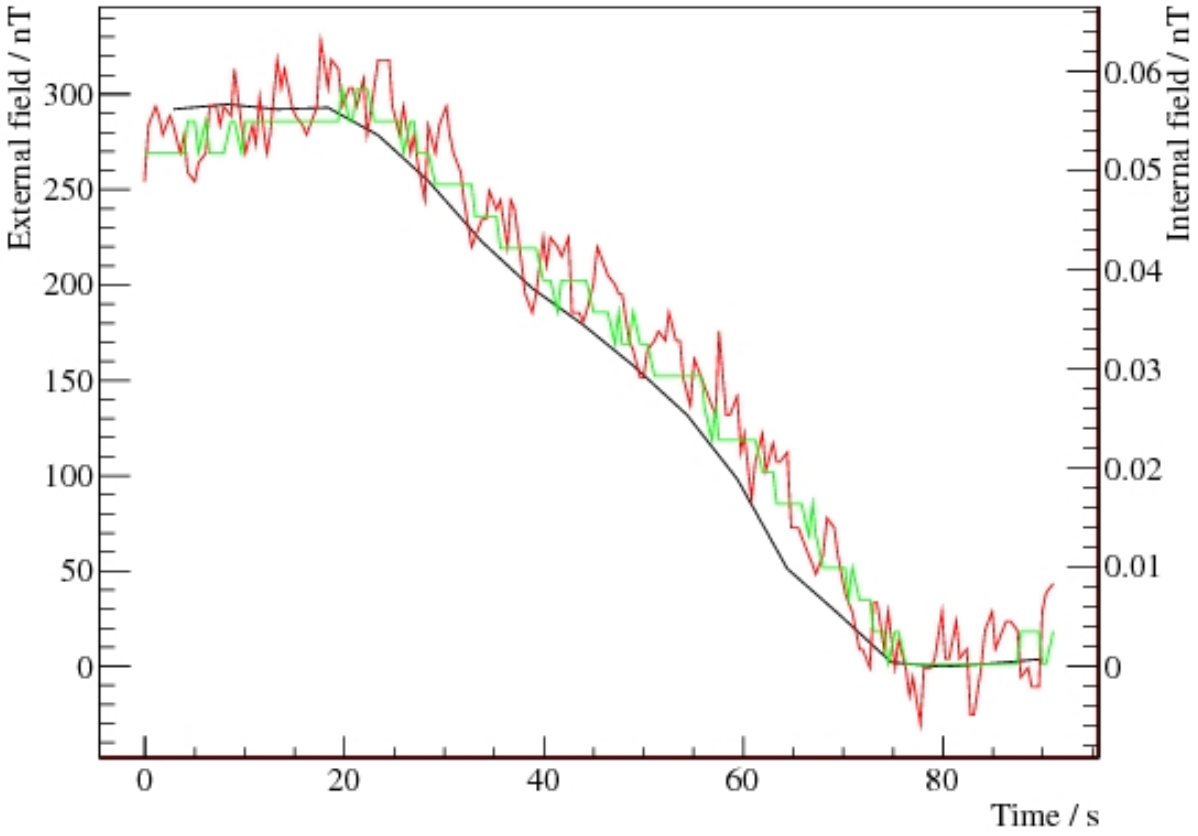

Figure 6.15: Shown is an example of an external field change observed outside the cryostat with a fluxgate (black), and inside all magnetic shields with SQUIDs (red and green), giving an axial shielding factor of about 5000.

estimate of the shielding factors of the individual shields than simply calculating the ratios. Figure 6.16 shows an example of the procedure required, finding the point in parameter space that minimises the square distance to all the surfaces defined by the measurements. Letting μ be the mu-metal shielding factor, ς the superconducting shielding factor, and κ the compensation coil shielding factor, the point $\boldsymbol{P} = \begin{pmatrix} \mu_{min} \\ \varsigma_{min} \\ \kappa_{min} \end{pmatrix}$ is the point that minimises

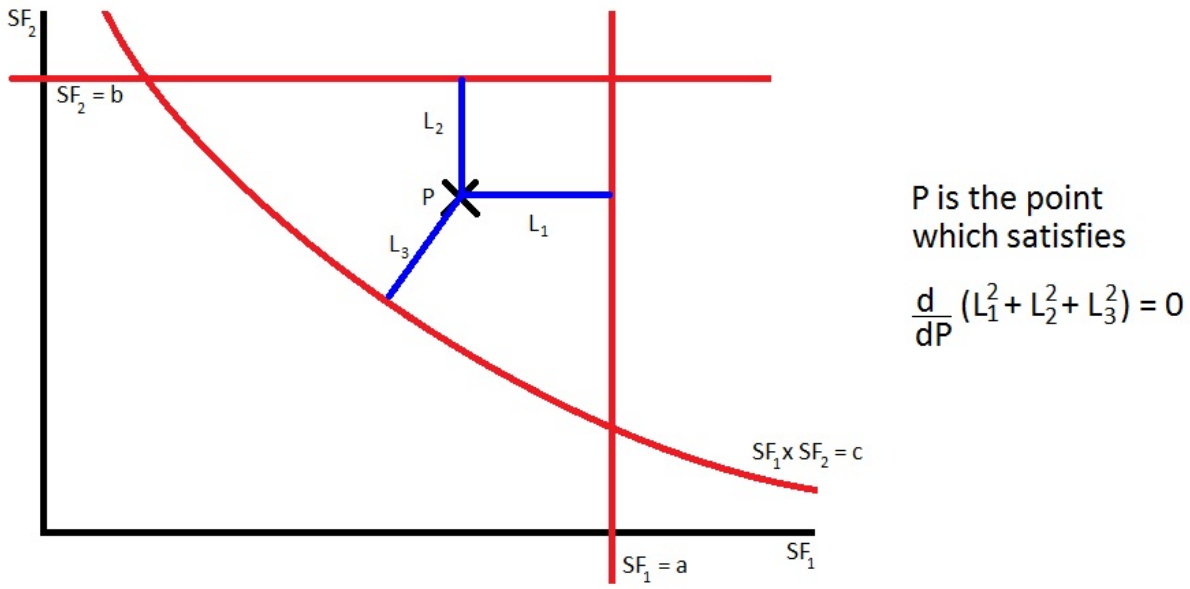

Figure 6.16: Shown is a 2D example of extracting individual shielding factors of two shields from three measurements. The shielding factors are the point in the parameter space that minimises the shortest distance to all three curves. This extends to the 3D case where the distances must be calculated to surfaces rather than curves.

the distance to the curves defined by

$$S_\mu = \mu \tag{6.3}$$

$$S_{\mu\varsigma} = \mu\varsigma \tag{6.4}$$

$$S_{\mu\kappa} = \mu\kappa \tag{6.5}$$

$$S_{\mu\varsigma\kappa} = \mu\varsigma\kappa \tag{6.6}$$

The square distance to the curve defined by equation 6.3 is simply $\delta_\mu^2 = (\mu_{min} - S_\mu)^2$. The square distance to the curve given in equation 6.4 is

$$\delta_{\mu\varsigma}^2 = (\mu_{min} - \mu)^2 + (\varsigma_{min} - \varsigma)^2 \tag{6.7}$$

or

$$\delta_{\mu\varsigma}^2 = (\mu_{min} - \mu)^2 + \left(\varsigma_{min} - \frac{S_{\mu\varsigma}}{\mu}\right)^2 \tag{6.8}$$

the minimum of this is at the point μ which is the root of the quartic

$$\mu^4 - \mu_{min}\mu^3 + \varsigma_{min}S_{\mu\varsigma}\mu - S^2_{\mu\varsigma} = 0 \tag{6.9}$$

There is an analogous quartic for the curve in equation 6.5. For the curve in equation 6.6 the roots of two coupled quartics must be found, these are

$$\mu^4 - \mu_{min}\mu^3 + \kappa_{min}\frac{S_{\mu\varsigma\kappa}}{\varsigma}\mu - \frac{S^2_{\mu\varsigma\kappa}}{\varsigma^2} = 0 \tag{6.10}$$

$$\varsigma^4 - \varsigma_{min}\varsigma^3 + \kappa_{min}\frac{S_{\mu\varsigma\kappa}}{\mu}\varsigma - \frac{S^2_{\mu\varsigma\kappa}}{\mu^2} = 0 \tag{6.11}$$

To find the roots of these quartics the GSL root finding library [113] is used. The Minuit package [114] was used to minimise the resulting distance. Finally, using the Monte Carlo method with the measured shielding factors and their uncertainties as sample distributions, allows the covariance matrix for the shielding factors to be computed readily. The individual shielding factors are tabulated in table 6.1, and the uncertainty ellipses plotted in figure 6.17.

Shield	Shielding factor	Uncertainty	Approximate expected value [115]
Mu-metal shield	36.4	1.0	50
Lead shield	-14.3	0.9	1×10^6
Compensation coil	8.5	1.3	15

Table 6.1: Table showing shielding factors and their uncertainties for individual magnetic shields around the experiment.

The measurement of the mu-metal shielding factor is quite close to the estimate of its shielding factor. Due to the complicated shape of the mu-metal shielding, estimates from simulations vary between 10 and 50 and so are consistent with the measurement. The compensation coils have a shielding factor of half the expected value, although the shielding from the coils is dependent on the gain in the feedback system, which was not optimised during this run so a deviation is not troubling. However, the superconducting shield has a shielding

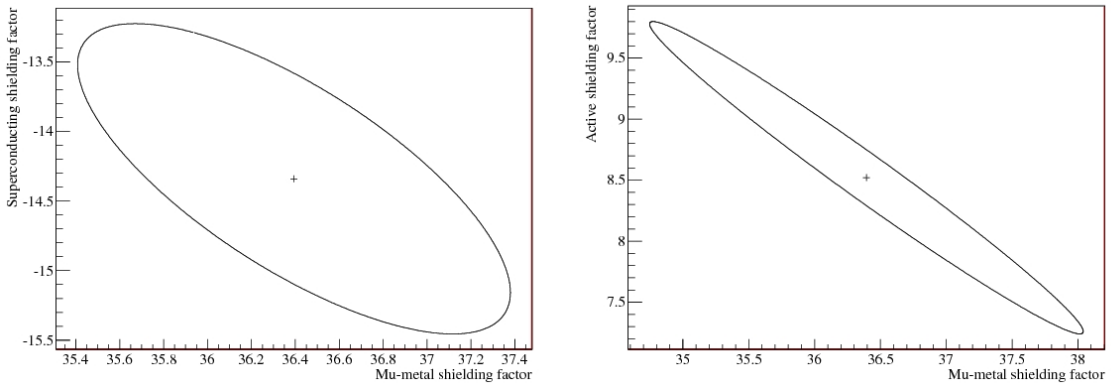

Figure 6.17: Shown are the uncertainty ellipses for the shielding factors plotted in the mu-metal-superconducting plane, and the mu-metal-compensation-coil plane respectively. These are plotted to show 68% confidence.

factor 10^5 smaller than expected. The disparity between the observed and expected shielding factors of the superconducting shield is due to the construction of the solenoid, which is nearly the same length as the superconducting shield, and inside the shield. Figure 6.18 shows the effect of magnetic field lines from the external field distorted by the end of the superconducting shield. The end effects of the shield cause the field lines to pass between the wires forming the solenoid. Changes in the external field change the field cutting the solenoid's wires and induce a current in the solenoid, which creates a field in the central part of the shield, i.e. the experimental volume. Without the solenoid the perturbations through the superconducting shield to the central volume would be much smaller, an expected value of approximately 1×10^6 would be achieved in this case.

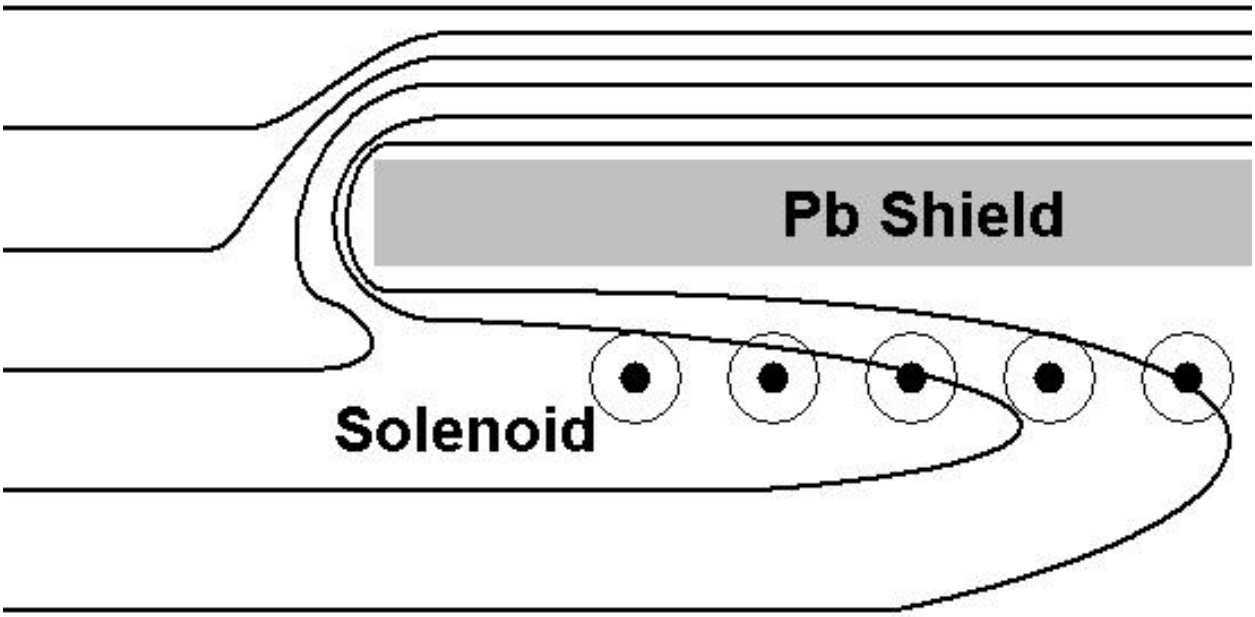

Figure 6.18: Shown is a schematic of magnetic field lines interacting with the superconducting shield and solenoid. End effects from the shield cut through the wires of the solenoid, so changes in the field outside the shield induce a current in the solenoid, and are therefore transmitted to the central region of the shield, the experimental volume.

6.5.2 AC shielding factor

Of course it is not only rapid changes in the field that are of concern, but also frequencies in the range of interest where it can significantly affect the measurements. It would be advantageous to have knowledge of the shielding the experiment has from these perturbations. In this section the shielding factors as a function of frequency will be calculated.

The most practical way of obtaining the frequency response is to Fourier transform the environmental noise. This assumes that the major sources of environmental noise are produced from outside the cryostat, this should be a reasonable assumption due to the magnetic homogeneity requirements of the experiment. Given that the environmental noise is smaller than the fields observed from IN15, larger uncertainties are expected than in the measurements in section 6.5.1. Here the shielding factor will be defined as

$$|\tilde{S}_f(f)| = \frac{|\tilde{B}_{ext}(f)|}{|\tilde{B}_{int}(f)|} \tag{6.12}$$

where an over tilde represents the Fourier transformed signal. The shielding factor can therefore be computed from the Fourier transform of two time domain signals, one shielded and one not.

6.5.2.1 Total and Mu-metal shielding factor

As there are no calibrated SQUIDs outside of the magnetic shielding, any AC shielding factors obtained for arrangements including the mu-metal shield would require the use of the fluxgate data. However, as shown in section 6.3 the fluxgates' response is dominated by the instrument's white noise above 1 Hz, and so at these frequencies the fluxgate is not measuring the true magnetic field it is experiencing and a measurement using the fluxgates would then not produce an interesting frequency range. An example measurement of the mu-metal shielding factor using a fluxgate and SQUID 5 is given in figure 6.19, there are artificial changes in the measurement caused by the fluxgates limited resolution. To avoid this issue all AC shielding factors have to be determined by the SQUIDs, this gives its own problems as all the calibrated SQUIDs are within the shielded regions. To obtain values of the shielding factors will require data taken with and without the shields operational, assuming the noise level remains constant over the measurement periods. This can be checked using uncalibrated SQUIDs that are installed outside of the shielded area. As the mu-metal can not be removed, this means that any shielding factors including the mu-metal shields can not be obtained.

