University of Sussex

A University of Sussex PhD thesis

Available online via Sussex Research Online:

http://sro.sussex.ac.uk/

This thesis is protected by copyright which belongs to the author.

This thesis cannot be reproduced or quoted extensively from without first obtaining permission in writing from the Author

The content must not be changed in any way or sold commercially in any format or medium without the formal permission of the Author

When referring to this work, full bibliographic details including the author, title, awarding institution and date of the thesis must be given

Please visit Sussex Research Online for more information and further details

University of Sussex

Data and Systematic Error Analysis for the Neutron Electric Dipole Moment Experiment at the Paul Scherrer Institute and Search for Axionlike Dark Matter

Nicholas Ayres

Submitted for the degree of Doctor of Philosophy University of Sussex 14 December 2018

Declaration

I hereby declare that this thesis has not been and will not be submitted in whole or in part to another University for the award of any other degree.

Signature:

Nicholas Ayres

UNIVERSITY OF SUSSEX

NICHOLAS AYRES, DOCTOR OF PHILOSOPHY

DATA AND SYSTEMATIC ERROR ANALYSIS FOR THE NEUTRON ELECTRIC DIPOLE MOMENT EXPERIMENT AT THE PAUL SCHERRER INSTITUTE AND SEARCH FOR AXIONLIKE DARK MATTER

This thesis details work conducted as part of the experimental collaboration responsible for the neutron electric dipole moment (nEDM) experiment based at the Paul Scherrer Institute (PSI). The nEDM is a sensitive probe of a broad range of new CP violating physics beyond the standard model, however it remains elusive: while historic experiments since 1951 have increased in sensitivity by over six orders of magnitude, a nonzero nEDM is yet to be detected. Many theories of physics beyond the standard model predict neutron EDMs of a size that would be detectable by current and next generation experiments, and it has been said that measurements of the neutron EDM have ruled out more theories than any other experiment. One explanation of the smallness of the neutron EDM, the Peccei-Quinn theory, invokes a novel particle, the axion, which is also a credible dark matter candidate. The axion is yet to be detected.

The work covers three main areas. First, contributions to the data analysis technique used to analyse the main dataset to produce a new world-leading limit on the neutron EDM. Second, an auxiliary measurement campaign to map the magnetic field within the experiment's magnetic shields is described, and the analysis of these datasets to provide corrections for several critical systematic effects is presented. Finally, a novel analysis of the data taken at a previous-generation nEDM experiment is used to derive the first experimental limits on the coupling of axion-like dark matter particles to gluons is described. These exclusions are up to 1000 times stronger than previous results for cosmologically interesting 10^{-22} eV axions.

Acknowledgements

I would first like to thank my supervisor, Philip Harris, for his excellent guidance and support throughout my time at Sussex. Ever since I first came to the group as an summer student in the CryoEDM days, he has always found time for me, and been my greatest advocate.

I also offer thanks to all of the Sussex nEDM group, who have all contributed to making time here a success. In particular, my fellow PhD students Jacob and Chris have always been there for me, especially in making sure that we never had a collaboration meeting Saturday without a hangover.

The mapping seemed like a gargantuan task, but with the help of many people it became manageable. I would like to thank everyone involved in the mapping of 2017 for the excellent teamwork, especially Philipp, Georg, Michi, Peter, Allard and Jochen. I also thank Elise and again Philipp for collaborating on the nEDM analysis. For the support and collaboration on the axion analysis, I would like to thank Michał.

I am very grateful to the entire team at PSI, who made great effort to welcome me during my LTA. Thank you Sybille for the all late night craziness, t-shirt buying and tortellini with Pesto-Rahm sauce that got me through the mapping. Without Nicolas, I probably wouldn't have ever found my way back from the Niederdörfli, and probably I would have gotten much more work done and drank a lot less coffee. Thanks again to Michi, for so much help with everything, and also for the help with my German, even back in the beginning when it was quite slow!

The experiment was made possible only by the hard work of many people at many institutions. The PSI nEDM collaboration has always strived to do excellent work, and I couldn't have asked to work as part of a more enthusiastic international team.

Of course one also needs support from people who don't just measure zero all day to stay sane. The 24/7 discussion in *the group chat* on 'varied' topics has provided a very welcome distraction and procrastination tool.

Finally, I would like to offer my greatest thanks to my family for their encouragement and help all these years.

Contents

Li	List of Tables x			
Li	List of Figures xx			
1	Intr	ntroduction		
2	The	ory and Motivation	3	
	2.1	C, P and T Invariance in the Standard Model	3	
	2.2	Baryon Asymmetry and Baryogenesis	6	
	2.3	The Neutron Electric Dipole Moment	8	
	2.4	EDMs of other systems	9	
	2.5	Standard Model predictions for the nEDM	11	
		2.5.1 Weak Sector	11	
		2.5.2 Strong CP	11	
	2.6	Relation of EDMs to Underlying Physics Through Effective Field Theories .	11	
	2.7	Historic and Competing nEDM Measurements	13	
	2.8	Ultracold Neutrons	16	
3	The	Search for the Neutron Electric Dipole Moment at the Paul Scherrer		
	Inst	itute	20	
	3.1	Introduction and Principle of Operation	20	
	3.2	The PSI UCN Source	29	
	3.3	UCN Systems	31	
	3.4	Magnetic Field Environment	36	
	3.5	Magnetometry	42	
	3.6	Data Acquisition	44	
	3.7	High Voltage	46	
	3.8	Conclusion: Expected Sensitivity and Outlook to n2EDM	47	

4	Dat	ta Analysis	5	52
	4.1	Structure of the Analysis	E	52
	4.2	Importing the Data	E	53
	4.3	Cycle Level Analysis	E	53
	4.4	Run Level Analysis	E	54
		4.4.1 Division of Sets into Subsets	E	5 4
		4.4.2 Caesium Gradients Fit	E	55
		4.4.3 Ramsey Fit	E	56
		4.4.4 'RvsE' and 'RvsT' fit \ldots	6	31
	4.5	Crossing Point Analysis: A First Exposition	6	35
	4.6	Corrections to the Crossing Point Analysis	6	37
		4.6.1 False EDMs and Gravitational Shifts	6	37
		4.6.2 The Earth's Rotation	6	39
		4.6.3 $< B_T^2 > $ Correction	6	39
	4.7	The Crossing Point Analysis: Recapitulation	7	70
	4.8	Sensitivity and Result	7	71
5	Ma	apping of the Magnetic Field within the nEDM appa	uratus 7	$^{\prime}2$
	5.1	Motivation	7	72
	5.2	Accuracy Requirements	7	73
	5.3	Mapper Hardware	7	75
	5.4	Datataking Summary	7	76
	5.5	Known Imperfections of the Mapper	7	77
	5.6	General Model of Fluxgate Mapper Misalignments, Nonort	thogonalities and	
		Offsets	7	79
	5.7	Determination of Mapper Misalignments, Nonorthogonalit	ties and Offsets 8	31
		5.7.1 Direct Measurement of Mapper Arm Profile	8	32
		5.7.2 Analysis of Inclinometer Measurements	8	35
		5.7.3 In Situ Determination of Fluxgate Nonorthogonali	ties and Offsets . §	<i>)</i> 2
		5.7.4 In Situ Mapper Misalignments Measurement		<i>)</i> 5
	5.8	Fit of Field Coefficients		<i>)</i> 7
		5.8.1 Decomposition into Rings	9) 8
		5.8.2 Choice of Fit Order and Goodness of Fit	10)0
		5.8.3 Fit of Magnetic Field Gradients to Ring Coefficien	ts: An Oversim-	
		plified First Exposition	10)4