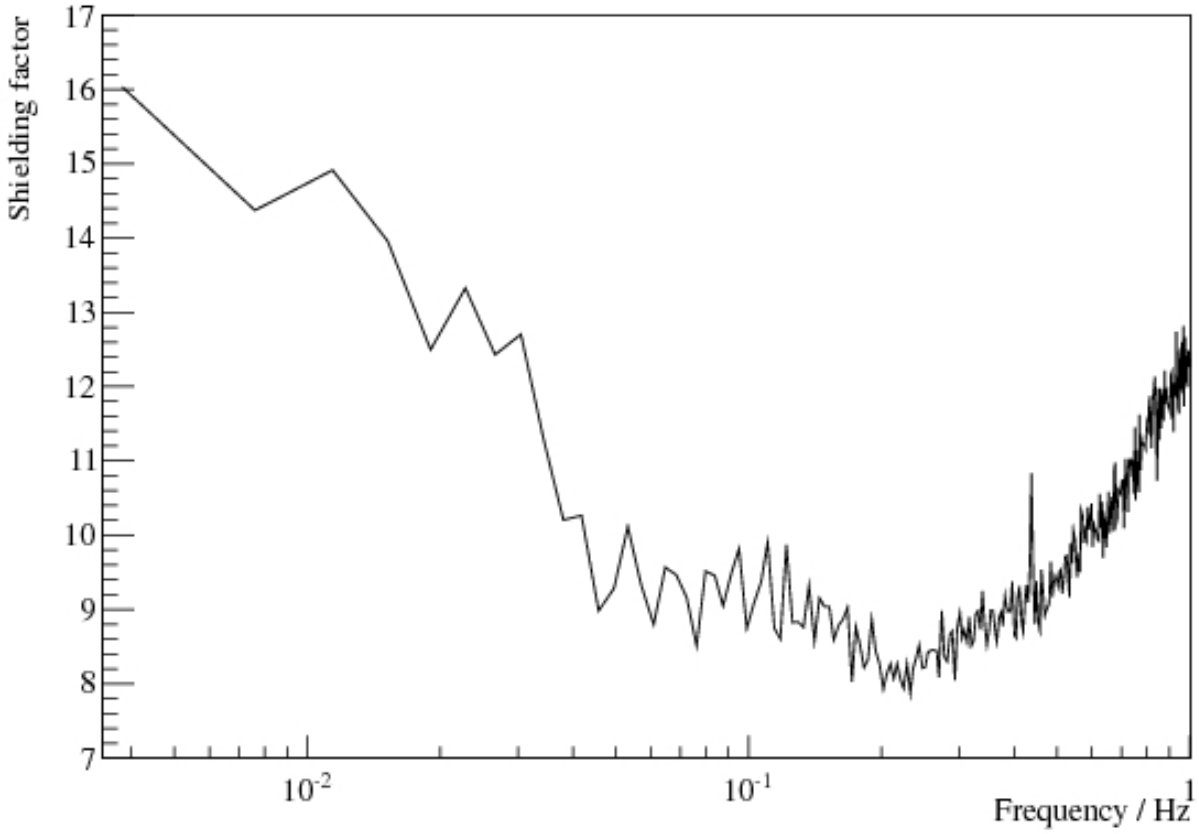

Figure 6.19: Shown is the frequency dependence of the mu-metal shields below 1 Hz taken from SQUID 5 and an external fluxgate. The shielding factor appears to decrease from 16 to 8 over the range. It is unclear how much the limited resolution of the fluxgates is affecting this measurement in this region. Near 1 Hz the loss in resolution of the fluxgates cause the shielding factor to artificially increase.

6.5.2.2 Superconducting shield shielding factor

The AC shielding factor of the superconducting shielding was computed from the SQUID response immediately before and after the superconducting shield became normal. This time displacement should not have a significant effect as an uncalibrated SQUID outside of the shielding showed no significant changes in noise between the two periods.

Shown in figure 6.20 are the individual noise spectra and in figure 6.21 is the frequency response of the shielding factor calculated between 4×10^{-3} Hz to 500 Hz. The shielding factor below 10 Hz seems quite flat and has a value of approximately 10, which is close to the value obtained in section 6.5.1. There is a dip in the shielding factor just above 0.1 Hz which is cause by a small noise peak on top of the background on the two spectra. There is

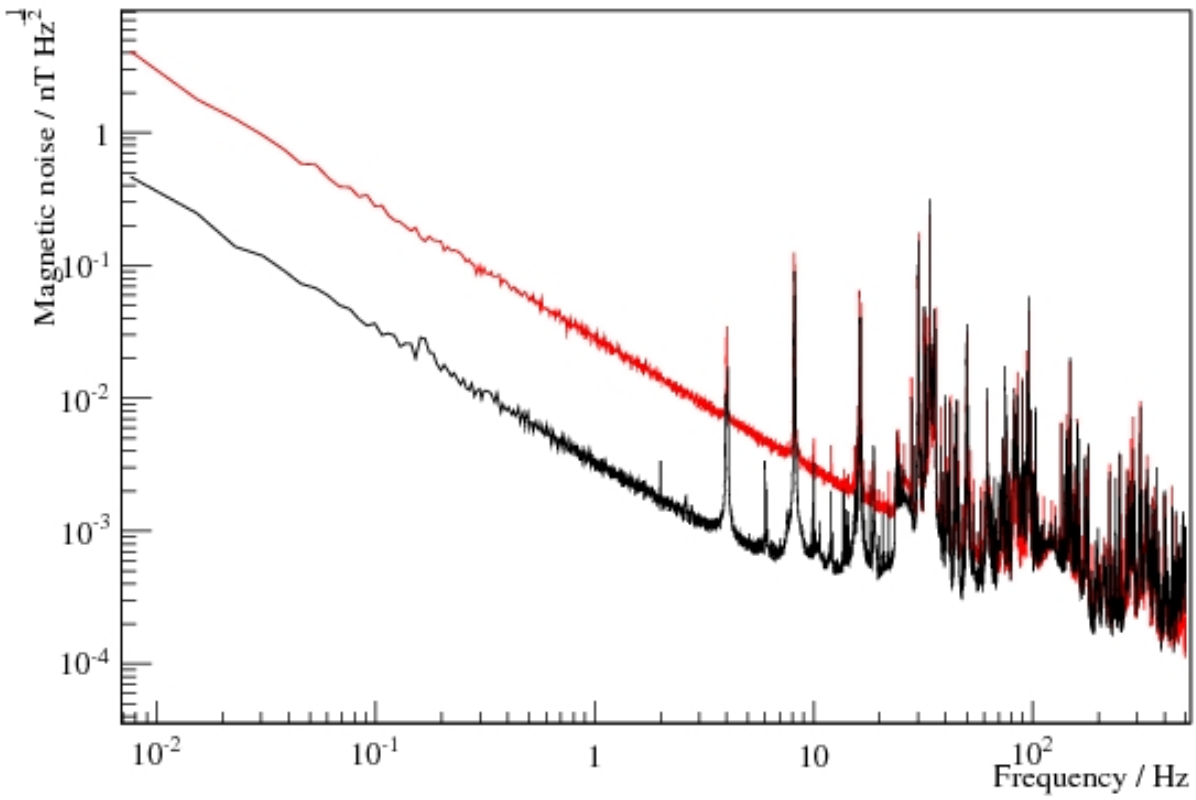

Figure 6.20: Shown are the noise spectra for SQUID 5 with superconducting shielding (black) and without superconducting shielding (red). In the high frequency region the spectra are dominated by noise peaks and both take very similar values. This would indicate that the source of the peaks must be within the shields to be unaffected by the shielding.

a sharp cut off above 4 Hz where the shielding falls to 1. This does not necessarily indicate that the superconducting shields transmit frequencies above this without attenuation. Given that in this region the spectra are dominated by noise peaks rather than the background, a more likely explanation is that the dominant source of the higher frequency magnetic noise is internal to the cryostat, and so cannot be removed by the shielding.

If the sources of these > 4 Hz noise peaks are within the shields, this narrows down the likely origin greatly. The only sources of large axial magnetic fields within the shields are the field trim coils, the AC Ramsey coil, the fluxgate cabling, and the SQUID calibration coils. It is possible that external sources of noise are being transferred inside the shields by these coils. The source is yet to be identified, but if it is from one of the above coils then a low pass filter with a 1 Hz cutoff on the cabling to these coils should remove this noise in most

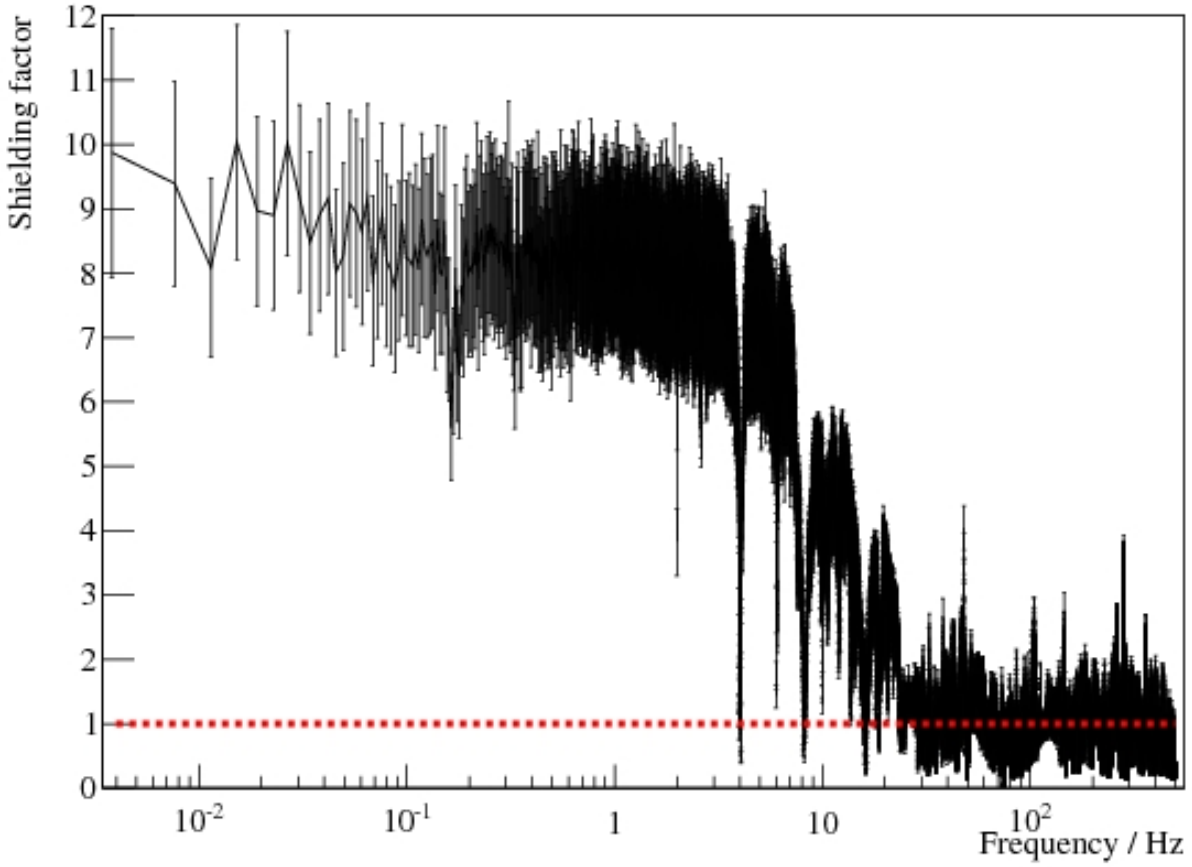

Figure 6.21: The shielding factor of the superconducting shield as a function of frequency. This has a value of approximately 10 and appears constant in frequency. The sharp drop in shielding above 10 Hz is likely caused by the sources of these frequencies being within the shield, rather than full transmission of these frequencies through the shield.

cases. It would be problematic if the Ramsey coils are introducing this noise, as frequencies of around 150 Hz are required to be transmitted as part of the measurement procedure.

Figure 6.22 shows the two spectra for a transverse SQUID, these produce the shielding factor given in figure 6.23. This shows identical features to the axial shielding factor, except the flat region has a value of approximately 30. This indicates that the transverse shielding of the superconducting shield is better than that of the axial shielding, which is likely due to the effect of the solenoid discussed above.

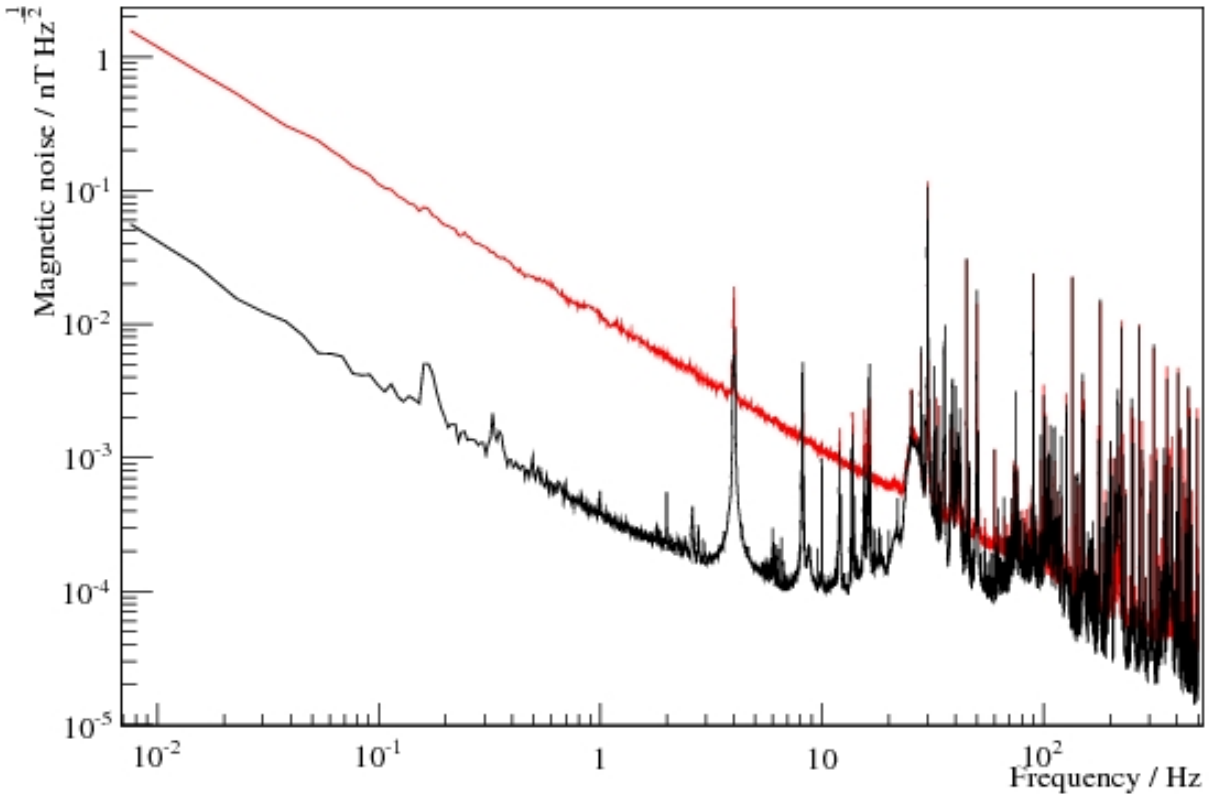

Figure 6.22: Shown are the noise spectra for a radially aligned SQUID within the shields, SQUID 1, with superconducting shielding (black) and without superconducting shielding (red). Again at higher frequencies the noise peaks dominate, causing the two spectra to be very similar in value.

6.5.2.3 Compensation coil shielding factor

The shielding factor of the compensation coils was computed from the SQUID response immediately before and after the shield was deactivated.

The response was so close to unity that figure 6.24 only shows a histogram of the shielding factors measured in the range 4×10^{-3} Hz to 500 Hz. The frequency dependence of the compensation coils' shielding factor is 1.0 ± 0.2, and so the compensation coils have no effect on AC frequencies. This is expected as the compensation system output is passed through a DC amplifier, which will act as a low pass filter before being applied to the coils.