		5.8.4	Compensating for Misalignments and Nonorthogonalities	105
		5.8.5	Assessment of fit accuracy and robustness	108
		5.8.6	Effect of Position Offsets	124
	5.9	Repro	ducibility of Magnetic Field	126
	5.10	Fit of	Trimcoil Fields	144
	5.11	Linear	rity	146
	5.12	Evalua	ation of Gradients During nEDM Datataking	149
	5.13	Concl	usion	151
6	Sea	rch for	Axion-Like Matter Through Nuclear Spin Precession in Elec	-
	tric	and M	Agnetic Fields	152
	6.1	Motiv	ation and Theory	153
		6.1.1	The Strong CP Problem and Axions	153
		6.1.2	Ultralight Axions as a Dark Matter Candidate	154
		6.1.3	Effects of Axionic Dark Matter in nEDM Experiments	155
		6.1.4	Other Experimental Probes of Axions and Axionic Dark Matter	156
	6.2	Exper	iment	159
		6.2.1	The Sussex-RAL-ILL Room Temperature EDM Experiment	159
	6.3	Analy	sis of the Sussex-RAL-ILL Dataset	159
		6.3.1	Preparation of a Time Series from the Sussex-RAL-ILL Experi-	
			mental Data	160
		6.3.2	Least Squares Spectral Analysis and the Lomb-Scargle Periodogram	160
		6.3.3	Estimation of the CDF at one Frequency Under the Null Hypothesis	162
		6.3.4	Search for Oscillations at any Frequency	162
		6.3.5	Exclusion of oscillations	165
	6.4	Analy	sis of the PSI dataset	167
		6.4.1	Fit Model	167
		6.4.2	Detection Analysis	168
		6.4.3	Exclusion Analysis	170
		6.4.4	Wind effect analysis	171
	6.5	Concl	usion	172
7	Clos	sing W	Vords	174
B	iblios	raphy		175
		, apiny		110

A The Standardised Adequate Basis of Polynomials

List of Tables

2.1	Transformation of some common quantities under \hat{C} , \hat{P} and \hat{T} 8
3.1	Forecast sensitivity parameters of n2EDM experiment in comparison with previous experiments. Calculations for n2EDM are from the n2EDM Pro- ject Overview Report (January 2018) [162], considering improvements achiev- able using only presently proven technology - the 'Baseline', and assuming no improvements in UCN source performance. Performance figures of ILL experiment from [20]
5.1	Results of fit to sets of calibration maps in different fields to determine angles described in text
5.2	Goodness of fit parameters comparing fit of $m = 0$ using only vertical (Y) and both radial and vertical $(Z \text{ and } Y)$ sensors
5.3	Goodness of fit parameters comparing fit of $m = 0$ using only vertical (Y) sensor, with and without misalignment corrections $\dots \dots \dots$
5.4	Goodness of fit parameters comparing fit of $m = 0$ using vertical and radial sensors (Y and Z) sensor, with and without misalignment corrections \dots 112
5.5	Goodness of fit parameters comparing fit of $m = 0$ using only vertical Y sensor, with and without drift corrections $\dots \dots \dots$
5.6	Goodness of fit parameters comparing fit of $m = 0$ using vertical Y and radial Z sensor, with and without drift corrections $\dots \dots \dots$
5.7	Goodness of fit parameters comparing fit of $m = 0$ using only the vertical (Y) sensor, using only measurement points with $R < 240$ mm
5.8	Goodness of fit parameters comparing fit of $m = 0$ using only the vertical (Y) sensor, using only height $-80 < 0 < 80$ mm

5.9	Goodness of fit parameters comparing fit of $m = 0$ using only the vertical
	(Y) sensor, using only measurement points with $R < 240$ mm and height
	-80 < 0 < 80mm
A.1	The basis of harmonic polynomials sorted by degree
A.2	The basis of harmonic polynomials sorted by degree in cylindrical coordin-
	ates