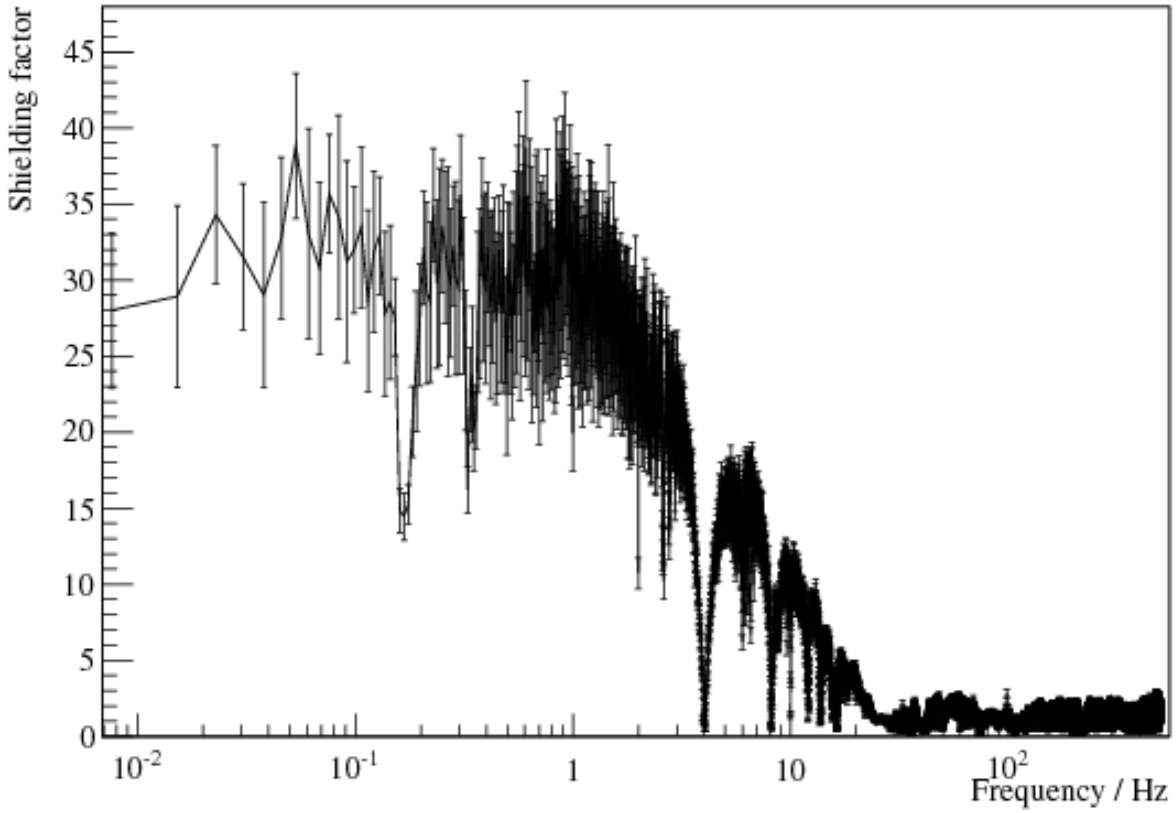

Figure 6.23: The radial shielding factor of the superconducting shield as a function of frequency. This has a value of approximately 30 and appears constant in frequency. The sharp drop in shielding above 4 Hz is likely caused by the sources of these frequencies being within the shield, rather than full transmission of these frequencies through the shield.

6.5.3 Conclusion

The above results show that the shielding factors of the current shields are far from the design intention, which was an axial shielding factor of 10^8, a factor of 2×10^4 smaller in fact. This is mostly due to the effect of the proximity of the solenoid to the superconducting shield as discussed in section 6.5.1.5. This is planned to be corrected in the near future by installation of a second superconducting cylinder inside the solenoid and surrounding the SCV, calculations show this should recover the missing shielding. It also shows that there is a significant source of high frequency $> 4\,\mathrm{Hz}$ noise produced within the cryostat, which cannot be shielded against. These noise peaks are present in the axial and transverse direction. This will require reexamination of the filtering of inputs to the magnetic coils in

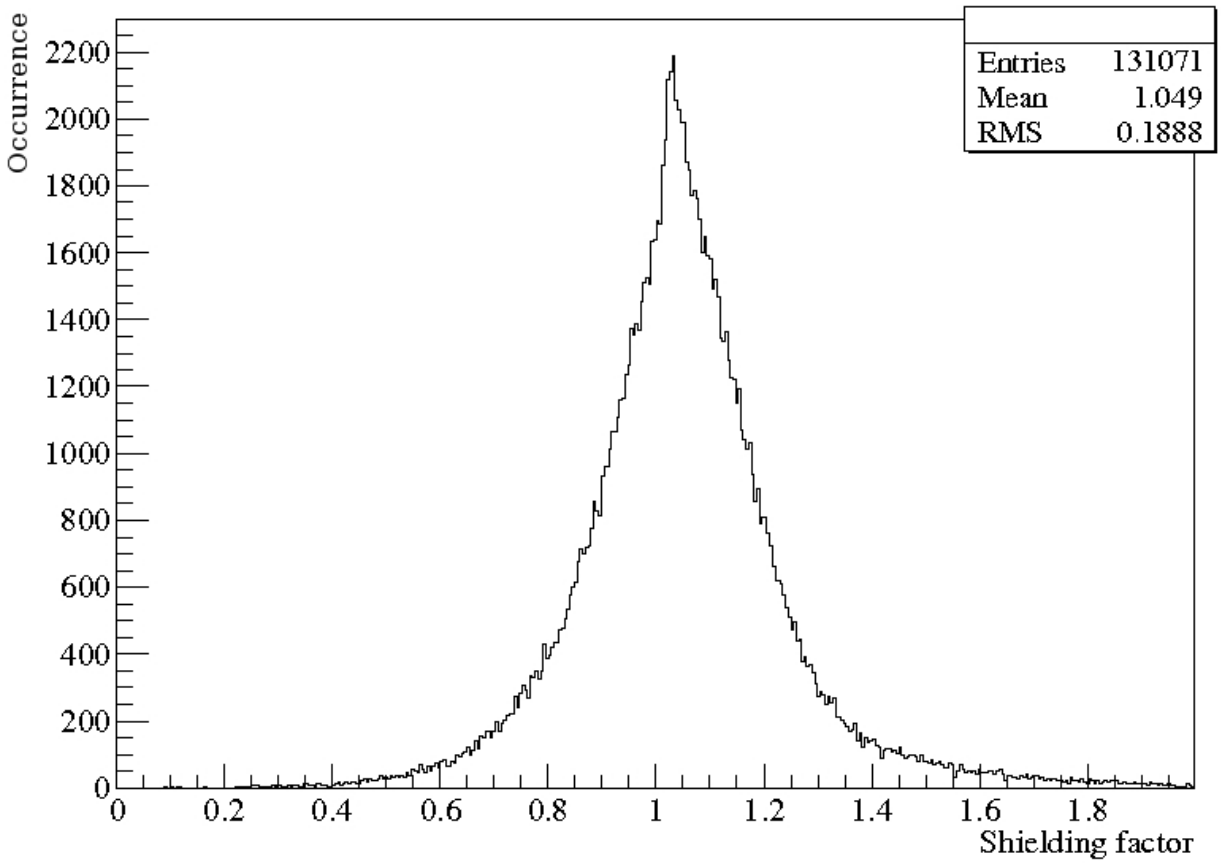

Figure 6.24: Shown is a histogram of the shielding caused by the compensation system for AC frequencies. The response was constant in frequency, and so just the amalgamated distribution for all frequencies is shown. This has a mean value of close to 1, and is significantly below the value of the DC shielding factor. This is not unsurprising as the compensation system is designed to offset the DC changes.

and around the SCV to determine the source.

6.6 SQUID slew rate

It was observed that some SQUIDs lost lock during the application of the AC fields used in a Ramsey resonance. This means that the rate of change of the AC field was exceeding the slew rate of the SQUID assembly, so the feedback system could not linearise the SQUID output correctly. In normal operation a y-axis aligned sinusoidal field of approximately 50 nT and 150 Hz will be applied to the neutrons, this will apply a peak change in field of around $47\,\mu\mathrm{T\ s^{-1}}$. Due to the geometry of the coil the applied field will be less than this

in the x and z directions and will be highly position dependent, but to be conservative this value will be taken as the minimum slew rate required by all the SQUIDs. The AC coil was used to determine approximate values for the slew rate of the SQUIDs arrangements in the magnetometry system.

A set of sinusoidal signals were sent through the AC coil at various frequencies and amplitudes. Figure 6.25 shows the response of SQUID 5 and 11 to a 1 Hz field applied by the coils. SQUID 5 shows a sinusoidal field of 5 nT as desired. The behaviour of SQUID 11 is different, in the high gradient areas of the sine curve, the slew rate of this SQUID is exceeded and so SQUID 11's response is deformed. To determine a value of the slew rate, the rate of change of field that causes the SQUID response to no longer vary sinusoidally must be determined, which in practice means finding the lowest frequency and amplitude settings that cause the deformation seen from SQUID 11 in figure 6.25.

A hardware failure interrupted the data taking so only limited data set was recorded, making determination of a precise value for the slew rate unachievable. However the upper and lower limits on the slew rate were found as the rate of change measured either side of the transition. Table 6.2 shows a summary of the slew rates found for each of the calibrated SQUIDs.

SQUID	Slew rate (lower limit) / nT s^{-1}	Slew rate (upper limit) / nT s^{-1}
1	100	200
4	113	226
5	5600	∞
11	79	158

Table 6.2: Table showing the upper and lower limits obtained on the slew rate of the calibrated SQUIDs.

SQUID 5 did not become unstable during any of the applied signals and so its upper limit was indeterminable. The large difference between the slew rates of SQUID 5 and SQUIDs 1, 4, & 11 is likely due to the smaller size of the pick-up loop attached to SQUID 5. This has

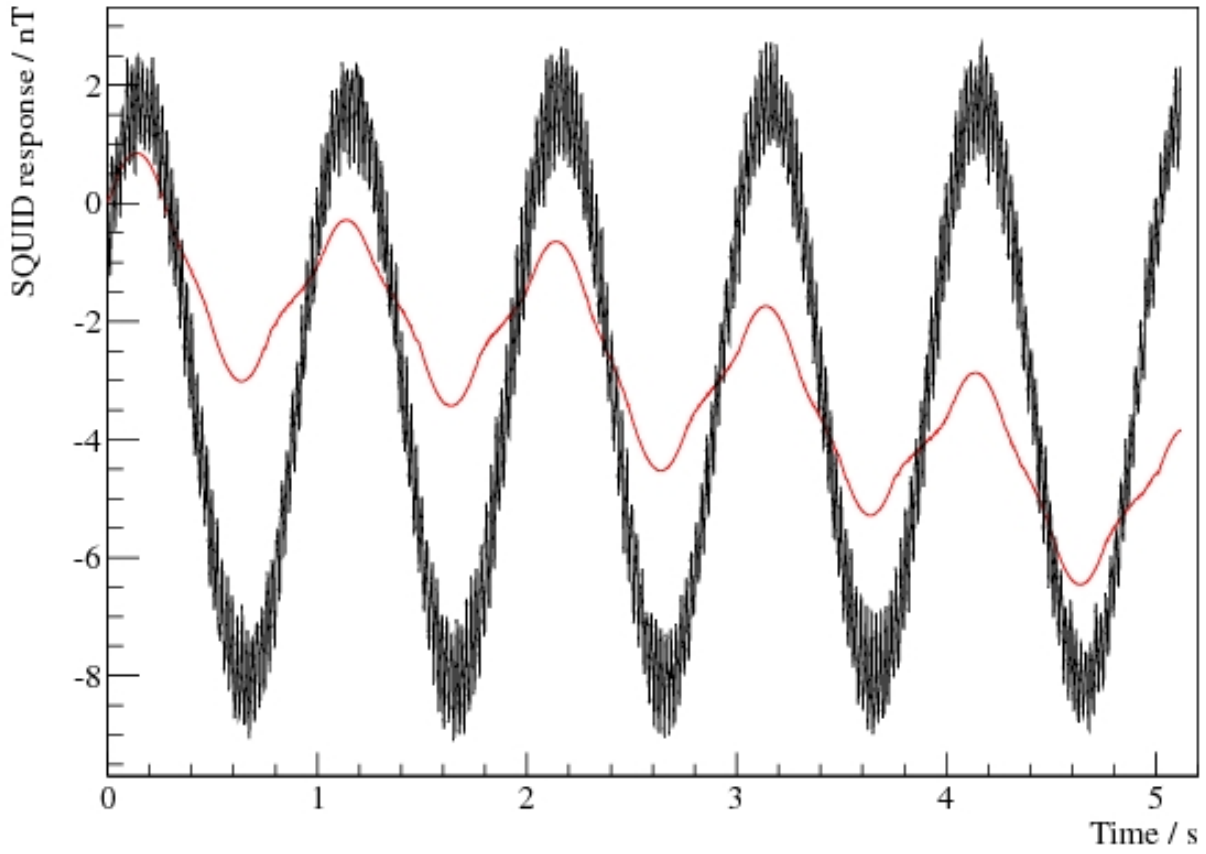

Figure 6.25: Shown is the response of two SQUIDs to the field produced from the AC coil. SQUID 5 (black) shows a well defined sine curve with an amplitude of 5 nT and frequency of 1 Hz. SQUID 11 (red) shows a deformed curve due to the instability caused in the quickly changing region of the curve where the slew rate has been exceeded.

two effects: a smaller pick-up loop would produce comparatively less current through the loop in response to an environmental field change, and hence a smaller change in flux being applied to the SQUID chip; perhaps a larger effect is that with the larger loop area, SQUID 1, 4, & 11 will be subjected to larger amounts of RF interference and, due to the known RF problems which prevent the other installed SQUIDs from functioning, these SQUIDs are likely operating close to their tolerance, which will reduce the slew rate. As the SQUIDs are manufactured in a standard way, it is likely the slew rates of the SQUIDs without pick-up loops are of a similar magnitude.

The upshot is that in this environment the current magnetometry system does not, in general, have a slew rate that will cope with the required AC signal applied during a

measurement cycle. The measured slew rate is more than two orders of magnitude away from the needed value. The magnetometry needs to be able to track the field during the whole of a Ramsey resonance measurement, and ideally be able to determine the magnitude and frequency of the applied rotation field, which is not possible if the SQUIDs become unstable during the AC field application. This situation requires urgent rectification as it makes the magnetometry system not fit for purpose. Reducing the size of the pick-up loops will reduce both the current and the RF interference being applied to the SQUID chip, and so increase the slew rate of the system. As noted earlier this will negatively impact on the resolution of the magnetometers, and so the elimination of the source of the RF interference is preferable. In the mean time, further investigation of pick-up loop designs to maximise resolution and minimise susceptibility to RF interference should be performed, along with confirmation of RF interference being the limiting factor in the system's performance.

Chapter 7

Conclusions

To place a limit on the electric dipole moment of the neutron, cryoEDM requires a high quality real-time magnetometry system to correct for false EDM signals and to act with control software to keep the magnetic field at the correct operating point.

To accomplish this a new analysis scheme was developed, which relates the magnetic signal observed by available magnetometers to the volume average of the modulus of the magnetic field experienced by the neutrons occupying the neutron cells.

Additionally two methods of calibration were presented: an updated version of a direct calibration from magnetic sources housed within the cryostat, which agrees well with calibrations obtained from independent sources; and a new scheme that calibrates the magnetometers using measurements of the neutron polarisation was developed, and tested in a similar environment to that in cryoEDM to good agreement.

Hardware limitations that introduce artefacts to the SQUIDs' magnetometry signals can be overcome in software, where corrections for these artefacts were performed and shown to correlate well with the applied signals when the noise is significantly smaller than the size of the artefacts.

Marrying each of these aspects together, a new software framework was created which can function as a data acquisition system, a real-time monitoring system, and operate as

part of an environmental control system.

Using the newly developed system the magnetic environment of cryoEDM could be characterised. The resolution of the SQUIDs in the current setup was found to be limited by the environmental noise. The magnetometer with the best resolution currently installed would give a RMS noise to the EDM signal of 2×10^{-23} e cm per batch of neutrons. This is 300 times larger than the uncertainty obtained from neutron counting statistics and is not ideal. Currently the RF environment is too harsh to increase the size of the sensing region, which would improve the resolution, however with the current environmental noise even this would not be sufficient. With the planned addition of extra superconducting shielding, both the RF interference and noise should be reduced, allowing the aimed for sensitivity to be achieved.

A significant amount of high frequency noise was observed inside the cryostat, the frequency of this noise is such that it could potentially interfere with the neutron precession at 150 Hz. The source of this high frequency noise is shown to come from within the cryostat itself, and is likely an unfiltered coil installed, and so should not be difficult to remove.

The performance of the magnetic shielding could also be determined. The design of the magnetic shielding was such that an axial shielding factor of 10^8 could be expected. An axial shielding factor closer to 5×10^3 was observed. Given the typical perturbations to the external field from neighbouring experiments, with all shields operational the magnetic field changes would be large enough to create a false EDM signal of 4×10^{-23} e cm per batch if left uncorrected. However, this will improve, once the additional shielding is installed, to a level close to the original design.

Sources of magnetic perturbations that correlate with cryogenic processes in the experiment were also investigated. The largest of which was caused by a build up of pressure in the surrounding nitrogen tanks. Release of this pressure causes large magnetic field changes, likely from movement of the experiment. If uncorrected these changes would yield a false EDM signal of 1×10^{-21} e cm per batch. The location of the source of the perturbation can be estimated assuming it is a dipole, placing it in the lower-east quadrant of the experiment.

Additional to removal of the source, the cryogenic operation should be evaluated to remove the sudden release of pressure triggering the perturbation.