List of Figures

2.1	Comparison of Feynman diagrams for the Fermi theory, and the underlying	
	high-energy process. From [1]	12
2.2	Schematic diagram illustrating the parameters coming into play to compute	
	different EDMs at different scales from the underlying BSM physics. Dashed	
	lines indicate a weaker dependence. In this notation, d_e is the electron	
	EDM, C_{qe} and C_{qq} are various 4-fermion couplings (quark-electron or quark-	
	quark), θ is the QCD $\overline{\theta}$, w is the Weinberg 3-gluon operator and d_q and	
	\tilde{d}_q and the quark EDMs and chromo-EDMs. From these, operators at	
	the nuclear scale $g_{\pi NN}$ (CP-violating pion-nucleon couplings) and $C_{S,P,T}$	
	(possible electron-nucleon couplings) can be derived and used to compute	
	atomic EDMs. From [79]	13
3.1	Sketch showing the layout of key components of the nEDM experiment at	
	the PSI, described in turn throughout this chapter. From [91]. \ldots	20
3.2	Aerial view of the Paul Scherrer Institute straddling the river Aare, looking	
	north. The round building in the foreground is the SLS. The HIPA acceler-	
	ator is located in the large white building behind it, and the UCN source,	
	cooling plant and nEDM experiment are located in the attached building	
	on the left.	22
3.3	Sketch of the neutron precession during a Ramsey measurement sequence	
	at the resonant frequency. Adapted from [121]	25
3.4	Ramsey fit of the Asymmetry $A = \frac{N_{\uparrow} - N_{\downarrow}}{N_{\uparrow} + N_{\downarrow}}$ to a typical section of EDM	
	datataking. Explanation of fit parameters indicated will be included in	
	chapter 4	26

- 3.5 The Ramsey fit of the asymmetry $A = \frac{N_{\uparrow} N_{\downarrow}}{N_{\uparrow} + N_{\downarrow}}$ from a special measurement covering a larger section of the Ramsey resonance curve beyond the central fringe. Shown above are the results over the full measurement range, while below is a zoomed view of the central few fringes. Note that at the extremes, the Rabi envelope becomes clearly visible over the tight Ramsey fringe pattern. 27
- View inside the vacuum tank, showing the precession chamber used for the 3.6 nEDM experiment. The castellated high voltage feed at the top connects to the upper electrode, which is separated from the lower electrode by the insulator ring. The UCN enter and leave the chamber from below. The mercury probe light travels across the chamber through two windows in the 28Diagram of the PSI UCN Source, from [131]. 3.730Comparison of the PSI UCN source performance [131] with several compet-3.8 30 itors worldwide, from [125]. 3.9 Photo of the UCN switch. The inner platform with different guides mounted can be rotated in order to make different connections between the various ports for filling and emptying of UCNs. 323.10 Storage curve showing measured neutron counts after a given storage time, measured at the beginning of 2016 data taking. The storage curve cannot be well described by a simple exponential because slow and fast UCN will have different storage time constants, and the population of stored UCN will typically contain a range of energies. The curve is better approximated by the sum of two exponentials, with fast and slow decay constants. . . . 33

xii

3.13	Photo of the NANOSC detector. On the front can be seen the nine scin-	
	tillator stacks, connected by the blue light guide assembly to nine PMTs	
	mounted at the rear. From [137]. \ldots \ldots \ldots \ldots \ldots \ldots	36
3.14	Photo of the experimental installation, with the μ -metal shield lifted by a	
	crane during installation. The six SFC coils are marked	38
3.15	Measured Ramsey visibility as a function of time for an optimised trimcoil	
	current configuration using the homogenisation algorithm [127], from spe-	
	cial datataking runs during the 2016 ramp up period. The measured T_2	
	of (1568 ± 19) s shows a vast improvement over that obtained in previous	
	experiments, which were of the order 600s [88]. This directly increases the	
	statistical sensitivity of the apparatus.	39
3.16	Measured Ramsey visibility after 180s compared to Spin-Echo measure-	
	ment[117] and longitudinal polarisation measurements, versus additional	
	vertical gradient applied on top of optimised configuration. The longitud-	
	inal polarisation measurements probe the baseline efficiency of the initial	
	polarisation and of the specificity and efficiency of the USSA spin sensitive	
	neutron detection system, as well as ' T_1 ' (or spin-lattice) type depolarisa-	
	tion such as that due to small scale magnetic contamination of the chamber	
	walls. The Spin-Echo measurements probe ' T_2 ' (spin-spin) type depolar-	
	is ation, such as that due to large scale $B\mbox{-field}$ inhomogeneities such as a	
	$\frac{\partial B_z}{\partial x}$. The Ramsey measurement also includes the effect of 'gravitationally	
	enhanced depolarisation', a T_2^* effect. This type of depolarisation occurs	
	when UCN are stored in a vertical gradient $\frac{\partial B_z}{\partial z}$, as UCN of different energies	
	have a different height distribution throughout the chamber. These results	
	provide clear validation of the measurements presented in [116], which are	