Further sources of magnetic perturbations come from operating certain aspects of the experiment. Operation of the necessary neutron control valves gives rise to both AC magnetic perturbations from the driving motors at a frequency of 100 Hz, and DC field changes from movement of magnetic components. These perturbations destabilise the SQUIDs, which can then no longer track field changes. Perhaps more importantly these changes can depolarise the neutrons as they move through the experiment, reducing the overall sensitivity.

Finally the dynamic performance of the SQUIDs was measured in response to typical AC fields, which are a necessity for the experiment to function. These fields cause some SQUIDs to become unstable during a measurement, meaning they cannot track the magnetic field during this time. As these fields are applied during a Ramsey measurement, during which excellent knowledge of the field is required, it is of critical importance to resolve this issue. The slew rate of the magnetometry system needs to be increased by at least two orders of magnitude. The reduced slew rate is caused by the SQUID pick-up loops being too large for the amount of RF interference in the experiment, which degrades performance. As the pick-up loop size needs to be increased to improve the magnetometer resolution the only option is, again, the elimination of the high levels of RF interference present.

In some respects the performance of the magnetometry system and the characterisation of the magnetic environment in cryoEDM paint an undesirable picture of the outlook for cryoEDM's magnetometry. However, it is important to note the steps taken in getting the magnetometry to this level. It has progressed from being viewed with suspicion, to being an important diagnostic tool within the experiment in the space of just two years. Additionally many of the problems outlined have common sources. For example if the magnetic environment is improved then the SQUIDs can be operated to reach their true resolution and drop the false EDM signal below other systematics.

A thorough overhaul of the magnetic items in cryoEDM must be performed urgently to ensure they meet the required specifications. Conveniently the experiment is scheduled to

relocate to a new neutron beam line within the ILL in 2012, this will require disassembly and reassembly of the whole experiment. This would provide the perfect opportunity to remove any current magnetic items and prevent the installation of any further items. It is hoped this process will improve the magnetic environment enough for the magnetometry to meet its ultimate performance.

Appendix A

Extracting the EDM from neutron measurements

The electric dipole moment of the neutron is inferred from measurements of the polarisation of bunches of neutrons after undergoing a Ramsey cycle. This appendix will outline how one goes from the polarisation measurements to a value of the EDM from the theoretically determined Ramsey curve.

A.1 Value of d_n

Starting from the line shape around the central minimum of a Ramsey resonance curve given in [99]. The number of neutrons expected with a precession frequency ω_0 rotated by an AC field with frequency ω, applied for a time τ with a period of free precession of T is given by

$$N_{counts} = N_{ave}\left(1 \mp \alpha \cos\left[(\omega - \omega_0)\left(T + \frac{2\tau}{\pi}\right)\right]\right), \tag{A.1}$$

where the negative indicates the spin-up and positive spin-down counts respectively, N_{ave} is the average population of the two states, and α is the product of the polarisation and the spin analyser efficiencies. In cryoEDM the measurement time is much longer than applied

field time, so we can apply the condition $T \gg \tau$. The difference between spin up and spin down is given by

$$\delta N = -2N_{ave}\alpha \cos\left[(\omega - \omega_0)T\right]. \tag{A.2}$$

Assuming the electric and magnetic fields are either parallel or antiparallel, the precession frequency is $\omega_0 = \omega_B \pm \omega_E$ respectively

$$\delta N = -2N_{ave}\alpha \cos\left[(\omega - (\omega_B \pm \omega_E))T\right]. \tag{A.3}$$

Using the cosine summation rules to expand this gives

$$\delta N = -2N_{ave}\alpha \left(\cos\left[(\omega - \omega_B)T\right]\cos\left[\omega_E T\right] \mp \sin\left[(\omega - \omega_B)T\right]\sin\left[\omega_E T\right]\right). \tag{A.4}$$

Subtracting these counts from aligned and antialigned fields gives

$$\delta N_{\uparrow\uparrow} - \delta N_{\uparrow\downarrow} = 4N_{ave}\alpha \sin\left[(\omega - \omega_B)T\right]\sin\left[\omega_E T\right]. \tag{A.5}$$

The experiment is carried out at the point with the largest gradient to maximise the difference in counts for change in frequency, the point where $(\omega - \omega_B)T = \frac{\pi}{2}$ and $\sin\left[(\omega - \omega_B)T\right] = 1$. It is also noted that the electric dipole moment is small, so $\omega_E T \ll 1$ and therefore $\sin\left[\omega_E T\right] \approx \omega_E T$. This leads to

$$\delta N_{\uparrow\uparrow} - \delta N_{\uparrow\downarrow} = 4N_{ave}\alpha\omega_E T. \tag{A.6}$$

Remembering that $\omega_E = \frac{2d_n E}{\hbar}$ this becomes

$$\delta N_{\uparrow\uparrow} - \delta N_{\uparrow\downarrow} = 4N_{ave}\alpha\frac{2d_n E}{\hbar}T. \tag{A.7}$$

Noting that $4N_{ave}$ is the total number of neutrons detected in both spin states and in both field orientations, denoted by N, equation A.7 can be rearranged to give the electric

dipole moment of the neutron

$$d_n = \frac{\hbar(\delta N_{\uparrow\uparrow} - \delta N_{\uparrow\downarrow})}{2EN\alpha T}. \tag{A.8}$$

This is not the only number of note, the uncertainty on this value must be determined as well.

A.2 Uncertainty in d_n

Assuming the uncertainties on the electric field, polarisation and measurement time are small, the dominating effect will be uncertainties in the neutron counting, which will be Poisson distributed. The uncertainty will be given by

$$\sigma^2(d_n) = \left(\frac{\hbar}{2E\alpha T}\right)^2 \left[\sigma^2(\delta N_{\uparrow\uparrow} - \delta N_{\uparrow\downarrow}) \left(\frac{dd_n}{d(\delta N_{\uparrow\uparrow} - \delta N_{\uparrow\downarrow})}\right)^2 + \sigma^2(N) \left(\frac{dd_n}{dN}\right)^2\right]. \tag{A.9}$$

Both $\sigma(\delta N_{\uparrow\uparrow} - \delta N_{\uparrow\downarrow})$ and $\sigma(N)$ equal $\sqrt{N}$. This leads to

$$\sigma^2(d_n) = \left(\frac{\hbar}{2E\alpha T}\right)^2 \left[N \left(\frac{1}{N}\right)^2 + N \left(\frac{\delta N_{\uparrow\uparrow} - \delta N_{\uparrow\downarrow}}{N^2}\right)^2\right]. \tag{A.10}$$

As the change in neutron counts is expected to be small, and N large, the second term in the bracket can be ignored giving

$$\sigma^2(d_n) = \left(\frac{\hbar}{2E\alpha T}\right)^2 \frac{1}{N}, \tag{A.11}$$

or

$$\sigma(d_n) = \frac{\hbar}{2E\alpha T\sqrt{N}}. \tag{A.12}$$

Appendix B

Calibration

B.1 Extracting a calibration pulse from noisy data

The magnetometer output during the calibration signal will be the sum of the signal and background field changes

$$V(t) = S(t) + B(t), \tag{B.1}$$

where $V(t)$ is the raw magnetometer output, $S(t)$ is the response to the applied signal, and $B(t)$ is the background field. This can be integrated and averaged, giving

$$\frac{1}{t_1 - t_0}\int_{t_0}^{t_1} V(t)dt = \frac{1}{t_1 - t_0}\int_{t_0}^{t_1} S(t)dt + \frac{1}{t_1 - t_0}\int_{t_0}^{t_1} B(t)dt, \tag{B.2}$$

where t_0 and t_1 delimit the region of interest during which we expect a calibration pulse.

Assuming that the background is stable, and the region of interest is short, the background can be approximated as a linear function of time, giving

$$B(t) = \frac{dV}{dt}t + V_0 \tag{B.3}$$

$$\frac{1}{t_1 - t_0}\int_{t_0}^{t_1} B(t)dt = \frac{1}{t_1 - t_0}\left[\frac{1}{2}\frac{dV}{dt}(t_1^2 - t_0^2) + V_0(t_1 - t_0)\right] \tag{B.4}$$

$$= \frac{1}{2}\frac{dV}{dt}(t_1 + t_0) + V_0, \tag{B.5}$$

where $\frac{dV}{dt}$ is the rate of change of the background field response and is assumed to be constant in time.

The value of the background during a calibration pulse can be estimated from the value of the background during times when the calibration signal will be zero, provided this period is close enough to the signal that the background has not changed significantly. If we take the average of the output during such a period, for example from a time δt before the region of interest until the region of interest we get

$$\begin{aligned} V(t) &= B(t) && \text{(B.6)} \\ \langle V \rangle_{\text{before}} &= \frac{1}{\delta t} \int_{t_0-\delta t}^{t_0} V(t) dt && \text{(B.7)} \\ &= \frac{1}{2} \frac{dV}{dt} (2t_0 - \delta t) + V_0 && \text{(B.8)} \end{aligned}$$

Similarly taking the average until a time δt after the region of interest we get

$$\begin{aligned} \langle V \rangle_{\text{after}} &= \frac{1}{\delta t} \int_{t_1}^{t_1+\delta t} V(t) dt && \text{(B.9)} \\ &= \frac{1}{2} \frac{dV}{dt} (2t_1 + \delta t) + V_0 && \text{(B.10)} \end{aligned}$$

Summing $\langle V \rangle_{\text{before}}$ and $\langle V \rangle_{\text{after}}$ gives

$$\begin{aligned} \langle V \rangle_{\text{before}} + \langle V \rangle_{\text{after}} &= \frac{1}{2} \frac{dV}{dt} (2t_0 - \delta t) + V_0 + \frac{1}{2} \frac{dV}{dt} (2t_1 + \delta t) + V_0 && \text{(B.11)} \\ &= \frac{1}{2} \frac{dV}{dt} (2t_0 + 2t_1) + 2V_0 && \text{(B.12)} \\ &= 2 \frac{1}{t_1 - t_0} \int_{t_0}^{t_1} B(t) dt && \text{(B.13)} \end{aligned}$$

We can then compute a value for the integral of the signal by subtracting off the back-

ground contribution from the magnetometer output

$$\int_{t_0}^{t_1} S(t)dt = (t_1 - t_0)\left(\langle V\rangle - \frac{\langle V\rangle_{\text{before}} + \langle V\rangle_{\text{after}}}{2}\right). \tag{B.14}$$

Appendix C

Material investigation

C.1 Calibration of Cernox thermometer

To calibrate the thermometer used in the measurements to determine the transition temperature of various materials, a set of samples with well known transition temperatures was needed. These samples had their transition temperatures determined in the apparatus in terms of thermometer resistance. These resistances could then be compared to the literature value of the transition temperature. Table C.1 shows this comparison for the superconductors Titanium, Molybdenum, Aluminium, Tin, Tantalum, and Lead. Also shown in the table are the resistances at other known temperatures, the boiling point of liquid Helium, the boiling point of liquid Nitrogen, and room temperature.

The functional form of the resistance vs temperature curve in the low temperature regions (< 10 K) is

$$\ln R = a(\ln T)^2 + b \ln T + c \tag{C.1}$$

where R is the thermometer resistance, T is the temperature of the device. The data in table C.1 was used to determine the values of the parameters a, b and c. The best fit values for these constants, and their associated uncertainty are shown in table C.2.

Sample	Resistance / Ω	Temperature / K
Titanium	899.8206	0.39
Molybdenum	507.92895	0.92
Aluminium	427.3158	1.2
Tin	257.62911	3.72
Tantalum	243.90093	4.48
Lead	208.6227	7.2
Liquid Helium	248	4.2
Liquid Nitrogen	94.6	77
Room Temperature	50.1	298

Table C.1: Table showing the calibration points used in determining the calibration of the Cernox thermometer used to measure the superconducting transition temperatures in the AC susceptometer.

Term	Value	Uncertainty
a	0.07468	0.01104
b	-0.57842	0.01025
c	6.19094	0.01227

Table C.2: Table showing the calibration factors for each term in equation C.1.

Appendix D

Software

In addition to removing false EDM signals from the final EDM analysis, the magnetometry system is to be used as a component in a system keeping the magnetic field at the experimental operating point. The former use can be done offline and only requires data files to be written while running, but the latter requires a full real-time analysis to be performed on the data as it is taken. Additionally, the magnetometry data must be made available to the other components in the field stabilisation system. This is done through communications over a network, with a master DAQ server issuing commands to the magnetometry system, and requesting data from it.

In order to have a working real-time analysis, a significant software framework had to be developed, this included data taking, data recording, signal processing, configuration, user display, etc. The following is a description of some important aspects of the framework that was developed.

D.1 Data analysis chain

The major missing component of the system was a data analysis chain. A network based command interface was already in place, as were device drivers to read ADC values, but nothing connecting the two systems. A large quantity of code was therefore needed to turn

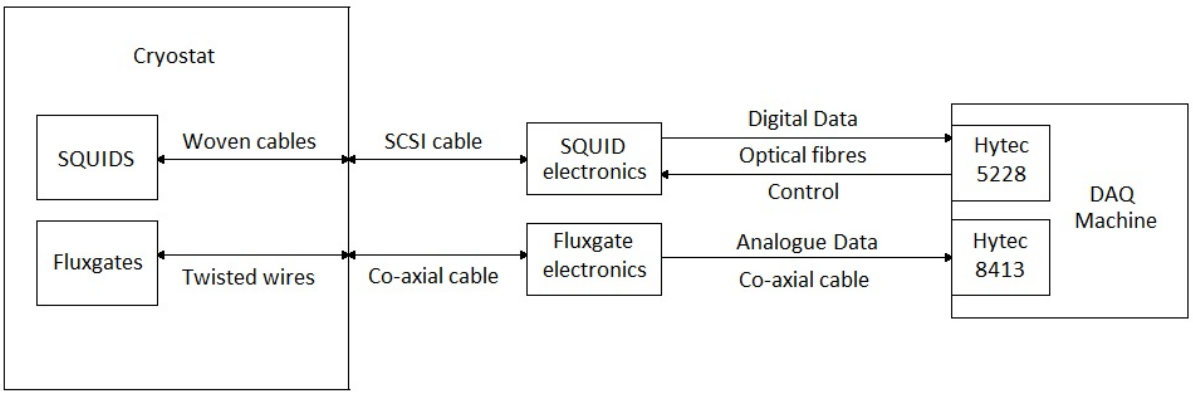

Figure D.1: Shown is a schematic overview of the electronics needed to acquire data from the SQUIDs and fluxgates. The SQUID data are digitised by the SQUID electronics as they exit the cryostat. The fluxgates are digitised on-board the DAQ machine.

the ADC outputs provided by the DAQ cards into the meaningful information required by the experiment upon server requests. It is the job of the data analysis chain to marshal the data from the DAQ cards, through the various subsystems developed to meet these requests. The framework can be broken down into two sections, a DAQ chain which is responsible for producing data useful in the analysis, and an analysis chain which analyses the data in response to commands from the master DAQ server.

Shown in figure D.1 is a schematic diagram of the read-out electronics for the magnetometry system on which the data analysis chain operates.

D.1.1 DAQ Chain

At a fundamental level the magnetometry is based on the outputs of the SQUIDs and fluxgates. Therefore the first step in the data chain is to acquire the raw voltage from the magnetometers. The SQUIDs and fluxgates have different sets of electronics, and operate with different bandwidths. The data for the different types of magnetometer are then taken using different DAQ systems. In software, these data are taken through various DAQ subsystems for the different types of DAQ used, but all must produce a floating-point representation of the voltage presented to the ADCs. Every DAQ system can then be treated

as identical further along the chain.

The DAQ subsystems have the responsibility of assessing the integrity of the data being collected, and fixing any DAQ errors if they arise, for example restarting data taking if the electronics freeze. Since there are multiple DAQs being operated for the various magnetometers and other equipment running at different data rates, these individual streams need to be collected, synchronised, and interleaved correctly before they can contribute to the analysis. This interleaving is done by the analysis chain described below. Once the data has been acquired a set of DAQ flags is set for each channel to indicate any issues on the channel.

After acquisition the data must be recorded, this is done to an arbitrary number of files in various formats e.g. the raw data storage format detailed in section D.2, and at various points in the DAQ chain e.g. raw voltage, or calibrated data. Each file provides an identical interface to the DAQ chain, requiring only a time stamp and a set of data, hiding the details of the format from the DAQ.

The raw data is not guaranteed to be directly relatable to the magnetic field. As described in chapter 4, there are a number of hardware artefacts that can alter the output. At this stage of the DAQ chain, the SQUID data is passed through a correction subsystem which implements one of the algorithms detailed in section 4.1.1.2. The returned data has any artefacts removed, and so is now a representation of the magnetic field. The correction subsystem sets a correction flag for each channel to indicate any corrections performed.