3.17 Sketch of entire n2EDM apparatus in area South. Grey blocks around area are concrete blocks for radiation protection. Magnetically shielded room (MSR) to contain vacuum tank (5) and precession chambers (4), mounted on large granite blocks (7) for stability. Superconducting magnet polariser (1) provides a polarised UCN beam from beam port (2). This feeds a Y-piece and switch (3) to feed each precession chamber. Upon emptying the chambers, the neutrons of each spin state are counted by a pair of USSAs (6) (as used in the previous generation experiment, described in section 3.3). 49

3.18	Close up of n2EDM precession chamber conceptual design: dual cham- bers surrounding a single high voltage electrode, with independent mercury magnetometry systems for each chamber. Ground electrodes on top and bottom to support yest cassium magnetometer array for online determina-	
	tion of field configuration.	50
4.1	Time structure of datataking scheme. Upper plot: yellow = USSA SF2a/b state (inverted every four cycles, see section 3.3), blue = SF1 state (inverted every 112 cycles). Lower plot: green = E field polarity (8-48-8-48 cycle of 8 cycles datataking at zero field followed by 48 cycles at alternating HV states), alternating light/dark grey = indicative division of single long run into 'ABBA' sections, combining both HV states. From [91]	55
4.2	HVCombi Ramsey fit to a typical subset, with no correction for gradient drifts.	58
4.3	Second stage Ramsey fit of only phase to one HV chunk of a typical subset. The blue and black points correspond to different states of the USSA spin flippers SF2a/b	58
4.4	Fitted relative vertical gradient drift over a typical run.	59
4.5	HVCombi Ramsey fit to typical subset, with CsM gradient compensation. The colour of each point corresponds to the high voltage state, with red corresponding to an applied voltage of $+132$ kV, blue corresponding to -132 kV, and black to cycles where no high voltage was applied	60
4.6	Second stage Ramsey fit of only phase to one HV chunk of typical subset, with CsM gradient compensation.	61
4.7	Typical RvsE fit, without CsM gradient drift correction.	62
4.8	Typical RvsE fit, with CsM gradient drift correction.	63
4.9	Typical RvsT fit, without CsM gradient drift correction. The red points correspond to cycles taken at positive high voltage, blue to cycles taken at	
	negative high voltage, and grey to cycles taken at zero high voltage	64
4.10	Typical RvsT fit, with CsM gradient drift correction	64