Once any corrections are performed the data streams represent the voltage produced by the magnetometers as a linear function of observed flux. The data are then converted to tesla using a simple linear gradient and offset, with the values of the SQUID calibration as obtained through methods detailed in chapter 3.2. After calibration the data can be filtered, as required, using a suite of digital filters as described in section E.1.

Finally, using the DAQ flags, correction flags, raw data and calibrated data, a data quality subsystem assesses the quality of the data. A quality flag is then set for each channel indicating any problems that may require exclusion from further analysis. This is the last

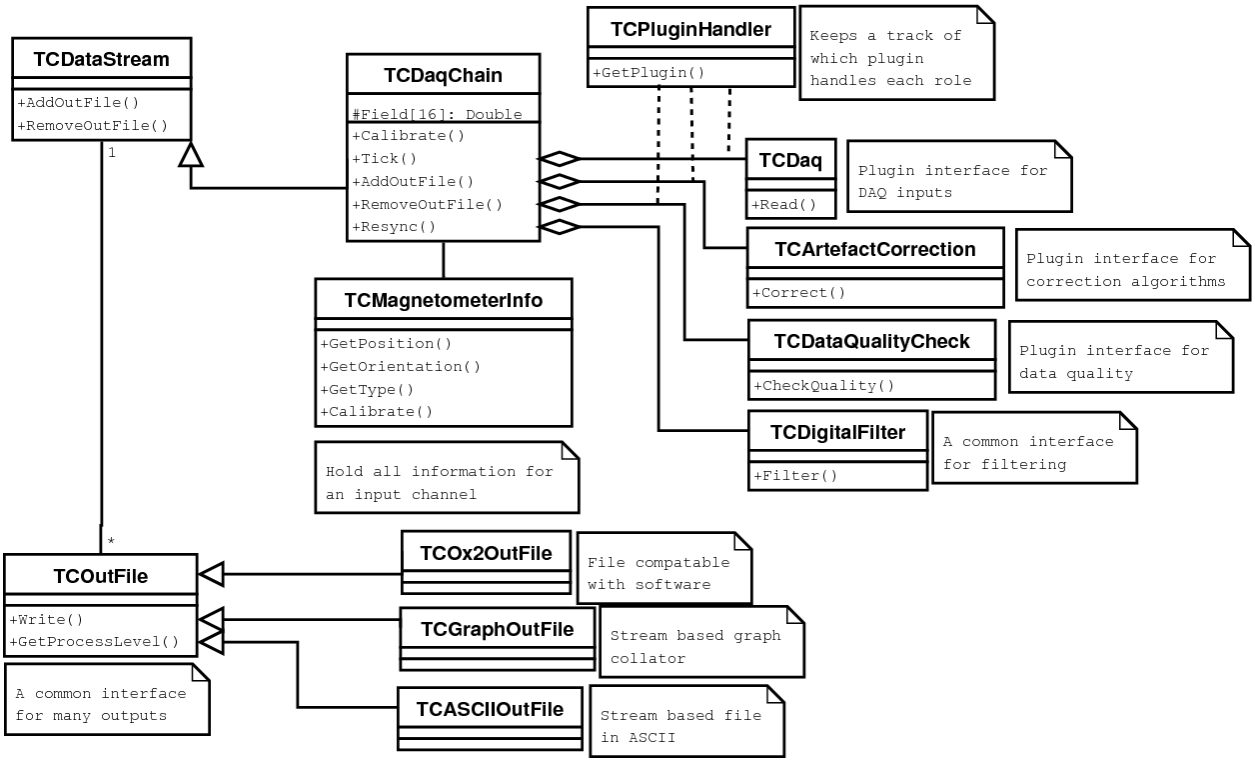

Figure D.2: Shown is the representation of the DAQ section of the data analysis chain. This shows the class TCDaqChain, which is responsible for producing a stream of meaningful magnetic data to the higher level subsystems of the analysis. Here the data is taken by TCDaq, corrected for artefacts by TCArtefactCorrection, calibrated by TCMagnetometerInfo, checked for data quality by TCDataQualityCheck, filtered by TCDigitalFilter and recorded to files implementing TCOutFile by TCDataStream.

action of the DAQ chain and the data should be ready for analysis. A UML diagram showing the structure of the DAQ chain can be seen in figure D.2.

D.1.2 Analysis Chain

The analysis chain consists of a set of DAQ chains, one for each DAQ system required. The analysis chain runs in two threads, one handling server requests and one controlling the DAQ chains. The DAQ thread is tasked with synchronising and interleaving the data from the various DAQs. This is done in a loop that ticks a master DAQ system, the fastest DAQ used, and uses this as a reference for the other DAQs. With knowledge of the clock rates, whether or not a DAQ should be providing a sample in the current iteration is calculated. If it should, the most recent sample is requested, otherwise the DAQ is ignored. This constructs

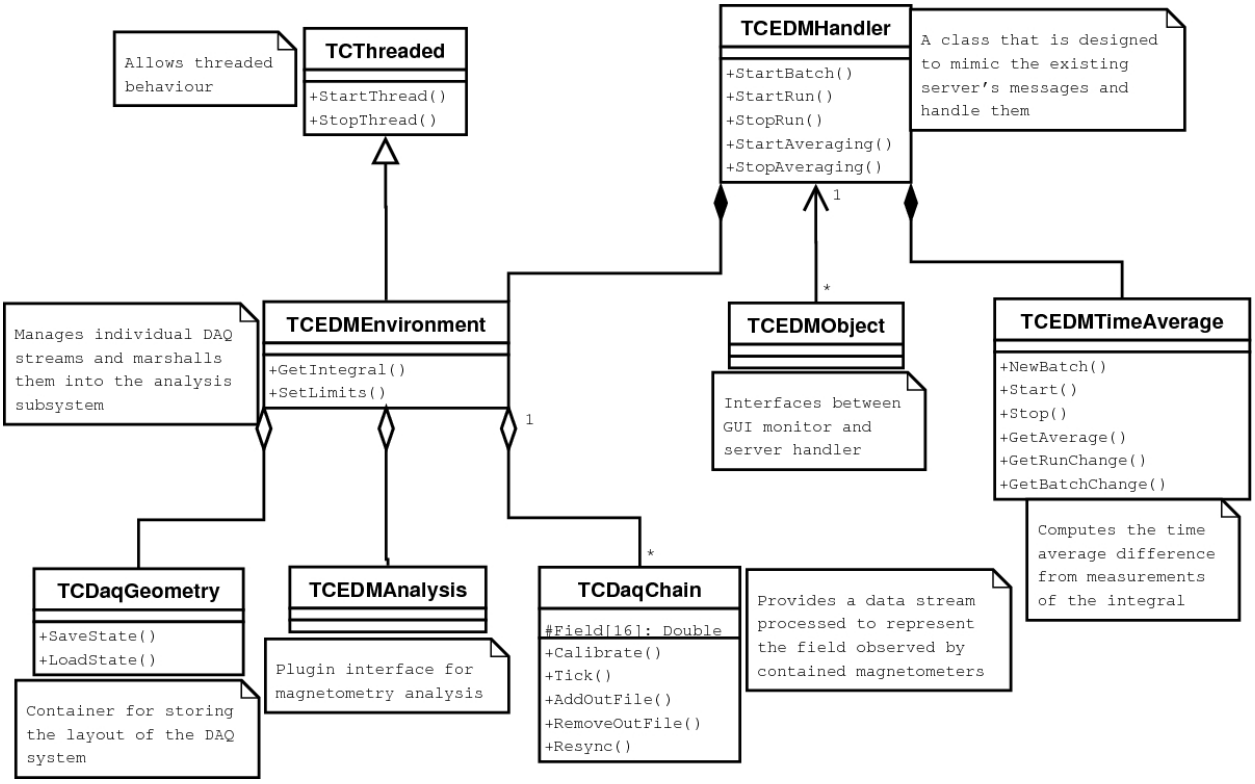

Figure D.3: Shown is the representation of the processing section of the data analysis chain. This shows the class TCEDMHandler that is responsible for presenting meaningful magnetic data to the rest of the experiment. It has connections to the monitoring GUI via TCEDMObject. It computes the required quantities from the data analysed by TCEDMEnvironment, which manages and synchronises the data streams from many TCDaqChains, and presents the data to be analysed by TCEDMAnalysis.

a set of the most recent data for each magnetometer. The timing comes from the master DAQ system, and so this must be able to clock itself. This constraint is not necessary for subordinate DAQs as the period between ticks will be very close to the correct sample spacing, but it is preferable. This approach limits the clock rates of the subordinate DAQs to be an integer multiple of the master DAQ sample spacing, but this causes no significant problems.

Once the DAQ section of the data chain is complete, the data are then ready for analysis, which is done via an analysis subsystem. For each tick of the master DAQ system the most recent set of data is passed to a subsystem, which implements the the algorithms as detailed in chapter 3. This subsystem is also passed the data quality flags from the DAQ chains, so that channels can be omitted if necessary. Returned is the analysed magnetic field and any

parameters calculated in the determination of the field, e.g the field gradient and DC value.

Once the data stream has been analysed, it is ready for formatting as requested by the rest of the experiment. This, in general, is the time average of the data sets between averaging commands received from the master DAQ server. This provides the required functionality to operate a real-time analysis. A UML diagram showing the structure analysis and formatting software can be seen in figure D.3.

D.2 Data file format

Before 2008 the SQUID DAQ hardware was event based, an array of values was recorded from the ADCs using an on-board clock, initiated by a software trigger linked to a simple timer. This left dead time between the events where the magnetic field was not being tracked. In 2008 the DAQ hardware was upgraded for continuous data taking, removing the dead time but causing problems for the data files, which were merely a representation on disk of the returned arrays, and were not suitable for continuous data. A new format for storing the data taken was needed, with the aim of handling different kinds of data from different DAQ systems, and to encapsulate DAQ, SQUID and other metadata with the data, which was not possible in the existing format.

The file format was intended to replace the existing format in other areas as well, so needed to accept both event based and continuous data. It was decided to address all of the above problems at once, and design a new file format independent of the experimental particulars that could handle both continuous and event based data taking, but reusing elements of the original data structures, where possible, to minimise the effort to produce reading/writing software.

To this end a new data file format was created with a simple philosophy, all internal data structures should conform to fixed specification as far as the file structure is concerned. Inspiration for this was taken from various audio formats [116][117] that must deal with a similar problem, but did not suit the needs of the experiment. This means that all data and

metadata must be placed inside a container that has a common and well-defined specification. Namely,

```
<Container> = <Identifier><Version><Content><Length>
```

where the identifier is a 16-byte field containing a string to identify the container type, version is a 4-byte integer to identify any container sub-type, the content is arbitrary and depends on the container type, and the length field is an 8-byte integer specifying the total length of the container.

This enables records and metadata to be identified within the file, with an appropriate version number for future changes to the data structures. The length field allows for simple fault checking of the container, and some recovery in case of corruption. It is placed after the content to act as a delimiter. Containers that are not recognised by the reading software can be ignored by the parser, which can then skip ahead until a recognisable container is found. This allows for extra containers to be added as a system progresses, without breaking forwards or backwards compatibility of the data files.

Having simply a collection of these containers alone make up a file would be functional in principle, however it would make certain tasks slow. To solve this an overall structure is applied to the file,

```
<File> = <Header><Version><Container>...<Container><Properties><Directory>
```

where the header is an 80-byte field containing a single string, the version is a 4-byte integer and all other content are containers, as defined above.

The header identifies the file type, and the version number makes future structural changes possible. The bulk of the file would then contain a collection of arbitrary containers, but to aid file reading, two specially designated containers are specified to follow the bulk: the properties container; and the directory container. The properties field contains information that might be useful for the file parser, e.g number of entries, but no information

that cannot be reconstructed from the previous sections in the file. The directory contains the position of the first container in a section e.g DAQ metadata, detector metadata, data etc, again containing information already available in the previous sections of the file. The directory container is required to be the last container in the file. In all implementations so far, the order of other containers is loosely specified to make parsing the file quicker, however this is not an explicit requirement of the format.

Implicit structures in the old file format were made explicit, for example in several analyses the concept of an event emerges consisting of coincident records. Events were formalised in the file format as coincident records, produced by the same DAQ system. It was then possible to move timing information into the event header using two 64-bit integers to record the "real"-time defined by the recording CPU, and the DAQ-time defined by the DAQ system's clock, both with microsecond precision, and at a net space-saving in typical cases. A per event header gave good opportunity for the length of the records within to be recorded, allowing arbitrary record lengths. The header also contained an ASCII metadata section used to define arbitrary metadata, if needed. However, the format of the ASCII metadata was specified to be similar to that used in the FITS standard [118] already used in some experiments.

D.2.1 Example specification

To illustrate the above design, a short example file will be given containing 2048 samples, split into two events, on a single channel. The header string used in this example file is "Ox2-Daq Std".

```
<"Ox2-Daq Std" 69 bytes 0x00>
<0x00000001>
<"DaqInfo" + 9 bytes 0x00>
<0x00000001>
<"Number of DaqInfos = 2
```

```
 >> Daq 0 <<
 Cutoff Frequency [Hz] =  1000
 Timebase [us] =          1000
 Digitizer minimum [V] =  -10
 Digitizer maximum [V] =  10
 Digitizer bits =          16
--------------------------------\n">
<0x00000000000000fa>
<"Event" + 11 bytes 0x00>
<0x00000001>
<0x0000000000000018>
<0x0000000000000000>
<0x0000000123456789>
<0x0000000000000001>
<"">
<"Pulse" + 11 bytes 0x00>
<0x00000001>
<0x0000000000000018>
<0x0000000000000400>
<0x0000000000000000>
<0x0000000000000000>
<"">
<1024 16-bit words data>
<0x000000000000083c>
<0x0000000000000878>
<"Event" + 11 bytes 0x00>
<0x00000001>
<0x0000000000000018>
<0x000000003d090000>
```

```
<0x0000000123551789>
<0x0000000000000001>
<"">
<"Pulse" + 11 bytes 0x00>
<0x00000001>
<0x0000000000000018>
<0x0000000000000400>
<0x0000000000000000>
<0x0000000000000000>
<"">
<1024 16-bit words data>
<0x000000000000083c>
<0x0000000000000878>
<"Statistics" + 6 bytes 0x00>
<0x00000001>
<"Number of Statistics = 2
 >> Statistics 0 <<
 Channel Number = 0
 Start Time [s us] = 1234 123456
 Stop Time [s us] = 1236 171456
 Number of Events = 2
--------------------------------\n">
<0x00000000000000d8>
<"File" + 12 bytes 0x00>
<0x00000001>
<0x0000000000000050>
<0x000000000000014a>
<0x000000000000123a>
<0x0000000000001312>
```

<0x000000000000003c>

D.2.2 Summary

In table D.1 the statistics of the new data format are shown. In general the format allows many new and useful features, but this comes at the price of additional overhead in forming the data structures. However, this additional overhead is far from limiting any DAQ system conceived to be used with it, and so does not constitute a problem.

Parameters	Value	Comment
Encapsulates meta data	Yes	
Maximum file size	16 exabytes	Enough for 100,000 years running at maximum data rate, for a full magnetometry system.
Typical size overhead	6 %	Higher than the 2 % of the replaced format.
Arbitrary record size	Yes	
Accepts ancillary data	Yes	Can include DAQ errors or flags, etc. with the data streams.

Table D.1: Statistics of the new data file format, and comments upon them in reference to the improvements or penalties of the format.

D.3 Plug-in system

D.3.1 Overview

As the experiment is still in the commissioning phase, a lot of development work is still occurring. New data acquisition systems are in active development, and the requirements of the system relating to the rest of the experiment often evolving. To handle this ever changing environment, without the need for constant redevelopment of the software, a plug-in framework was designed. This would allow a specific task required by the software to

be defined in a fixed fashion, but leave the implementation undecided at compile time. To accomplish this, a virtual specification was created for each task that was deemed 'plugable'. This served as the access point for interaction with the software. A separate library for each plug-in would then be created, containing an object that inherited from the virtual specification, and a number of functions used to register the plug-in. A plug-in manager would then handle the registration of each plug-in, checking compatibility with the current software version, and providing an instance of the object when requested. The correct library would then be linked into the software when the object is created giving a valid implementation.

Each of the subsystems described in the data analysis chain above are plug-ins. The DAQ systems, correction algorithms, data quality checking, and analysis algorithms can be updated, expanded, interchanged, or removed as needed without affecting the overall software framework. In addition to allowing flexibility, this introduced convenient ways of testing the software, for example if a DAQ plug-in was created that simply returned values designed to give expected results, all aspects of the system could be checked offline to see if these results are reproduced.

D.3.2 Registration

The plug-in manager needs to know certain information about the plug-in to function: general information about the plug-in e.g its name, a description, its type, etc required for the database entry; the version of software the plug-in was compiled against, used to ensure compatibility with the current version; and a method for creating an instance of an object, a factory function.

These are provided by global and predefined functions bundled with the plug-in. The functions are found in the library by searching for the names of their symbols. In C++, to accommodate functions with identical names but different signatures, a process of name mangling occurs [119]. This gives the symbol defining the location of each function a unique name. However, the name mangling process is not defined in the standard, thus the symbol

name given is compiler specific [120], which means it would be difficult to determine the name of a symbol, in advance, in a cross-compiler friendly way. The solution is to explicitly declare the functions to be made available as C symbols. This removes name mangling, but introduces a new issue: there can not be a suite of functions with different signatures to handle object creation via different constructors. The immediate solution to this is to pass a wrapper object to the factory function that details which constructor should be used, and its arguments.

D.3.3 Extensibility

An interesting side effect of this system is that it allows not only a simple way of accommodating changes, but enables extensions of functionality. By defining a class that operates on a data stream and returns a modified data stream, a complicated analysis can be reduced to a pipeline of simple tasks performed by plug-ins. The pipelines can be modified easily and extended by the creation of new plug-ins for new operations. A good example of this is the script system designed for the data monitoring software detailed in section D.5.1.

D.4 Settings system

The system is designed to interact with the rest of the data acquisition by receiving commands through the network, and therefore requires a large amount of automation. To accomplish this, a robust configuration system was needed to manage settings for multiple DAQ systems, positions, orientations and calibrations of magnetometers, along with signal processing, data recording, and data analysis options. Nearly 500 items need to be handled by the settings system in total.

The format chosen to represent the settings for storage on disk is an XML based format. This neatly captures the structure of the tree, and allows for both reliable software based access and for editing by hand. The ability to edit by hand is useful for certain cases where

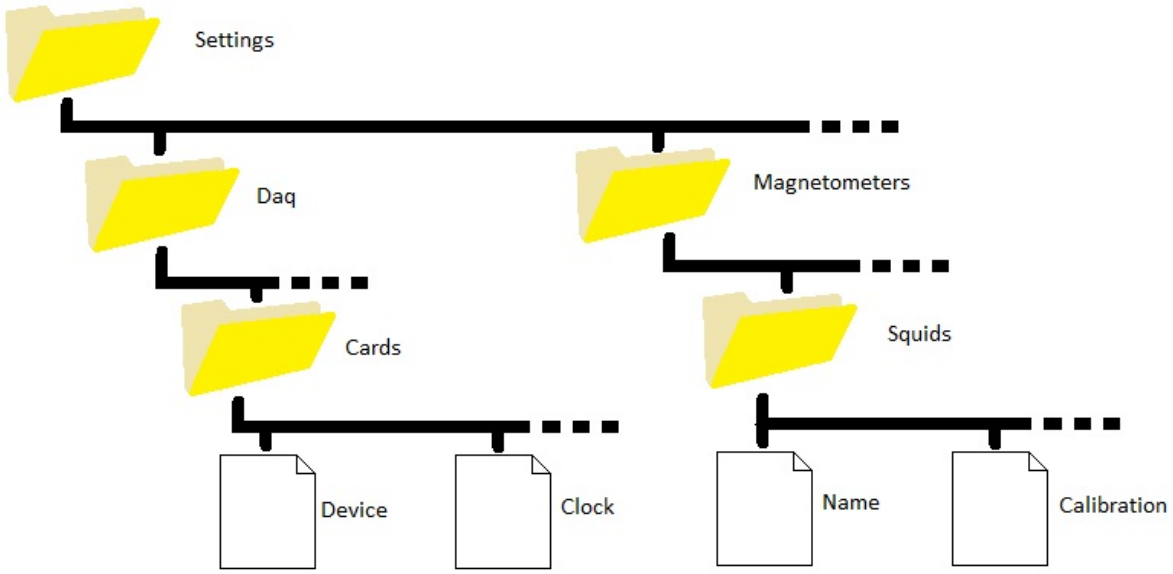

Figure D.4: Shown is an example of how the settings are represented internally. Directories group related settings, and individual 'files' contain the value of the settings.

dangerous specialist settings need to be enabled that the GUI configuration program does not offer to the user for safety reasons.

Internally the settings are represented by a virtual file system, accessed through typical directory structure methods, each "file" containing the value of the setting. Settings can be accessed by absolute or relative paths, and the directory tree traversed by changing working directories or through iterators. Figure D.4 shows a visual representation of the internal settings structure.

This structure is represented in XML by using the tag

```
<dir name="dir_name">
```

to define the directory tree, and

```
<key name="key_name" value="key_value">
```

to define the settings and their values.

The settings are stored as ASCII strings no matter the data type of the value. The settings

can be converted between strings and the required data types automatically through the use of templated functions that make use of stream operators.

D.5 Monitoring software

To complement the data analysis software detailed above, a graphical environment in which to monitor the magnetic environment remotely, and in real time was requested. As the data-taking framework allows multiple data outputs to be attached to a data stream, this seemed the logical place to extract portions of the data to transmit to a client. Both raw and corrected/calibrated data are available in the client program.

Multiple clients are supported, up to a user set limit. Clients can choose arbitrary length data sets, again up to a user set limit, and can reduce the data rate by any integer factor. Once acquired, the user can display the data set in various dynamically created tabs with various display options. Data sets may be stored for future reference, or combined to produce new data sets via a scripting system (detailed in section D.5.1). An example of the monitoring software in use is shown in figure D.5.

In addition to the passive features, the client program may also take the place of the server requests from the rest of the experiment. It can start and stop data taking, emulate data runs and neutron batches, and perform the automated calibrations described in section 3.3. These actions can be dangerous and are only permitted if the client has been granted administrative privileges. The administration mode can only be accessed if enabled by the local configuration, and upon presentation of the correct password, which is hashed locally as an MD5 check sum [121]. Further security measures are not seen as necessary at this time, as access to the local network is required to counteract these measures, and in such cases direct damage could be performed anyway.

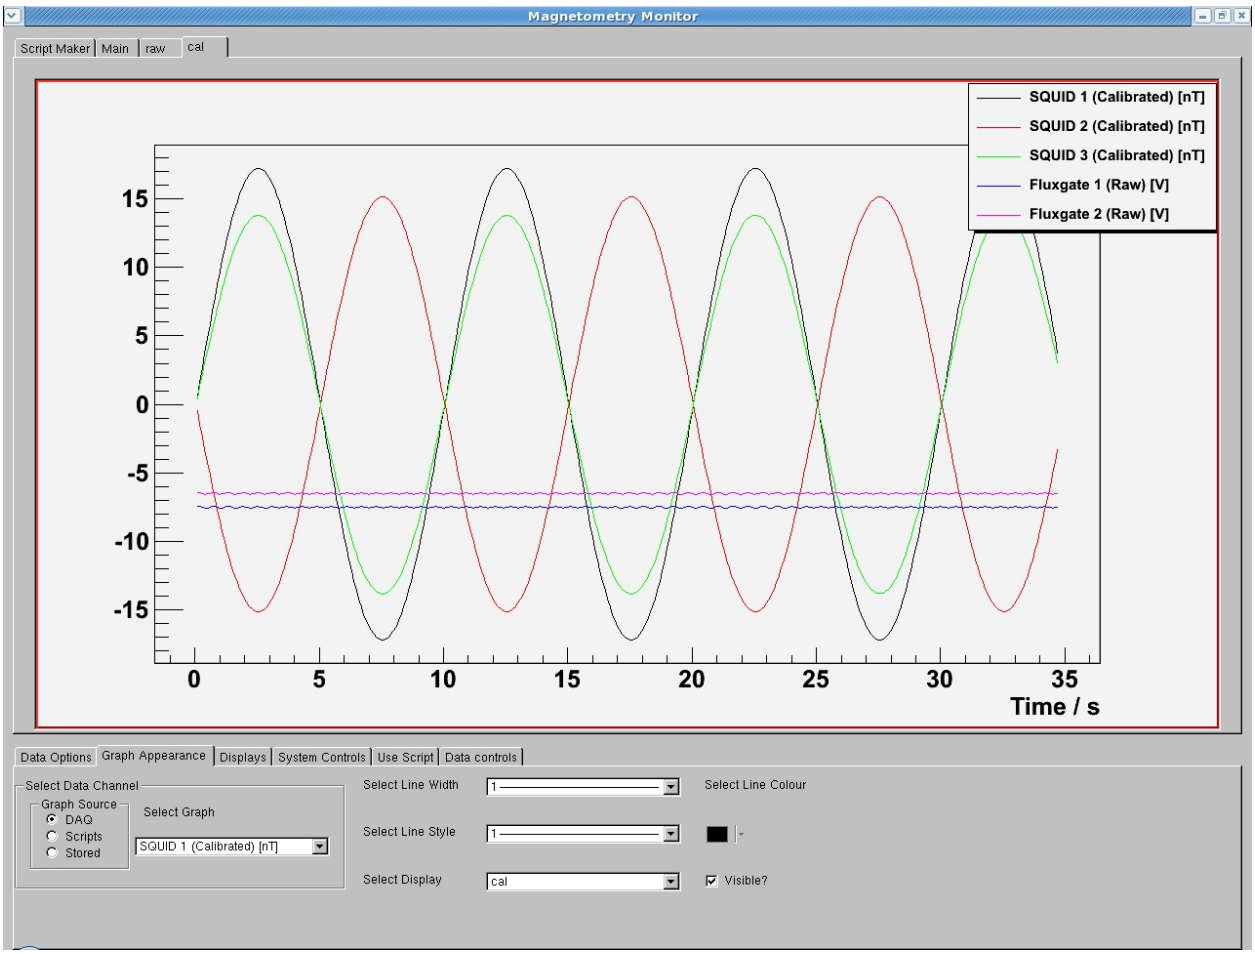

Figure D.5: Shown is a screen shot of the magnetometry monitoring software displaying the corrected and calibrated output of three example SQUIDs, and the raw output of two example fluxgates at the same time. The display properties tab at the bottom of the figure is open to allow display changes. Other tabs under the drawing region allow access to other functionality, e.g. changing incoming data rates or sending commands to the server.

D.5.1 Script system

Various collaborators have often requested the ability to observe certain parameters of the magnetic environment in real time for data quality checking. Many of these parameters are not directly available from the measurements, but are simple manipulations of them, for example DC shielding factors. To this end, a scripting system was developed to allow a user to perform simple analysis of the magnetic environment in real time. This was created using the plug-in system to provide operators that are defined to take a number of data sets and return a new data set. These operators could be combined into any expression in reverse-polish notation.

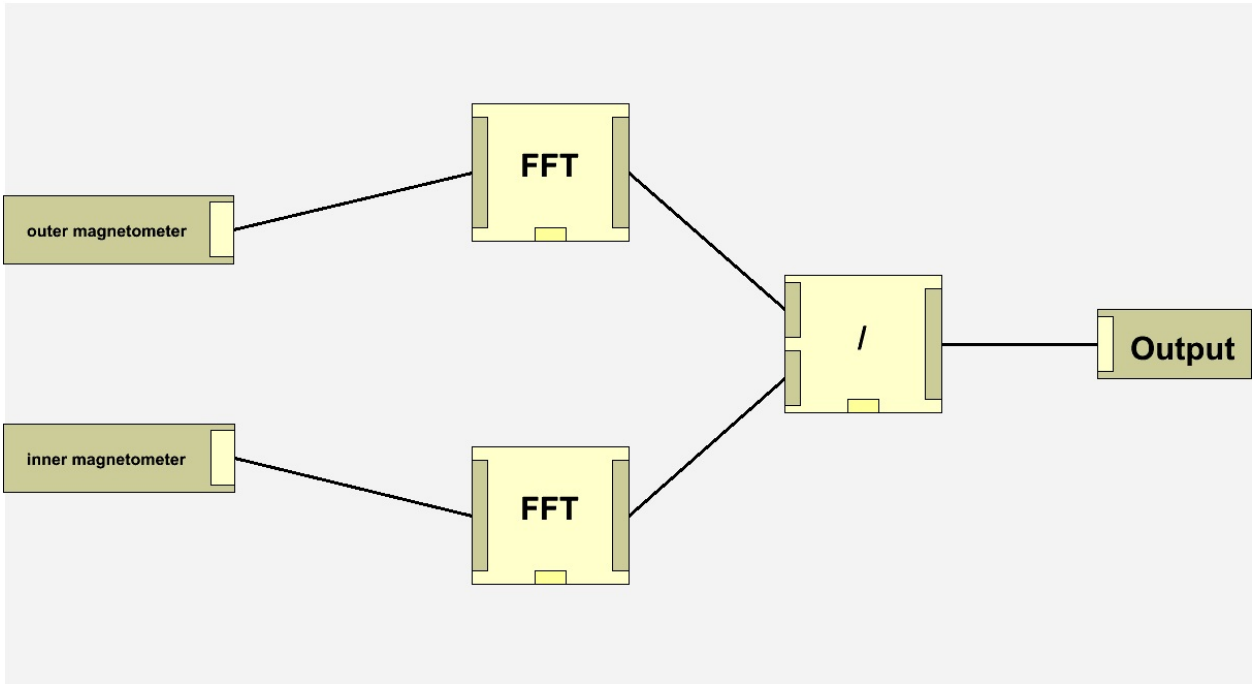

Figure D.6: An example of the graphical representation of an analysis script as created by the script editor in the monitoring software.

To make the scripts easy to write, a graphical user interface was constructed which allowed expressions to be created in a block diagram style. An example of such a script can be seen in figure D.6 to display the ratio of two Fourier transforms on input waveforms. Inputs to the scripts can be raw or calibrated data, previously stored data, other scripts, constant values, or the previous version of the scripts output, allowing a large range of analysis to be performed. This is useful as a large number of diagnostics are required while the experiment is running, e.g. stability or shielding factors, and this system provides a way to easily access this information in real time.

There are a number of built in operators including: basic arithmetic; parameter calculations e.g. averages; comparisons; conditionals; concatenation; application of c-style functions to data points; and axis manipulation. As the script system is based on the plug-in system, a user may create their own operators, or use operators which may not be appropriate on all systems, with which to extend the basic functionality, e.g. a Fourier transform operator. This makes the script system very powerful.

The script system was used to great effect during the susceptibility measurements de-

scribed in chapter 5. The script system was used to plot the calibrated superconducting transition curves of the samples as the data was being taken, and to monitor various parameters. This allowed a rapid check of the data quality, meaning experimental issues could be spotted and corrected before the sample was removed. The system will play a similar role in the main experiment.

Appendix E

Signal processing

E.1 Digital filters

The recorded and analysed data streams may not require the full amount of data required by the artefact correction algorithms. Before the data rate can be reduced, it must be filtered to remove any high frequency components above the Nyquist frequency [122] of the new sampling rate to avoid aliasing effects. To achieve this a suite of digital filters was created.

The implementation of digital filters is rather simple, just a sum of the product of a coefficient and an input [123]. They fall into two broad categories depending on the types of inputs required;

Finite impulse response (FIR) are the simplest type of digital filters, and require no feedback, using only the raw data as inputs. They are inherently stable but at the price of performance. FIR filters have the following implementation:

$$y_k = \sum_{n=0}^{k-1} a_n x_{k-n}, \tag{E.1}$$

where y_k is the output of the filter at time k, x_k the input at time k, and a_n coefficients that determine the filter response.

Infinite impulse response (IIR) require feedback of the previous filter outputs in addition to the raw data as inputs, they can produce sharper cut-off frequencies at the cost of stability. IIR filters have the following implementation:

$$y_k = \sum_{n=0}^{k-1} a_n x_{k-n} - \sum_{n=1}^{k-1} b_n y_{k-n}, \tag{E.2}$$

where y_k is the output of the filter at time k, x_k the input at time k, and a_n and b_n coefficients that determine the filter response.

The calculation of the coefficients is the more interesting problem. The coefficients are the time domain representation of the desired frequency response, which for the FIR filters is simply the Fourier transform of a rectangular function with edges at the cut-off frequency, $\pm f_0$. This gives a sinc function. The problem is complicated slightly by the fact that in practice there is a limit to the number of data points that can be used to calculate a new output point. This effectively means the sinc function is truncated. To accomplish this a windowing function is multiplied on to the sinc function, clearly leading the frequency response to be the convolution of the rectangular function with the Fourier transform of the windowing function. For the trivial windowing function, itself a rectangular function, this is a sinc in the frequency domain, causing the response to oscillate. Various windowing functions with particular properties have been previously investigated and tabulated [124]. The coefficients $a_n = 2f_0 \mathrm{W}(t_n)\mathrm{sinc}(2f_0 t_n)$, where t_n is the relative time of the point associated with the coefficient, and W(t) is the windowing function, are thus readily calculable for any order and any cut-off provided the window is known.

IIR filters are generally based on physical filters, and thus the coefficients are given by the solutions of the differential equations governing their behaviour. These equations are often solved in the frequency domain using a Laplace transformation, and used to produce to a transfer function, the ratio of Laplace transformed outputs and inputs $\mathrm{H(s)} = \frac{\mathrm{Y}(s)}{\mathrm{X(s)}}$, where s is the complex angular frequency. Two popular IIR filters are based on the Bessel [125] and Butterworth [126] polynomials, with the transfer functions for a particular order of filter

derived from that order of polynomial. These polynomials are defined elsewhere [127][128] and easily computed, and so, therefore, are the transfer functions for arbitrary filter orders in powers of s.

For a Butterworth filter the transfer function is

$$\mathrm{H(s)} = \frac{1}{B_n(s/\omega_c)}, \tag{E.3}$$

where $B_n(s)$ is the nth normalised Butterworth polynomial given by

$$B_n(s) = \begin{cases} \prod_{k=1}^{n/2} \left(s^2 - 2s\cos\left(\frac{2k+n-1}{2n}\pi\right) + 1 \right), \text{ for even } n \\ (s+1)\prod_{k=1}^{\frac{n-1}{2}} \left(s^2 - 2s\cos\left(\frac{2k+n-1}{2n}\pi\right) + 1 \right), \text{ for odd } n \end{cases} \tag{E.4}$$

and ω_c is the cutoff frequency of the filter.

For a Bessel filter the transfer function is

$$\mathrm{H(s)} = \frac{\theta_n(0)}{\theta_n(s/\omega_c)}, \tag{E.5}$$

where $\theta_n(s)$ is the nth normalised Bessel polynomial given by

$$\theta_n(s) = \sum_{k=0}^{n} a_k s^k \tag{E.6}$$

where

$$a_k = \frac{(2n-k)!}{2^{n-k}k!(n-k)!} \tag{E.7}$$

and ω_c is the cutoff frequency of the filter.

E.1.1 Bilinear transform

Of course, these transfer functions are only valid for continuous time, to convert the continuous transfer function into a discrete transfer function a bilinear transform is used [129],

with

$$s \to \frac{2}{T}\frac{z-1}{z+1}, \tag{E.8}$$

where z is the discrete equivalent of complex frequency and T the sampling time, giving $\mathrm{H}_d(z) = \mathrm{H}(s)$. This needs to be rewritten in the form

$$\mathrm{H}_d(z) = \frac{\sum_n^k a_n z^{-n}}{\sum_n^k b_n z^{-n}} = \frac{\mathrm{Y}(z)}{\mathrm{X}(z)} \tag{E.9}$$

to give the coefficients used in equation E.2 (See appendix E.2 for details of why this is true). Unfortunately it is non-trivial to calculate this transform directly, so an algorithm was written that would compute the transform for an order k filter

$$\mathrm{H}(s) = \frac{\sum_n^k a'_n s^n}{\sum_n^k b'_n s^n} \to \mathrm{H}_d(z) = \frac{\sum_n^k a_n z^{-n}}{\sum_n^k b_n z^{-n}} \tag{E.10}$$

as follows (for brevity only showing the calculation of the numerator, the method to determine the denominator is identical),

$$\begin{aligned}
\mathrm{Num}\left[\mathrm{H}_d(z)\right] &= \sum_n^k a'_n \left(\frac{z-1}{z+1}\right)^n && \text{(E.11)}\\
&= \sum_n^k a'_n (z-1)^n (z+1)^{k-n} && \text{(E.12)}\\
&= \sum_n^k a'_n (z-1)^n (z+1)^{k-n} z^{-k} && \text{(E.13)}
\end{aligned}$$

where in equation E.12 both numerator and denominator are multiplied by $(z+1)^k$ and in in equation E.13 both numerator and denominator are multiplied by z^{-k}. Using the binomial

theorem

$$\text{Num}\left[\text{H}_d(z)\right] = \sum_n^k a'_n \left(\sum_m^n \begin{pmatrix} n \\ m \end{pmatrix} z^m (-1)^{n-m} \right) \left(\sum_l^{k-n} \begin{pmatrix} k-n \\ l \end{pmatrix} z^l \right) z^{-k} \quad \text{(E.14)}$$

$$= \sum_n^k \sum_m^n \sum_l^{k-n} (-1)^{n-m} a'_n \begin{pmatrix} n \\ m \end{pmatrix} \begin{pmatrix} k-n \\ l \end{pmatrix} z^{m+l-k} \quad \text{(E.15)}$$

The powers of z can then be summed and normalised such that $b_0 = 1$ to find the correct coefficients to be used in equation E.2.

E.1.2 Performance

Shown in figure E.1 is the frequency response of the digital filters to applied random Gaussian noise. A Bessel, a Butterworth and a FIR filter with 6 different windows with a cut-off frequency 0.1 (normalised so the sampling frequency is 1), all filters have the same order. Clearly visible is the oscillation of the rectangular window FIR filter. The sharpest cut-off is from the Butterworth filter, as expected. The choice of window has significant effects on the frequency response and so care in selection must be taken. Typical attenuation for signals in the stop band of the filters is seen to be between 10^3 and 10^4 depending on the filter type.

E.1.3 A note on numerical stability

Naive implementations of calculations to find the above coefficients can lead to numerical instabilities in calculating values due to terms such as $\frac{n!}{m!(n-m)!}\frac{(k-n)!}{l!(k-n-l)!}$ from the binomial coefficients. The calculation of which, if computed carelessly, will cause instabilities in calculations for order 11 filters due to the multiplication and division of very large integers to produce typically order unity coefficients.

To combat this, the algorithm implemented splits the calculations into sections of approximately order unity, producing better stability at the cost of computational speed, but

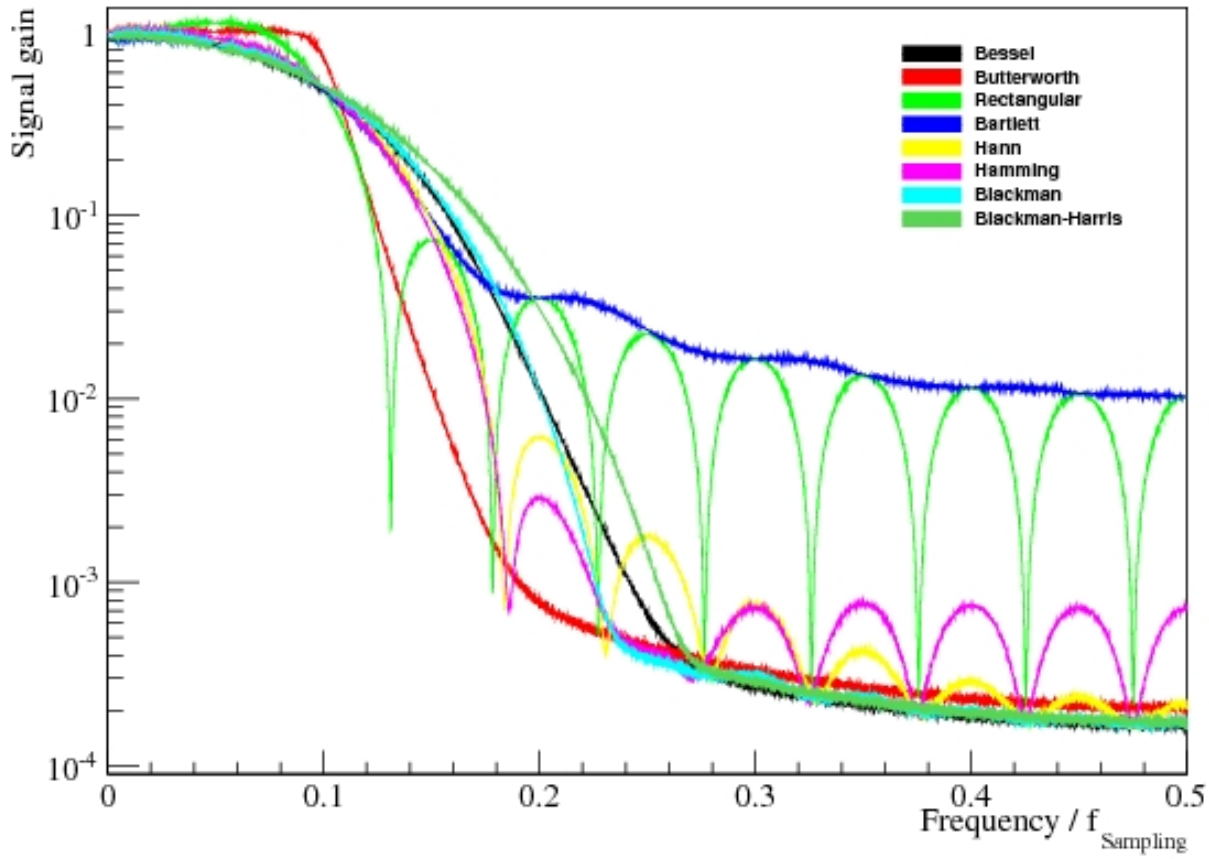

Figure E.1: Frequency response of the Bessel, Butterworth, and sinc filters with a cut-off frequency of 0.1 normalised to the sampling frequency. The response for a sinc with rectangular, Bartlett, Hann, Hamming, Blackman, and Blackman-Harris windows are shown.

as the coefficients are calculated only once for the lifetime of the filter this not a significant issue.

E.2 Digital Filters

Digital filters are a defined by a transfer function that maps input values $x[n]$, into output values $y[n]$ in the Z-domain.

$$\mathrm{H}(z) = \frac{\mathrm{Y}(z)}{\mathrm{X}(z)} \tag{E.16}$$

where the Z-transform is defined as

$$\begin{aligned} \mathrm{F}(z) &= \mathcal{Z}\{f[n]\} && \text{(E.17)} \\ &= \sum_{n=0}^{\infty} f[n] z^{-n} && \text{(E.18)} \end{aligned}$$

If the filter is causal (which it must be to act in real-time) the transfer function may be expressed as follows

$$\mathrm{H}(z) = \frac{\sum_{k=0}^{N-1} a_k z^{-k}}{\sum_{k=0}^{M-1} b_k z^{-k}} \qquad \text{(E.19)}$$

This leads to

$$\sum_{k=0}^{M-1} b_k z^{-k} \mathrm{Y}(z) = \sum_{k=0}^{N-1} a_k z^{-k} \mathrm{X}(z) \qquad \text{(E.20)}$$

Noting from equation E.18 a time delayed signal $f[n-k]$ would Z-transform as

$$\begin{aligned} \mathcal{Z}\{f[n-k]\} &= \sum_{n=0}^{\infty} f[n-k] z^{-n} && \text{(E.21)} \\ &= \sum_{j=-k}^{\infty} f[j] z^{-j-k} && \text{(E.22)} \\ &= \sum_{j=0}^{\infty} f[j] z^{-j} z^{-k}, \ (f[j] = 0 \text{ if } j < 0) && \text{(E.23)} \\ &= z^{-k} \sum_{j=0}^{\infty} f[j] z^{-j} && \text{(E.24)} \\ &= z^{-k} \mathrm{F}(z) && \text{(E.25)} \\ & && \text{(E.26)} \end{aligned}$$

Therefore the inverse Z-transform of equation E.20 will be

$$\mathcal{Z}^{-1}\{\sum_{k=0}^{M-1} b_k z^{-k} \mathrm{Y}(z)\} = \mathcal{Z}^{-1}\{\sum_{k=0}^{N-1} a_k z^{-k} \mathrm{X}(z)\} \tag{E.27}$$

$$\sum_{k=0}^{M-1} b_k \mathcal{Z}^{-1}\{z^{-k} \mathrm{Y}(z)\} = \sum_{k=0}^{N-1} a_k \mathcal{Z}^{-1}\{z^{-k} \mathrm{X}(z)\} \tag{E.28}$$

$$\sum_{k=0}^{M-1} b_k y[n-k] = \sum_{k=0}^{N-1} a_k x[n-k] \tag{E.29}$$

If this is normalised such that $b_0 = 1$, then we may write

$$y[n] = \sum_{k=0}^{N-1} a_k x[n-k] - \sum_{k=1}^{M-1} b_k y[n-k] \tag{E.30}$$

giving the equation of the next output as a function of previous inputs and output.

For an FIR filter where there is no feedback of previous outputs so the value of the feedback coefficients will be,

$$b_k = \begin{cases} 1, \, k = 0 \\ 0, \, k > 0 \end{cases} \tag{E.31}$$

Bibliography

[1] F. W. Stecker. *Nature*, **273**, 5663 (1978) 493–497.

[2] C. L. Bennett, M. Halpern, G. Hinshaw, N. Jarosik, M. Limon, J. Mather, S. S. Meyer, L. Page, D. N. Spergel, G. Tucker, D. T. Wilkinson, E. Wollack & E. L. Wright. "The Microwave Anisotropy Probe (MAP) Mission". In "American Astronomical Society Meeting Abstracts", volume 29 of *Bulletin of the American Astronomical Society*, pages 1353–+ (1997).

[3] D. N. Spergel et al. *The Astrophysical Journal Supplement Series*, **170**, 2 (2007) 377.

[4] A.D. Dolgov. "Baryogenesis, 30 years after" (1997). ArXiv:hep-ph/9707419v1.

[5] A. D. Sakharov. *Soviet Journal of Experimental and Theoretical Physics Letters*, **5** (1967) 24–27.

[6] G. 't Hooft. *Phys. Rev. Lett.*, **37**, 1 (1976) 8.

[7] A. A. Belavin, A. M. Polyakov, A. S. Schwartz & Yu. S. Tyupkin. *Physics Letters B*, **59**, 1 (1975) 85 – 87. ISSN 0370-2693.

[8] V. A. Kuzmin, V. A. Rubakov & M. E. Shaposhnikov. *Physics Letters B*, **155**, 1-2 (1985) 36 – 42. ISSN 0370-2693.

[9] C. S. Wu, E. Ambler, R. W. Hayward, D. D. Hoppes & R. P. Hudson. *Phys. Rev.*, **105**, 4 (1957) 1413–1415.

[10] T. D. Lee & C. N. Yang. *Phys. Rev.*, **104**, 1 (1956) 254–258.

[11] Maxim Pospelov & Adam Ritz. *Annals of Physics*, **318**, 1 (2005) 119 – 169. ISSN 0003-4916. Special Issue.

[12] Jihn E. Kim & Gianpaolo Carosi. *Rev. Mod. Phys.*, **82** (2010) 557–601.

[13] R. D. Peccei & Helen R. Quinn. *Phys. Rev. D*, **16** (1977) 1791–1797.

[14] C. Robilliard, R. Battesti, M. Fouché, J. Mauchain, A.-M. Sautivet, F. Amiranoff & C. Rizzo. *Phys. Rev. Lett.*, **99** (2007) 190403.

[15] E. Zavattini, G. Zavattini, G. Ruoso, E. Polacco, E. Milotti, M. Karuza, U. Gastaldi, G. Di Domenico, F. Della Valle, R. Cimino, S. Carusotto, G. Cantatore & M. Bregant. *Phys. Rev. Lett.*, **96** (2006) 110406.

[16] G. Chardin. *Nuclear Physics A*, **558** (1993) 477 – 495. ISSN 0375-9474.

[17] Michael Martin Nieto & T. Goldman. *Physics Reports*, **205**, 5 (1991) 221 – 281. ISSN 0370-1573.

[18] M.H. Holzscheiter, G. Bendiscioli, A. Bertin, G. Bollen, M. Bruschi, C. Cesar, M. Charlton, M. Corradini, D. DePedis, M. Doser, J. Eades, R. Fedele, X. Feng, F. Galluccio, T. Goldman, J.S. Hangst, R. Hayano, D. Horvth, R.J. Hughes, N.S.P. King, K. Kirsebom, H. Knudsen, V. Lagomarsino, R. Landua, G. Laricchia, R.A. Lewis, E. Lodi-Rizzini, M. Macri, G. Manuzio, U. Marconi, M.R. Masullo, J.P. Merrison, S.P. Mller, G.L. Morgan, M.M. Nieto, M. Piccinini, R. Poggiani, A. Rotondi, G. Rouleau, P. Salvini, N. Semprini-Cesari, G.A. Smith, C.M. Surko, G. Testera, G. Torelli, E. Uggerhj, V.G. Vaccaro, L. Venturelli, A. Vitale, E. Widmann, T. Yamazaki, Y. Yamazaki, D. Zanello & A. Zoccoli. *Nuclear Physics B - Proceedings Supplements*, **56**, 1-2 (1997) 336 – 348. ISSN 0920-5632. Proceedings of the Fourth Biennial Conference on Low Energy Antiproton Physics.

[19] G. Gabrielse, P. Larochelle, D. Le Sage, B. Levitt, W. S. Kolthammer, R. McConnell, P. Richerme, J. Wrubel, A. Speck, M. C. George, D. Grzonka, W. Oelert, T. Sefzick, Z. Zhang, A. Carew, D. Comeau, E. A. Hessels, C. H. Storry, M. Weel & J. Walz. *Phys. Rev. Lett.*, **100** (2008) 113001.

[20] J. H. Christenson, J. W. Cronin, V. L. Fitch & R. Turlay. *Phys. Rev. Lett.*, **13**, 4 (1964) 138–140.

[21] V. Fanti et al. *Physics Letters B*, **465**, 1-4 (1999) 335 – 348. ISSN 0370-2693.

[22] A. Alavi-Harati et al. *Phys. Rev. Lett.*, **83**, 1 (1999) 22–27.

[23] B. Aubert et al. *Phys. Rev. D*, **65**, 5 (2002) 051101.

[24] A. Abashian et al. *Phys. Rev. Lett.*, **86**, 12 (2001) 2509–2514.

[25] R. Aaij et al. *Phys. Rev. Lett.*, **108** (2012) 111602.

[26] Makoto Kobayashi & Toshihide Maskawa. *Progress of Theoretical Physics*, **49**, 2 (1973) 652–657.

[27] A G Cohen, D B Kaplan & A E Nelson. *Annual Review of Nuclear and Particle Science*, **43**, 1 (1993) 27–70.

[28] Patrick Huet & Eric Sather. *Phys. Rev. D*, **51**, 2 (1995) 379–394.

[29] Ziro Maki, Masami Nakagawa & Shoichi Sakata. *Progress of Theoretical Physics*, **28**, 5 (1962) 870–880.

[30] K. Kimura & A. Takamura. *Modern Physics Letters A*, **14** (1999) 1625–1635.

[31] M. Doi, T. Kotani, H. Nishiura, K. Okuda & E. Takasugi. *Physics Letters B*, **102**, 5 (1981) 323 – 326. ISSN 0370-2693.

[32] M.C. Gonzalez-Garcia & Michele Maltoni. *Physics Reports*, **460**, 1-3 (2008) 1 – 129. ISSN 0370-1573.

[33] G. Karagiorgi, A. Aguilar-Arevalo, J. M. Conrad, M. H. Shaevitz, K. Whisnant, M. Sorel & V. Barger. *Phys. Rev. D*, **75** (2007) 013011.

[34] E.Kh. Akhmedov, G.C. Branco & M.N. Rebelo. *Physics Letters B*, **478**, 13 (2000) 215 – 223. ISSN 0370-2693.

[35] G. C. Branco, P. M. Ferreira, L. Lavoura, M. N. Rebelo, Marc Sher & Joao P. Silva. "Theory and phenomenology of two-Higgs-doublet models" (2011).

[36] N. G. Deshpande, M. Gupta & Palash B. Pal. *Phys. Rev. D*, **45** (1992) 953–957.

[37] N. G. Deshpande & Xiao-Gang He. *Phys. Rev. D*, **49** (1994) 4812–4819.

[38] M.N. Rebelo. *Fortschritte der Physik*. ISSN 1521-3978.

[39] G. Senjanovic & R. N. Mohapatra. *Phys. Rev. D*, **12** (1975) 1502–1505.

[40] T. Aaltonen et al. *Phys. Rev. Lett.*, **103** (2009) 041801.

[41] A. Abulencia et al. *Phys. Rev. Lett.*, **96** (2006) 211801.

[42] Yue Zhang, Haipeng An, Xiangdong Ji & R. N. Mohapatra. *Phys. Rev. D*, **76** (2007) 091301.

[43] H. P. Nilles. *Physics Reports*, **110**, 1-2 (1984) 1 – 162. ISSN 0370-1573.

[44] S. Abel, S. Khalil & O. Lebedev. *Nuclear Physics B*, **606**, 1-2 (2001) 151 – 182. ISSN 0550-3213.

[45] M. Rohini Godbole. *Pramana Journal of Physics*, **67** (2006) 835–847.

[46] Tarek Ibrahim & Pran Nath. *Physics Faculty Publications*, **89**. `http://hdl.handle.net/2047/d20000442`

[47] Savas Dimopoulos & Howard Georgi. *Physics Letters B*, **117**, 5 (1982) 287 – 290. ISSN 0370-2693.

[48] Res Jost. *Helvetica Physcia Acta*, **30** (1957) 409–416.

[49] M. B. Gavela, A. Le Yaouanc, L. Oliver, O. Pne, J. C. RaynaL & T. N. Pham. *Physics Letters B*, **109**, 3 (1982) 215 – 220. ISSN 0370-2693.

[50] Bruce H. J. McKellar, S. R. Choudhury, Xiao-Gang He & Sandip Pakvasa. *Physics Letters B*, **197**, 4 (1987) 556 – 560. ISSN 0370-2693.

[51] Steven A. Abel & Oleg Lebedev. *JHEP*, **0601** (2006) 133.

[52] J. M. Pendlebury & E. A. Hinds. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, **440**, 3 (2000) 471 – 478. ISSN 0168-9002.

[53] C. A. Baker et al. *Phys. Rev. Lett.*, **97**, 13 (2006) 131801.

[54] V. F. Dmitriev & R. A. Sen'kov. *Phys. Rev. Lett.*, **91** (2003) 212303.

[55] J. J. Hudson, D. M. Kara, I. J. Smallman, B. E. Sauer, M. R. Tarbutt & E. A. Hinds. *Nature*, **473** (2011) 493–496.

[56] G. W. Bennett et al. *Phys. Rev. D*, **80** (2009) 052008.

[57] K. Inami et al. *Physics Letters B*, **551**, 12 (2003) 16 – 26. ISSN 0370-2693.

[58] W. C. Griffith, M. D. Swallows, T. H. Loftus, M. V. Romalis, B. R. Heckel & E. N. Fortson. *Phys. Rev. Lett.*, **102** (2009) 101601.

[59] F. del Aguila & Marc Sher. *Physics Letters B*, **252**, 1 (1990) 116 – 118. ISSN 0370-2693.

[60] John Ellis. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, **284**, 1 (1989) 33 – 39. ISSN 0168-9002.

[61] E. Fermi & L. Marshall. *Phys. Rev.*, **72**, 12 (1947) 1139–1146.

[62] E. M. Purcell & N. F. Ramsey. *Phys. Rev.*, **78**, 6 (1950) 807.

[63] Norman F. Ramsey. *Phys. Rev.*, **78**, 6 (1950) 695–699.

[64] I. I. Rabi, J. R. Zacharias, S. Millman & P. Kusch. *Phys. Rev.*, **53**, 4 (1938) 318.

[65] J. H. Smith, E. M. Purcell & N. F. Ramsey. *Phys. Rev.*, **108**, 1 (1957) 120–122.

[66] I. S. Altarev et al. *Nuclear Physics A*, **341**, 2 (1980) 269 – 283. ISSN 0375-9474.

[67] Ya.B. Zel'dovich. *Soviet Phys. JETP*, **9** (1959) 1389.

[68] P. D. Miller, W. B. Dress, J. K. Baird & Norman F. Ramsey. *Phys. Rev. Lett.*, **19**, 7 (1967) 381–384.

[69] W. B. Dress, P. D. Miller & N. F. Ramsey. *Phys. Rev. D*, **7**, 11 (1973) 3147–3149.

[70] W. B. Dress, P. D. Miller, J. M. Pendlebury, Paul Perrin & Norman F. Ramsey. *Phys. Rev. D*, **15**, 1 (1977) 9.

[71] I. S. Altarev et al. *Physics Letters B*, **102**, 1 (1981) 13 – 16. ISSN 0370-2693.

[72] Altarev I. S. et al. *JETP Letters*, **44**, 8 (1986) 460.

[73] K. F. Smith et al. *Physics Letters B*, **234**, 1-2 (1990) 191 – 196. ISSN 0370-2693.

[74] I. S. Altarev et al. *Physics Letters B*, **276**, 1-2 (1992) 242 – 246. ISSN 0370-2693.

[75] P. G. Harris et al. *Phys. Rev. Lett.*, **82**, 5 (1999) 904–907.

[76] M.G.D. van der Grinten. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, **611**, 2-3 (2009) 129 – 132. ISSN 0168-9002. Particle Physics with Slow Neutrons.

[77] C. A. Baker, S. N. Balashov, J. Butterworth, P. Geltenbort, K. Green, P. G. Harris, M. G. D. van der Grinten, P. S. Iaydjiev, S. N. Ivanov, J. M. Pendlebury, D. B. Shiers, M. A. H. Tucker & H. Yoshiki. *Physics Letters A*, **308**, 1 (2003) 67 – 74. ISSN 0375-9601.

[78] J. Gerhold. *Cryogenics*, **38**, 11 (1998) 1063 – 1081. ISSN 0011-2275.

[79] W. Mampe, P. Ageron & R. Ghler. *Zeitschrift fr Physik B Condensed Matter*, **45** (1981) 1–14. ISSN 0722-3277.

[80] C. A. Baker, S. N. Balashov, K. Green, M. G. D. van der Grinten, P. S. Iaydjiev, S. N. Ivanov, J. M. Pendlebury, D. B. Shiers, M. A. H. Tucker, H. Yoshiki & P. Geltenbort. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, **501**, 2-3 (2003) 517 – 523. ISSN 0168-9002.

[81] W. Meissner & R. Ochsenfeld. *Naturwissenschaften*, **21** (1933) 787–788. ISSN 0028-1042. 10.1007/BF01504252.

[82] K. Grohmann & D. Hechtfischer. *Cryogenics*, **17**, 10 (1977) 579 – 581. ISSN 0011-2275.

[83] A. Katsika, M. Hardiman, D. Shiers & D. K. Hill. "Improving the Cryogenic SF". Presentation at collaboration meeting (2008).

[84] J. D. Cockcroft & E. T. S. Walton. *Proceedings of the Royal Society of London. Series A*, **129**, 811 (1930) 477–489.

[85] Norman F. Ramsey. *Phys. Rev.*, **100** (1955) 1191–1194.

[86] J. M. Pendlebury, W. Heil, Yu. Sobolev, P. G. Harris, J. D. Richardson, R. J. Baskin, D. D. Doyle, P. Geltenbort, K. Green, M. G. D. van der Grinten, P. S. Iaydjiev, S. N. Ivanov, D. J. R. May & K. F. Smith. *Phys. Rev. A*, **70** (2004) 032102.

[87] K. Green, P.G. Harris, P. Iaydjiev, D.J.R. May, J.M. Pendlebury, K.F. Smith, M. van der Grinten, P. Geltenbort & S. Ivanov. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, **404**, 2-3 (1998) 381 – 393. ISSN 0168-9002.

[88] The EDM Collaboration. "A New Search for the Neutron Electric Dipole Moment". Funding Pre-proposal submitted to The Department of Energy (2002).

[89] M Schmelz, R Stolz, V Zakosarenko, T Schnau, S Anders, L Fritzsch, M Mck & H-G Meyer. *Superconductor Science and Technology*, **24**, 6 (2011) 065009.

[90] S. K'doah. "Gravity Probe B: Exploring Einstein's Universe with Gyroscopes". NASA (2004).

[91] Bartington Intruments. "Operation Manual for Mag-01 & Mag-01H Single Axis Fluxgate Magnetometers".

[92] B. D. Josephson. *Rev. Mod. Phys.*, **46**, 2 (1974) 251–254.

[93] Y. Aharonov & D. Bohm. *Phys. Rev.*, **115** (1959) 485–491.

[94] John Clarke & Alex I. Braginski. "The SQUID Handbook: Fundamentals and Technology of SQUIDs and SQUID Systems". Wiley-VCH, 1^{st} edition (2004). ISBN 3527402292.

[95] P. Dutta & P. M. Horn. *Rev. Mod. Phys.*, **53** (1981) 497–516.

[96] M. B. Weissman. *Rev. Mod. Phys.*, **60** (1988) 537–571.

[97] C. T. Rogers & R. A. Buhrman. *Phys. Rev. Lett.*, **53** (1984) 1272–1275.

[98] Stuart Ingleby. "Cryodetector readout for direct dark matter searches". Ph.D. thesis, University of Oxford (2011).

[99] K. Green et al. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, **404**, 2-3 (1998) 381 – 393. ISSN 0168-9002.

[100] G. Golub & W. Kahan. *Journal of the Society for Industrial and Applied Mathematics: Series B, Numerical Analysis*, **2**, 2 (1965) 205–224. ISSN 00361429.

[101] Press W. H., Flannery B. P., S. A. Teukolsky & W. T. Vetterling. "Numerical Recipes in C: The Art of Scientific Computing". Cambridge University Press, 2^{nd} edition (1992). ISBN 978-0521431088.

[102] E. W. Collings. "Materials Properties Handbook: Titanium Alloys". ASM International (1994). ISBN 978-0871704818.

[103] ASM. "Asm Handbook: Properties and Selection : Nonferrous Alloys and Special-Purpose Materials". ASM International (1990). ISBN 978-0871703781.

[104] Y. Oda, G. Fujii & H. Nagano. *Cryogenics*, **14**, 2 (1974) 84 – 87. ISSN 0011-2275.

[105] Hiroshi Akima. *J. ACM*, **17** (1970) 589–602. ISSN 0004-5411.

[106] T Francavilla R. Hein & D. Liebenberg, editors. "Magnetic Susceptibility of Superconductors and Other Spin Systems", pages 49–80. Plenum (1991).

[107] U Divakar, S Henry, H Kraus & A J B Tolhurst. *Superconductor Science and Technology*, **21**, 6 (2008) 065021.

[108] N. Ishikawa, K. Nagata, H. Sato, N. Kasai & S. Kiryu. *Applied Superconductivity, IEEE Transactions on*, **3**, 1 (1993) 1910 –1913. ISSN 1051-8223.

[109] D. Drung, E. Dantsker, F. Ludwig, H. Koch, R. Kleiner, John Clarke, S. Krey, D. Reimer, B. David & O. Doessel. *Applied Physics Letters*, **68**, 13 (1996) 1856–1858.

[110] Marc-Antoine Parseval des Chnes. *Mmoires prsents lInstitut des Sciences, Lettres et Arts, par divers savans, et lus dans ses assembles. Sciences, mathmatiques et physiques. (Savans trangers.)*, **1** (1806) 638–648.

[111] John Clarke & Alex I. Braginski. "The SQUID Handbook: Fundamentals and Technology of SQUIDs and SQUID Systems", pages 16–20. Wiley-VCH, 1^{st} edition (2004). ISBN 3527402292.

[112] S Denis, L Dusoulier, M Dirickx, Ph Vanderbemden, R Cloots, M Ausloos & B Vanderheyden. *Superconductor Science and Technology*, **20**, 3 (2007) 192.

[113] The GNU Scientific Library. "Multidimensional Root-Finding".

[114] F. James. "MINUIT – Function Minimization and Error Analysis Reference Manual". CERN Program Library Long Writeup D506.

[115] J. M. Pendlebury. "CryoEDM magnetic shielding". Presentation at collaboration meeting (2009).

[116] S. Pfeiffer. "The Ogg Encapsulation Format Version 0" (2003).

[117] ISO/IEC. "ISO/IEC 11172-3:1993 - Information technology – Coding of moving pictures and associated audio for digital storage media at up to about 1,5 Mbit/s – Part 3: Audio". Padrão (1993).

[118] Pence, W. D., Chiappetti, L., Page, C. G., Shaw, R. A. & Stobie, E. *A&A*, **524** (2010) A42. `http://dx.doi.org/10.1051/0004-6361/201015362`

[119] ISO/IEC. "ISO/IEC 14882:2003 - Programming languages – c++" (2003).

[120] Margaret Ellis. "The annotated C++ reference manual". Addison-Wesley, Reading, Mass (1990). ISBN 0201514591.

[121] R. Rivest. "The MD5 Message-Digest Algorithm" (1992).

[122] H. Nyquist. *Proceedings of the IEEE*, **90**, 2 (2002) 280 –305. ISSN 0018-9219.

[123] S. W. Smith. "The Scientist & Engineer's Guide to Digital Signal Processing". California Technical Pub, 1^{st} edition (1997). ISBN 978-0966017632.

[124] A. Nuttall. *Acoustics, Speech and Signal Processing, IEEE Transactions on*, **29**, 1 (1981) 84 – 91. ISSN 0096-3518.

[125] W. E. Thomson. *Proceedings of the Institution of Electrical Engineers, Part III*, **96** (1949) 487–490.

[126] S. Butterworth. *Experimental Wireless and the Wireless Engineer*, **7** (1930) 536–541.

[127] Krall, H. L. & Fink, O. *Trans. Amer. Math. Soc.*, **65** (1948) 100–115.

[128] Godse, A. P. & Bakshi, U. A. "Linear Integrated Circuits". Technical Publications, 1^{st} edition (2008). ISBN 978-8184313253.

[129] Brown, J. W. & Churchill, R. V. "Complex Variables and Applications". McGraw-Hill, 6^{th} edition (1995). ISBN 0079121470.