4.11	Result of the crossing point analysis for all data of 2015 and 16. A correction is applied for the Earth's rotation, but corrections for $G_{3,0}$ and $\langle B_T^2 \rangle$ are not applied. In the upper plot, the result from each run taken with B up is marked by a red point, and in the lower plot the result taken with B down is marked by a blue point. The slopes for each of the two lines are equal in magnitude but of opposite sign, and fixed	66
5.1	Photo of the mapper as installed inside the nEDM vacuum chamber, with the radial \hat{r} , tangential $\hat{\phi}$ and vertical \hat{z} directions annotated	72
5.2	Uncorrected measured height profile of the mapper rails. This shows a clear deviation from straightness, on the order 1mm.	82
5.3	Misalignment α versus cart radius, computed in the cases where the slope is defined by the two front wheels, and the two rear wheels	84
5.4	Misalignment β versus cart radius, computed in the cases where the slope is defined by the two left wheels, and the two right wheels	84
5.5	Projection of global tilt on xz and yz axes, presented as a function of height z , as measured at the start of the 2017 mapping campaign. The lack of z -dependence of the global tilt indicates that the vertical axle of the mapper is properly aligned through the apparatus	87
5.6	Raw tilt and roll recorded during a 'finestar' map taken at the start of the 2017 mapping campaign. The apparant 'oscillations' arise from imperfections in the bearings which function as wheels for the mapper cart. Their apparant amplitudes of 0.5 mrad in the roll axis and 0.2 mrad in the tilt axis correspond to an imperfection (deviation from roundness of the outer housing, or radial run-out) of order 50 microns.	88
5.7	Computed α and β in 2017, from a fit to each ring of Figure 5.6	89
5.8	Difference between α and β as measured by inclinometers mounted in the base and tube of the mapper cart. A constant offset is expected as no special effort was made to calibrate the absolute readings. In the 2014 mapping campaign the mapper cart was overconstrained between two distorted rails, leading to significant bending and twisting of the card varying by up to 2-3	
	mrad across the radial motion range as shown above	90

5.9	Difference between α and β in 2017 as measured by inclinometers mounted
	in the base and tube of the mapper cart. A constant offset is expected as
	no special effort was made to calibrate the absolute readings. The bend of
	the cart is found to be constant to better than 0.1 mrad, while the twist
	shows some structure of approx 0.5 mrad in size. $\ldots \ldots \ldots \ldots \ldots $ 91
5.10	Typical first-stage fit result for a calibration map taken in the remnant field. 94
5.11	Example ring measurement and fit recorded towards the edge of mapped
	volume, where gradients are typically larger, including plots of the fitted
	field, truncated at order 2 or order 6
5.12	Typical Allan standard deviation measured during data taking, with the
	fluxgate magnetometer stationary at the centre of the shield. \ldots \ldots \ldots 102
5.13	Colourmap of the RMS error for each ring fit as a function of position,
	computed for fits to up to orders 2,4 and 6. Higher order terms grow in size
	very strongly with radius and height, corresponding to a poorer goodness
	of fit at larger radius when truncating the fit too soon 103
5.14	RMS fit error from a ring decomposition fit of all points within all rings of
	one B_0 map. For comparison, the measured Allan standard deviation at
	this timescale is approximately 20pT
5.15	Fitted value of terms $G_{0,m}$ for different choices of misalignment values, drift
	correction and sensor inclusion
5.16	Fitted value of terms $G_{1,m}$ for different choices of misalignment values, drift
	correction and sensor inclusion
5.17	Fitted value of terms $G_{2,m}$ for different choices of misalignment values, drift
	correction and sensor inclusion
5.18	Fitted value of terms $G_{3,m}$ for different choices of misalignment values, drift
	correction and sensor inclusion $\ldots \ldots \ldots$
5.19	Fitted value of terms $G_{4,m}$ for different choices of misalignment values, drift
	correction and sensor inclusion
5.20	Fitted value of terms $G_{5,m}$ for different choices of misalignment values, drift
	correction and sensor inclusion
5.21	Fitted value of terms $G_{6,m}$ for different choices of misalignment values, drift
	correction and sensor inclusion
5.22	Fitted value of terms $G_{0,m}$ for restricted domains of points

5.24	Fitted value of terms $G_{2,m}$ for restricted domains of points
5.25	Fitted value of terms $G_{3,m}$ for restricted domains of points
5.26	Fitted value of terms $G_{4,m}$ for restricted domains of points
5.27	Fitted value of terms $G_{5,m}$ for restricted domains of points
5.28	Fitted value of terms $G_{6,m}$ for restricted domains of points
5.29	Fitted value of $G_{3,0}$ from map 12112X15 and RMS fit error for small offsets
	in radius
5.30	Fitted value of $\left< B_T^2 \right>$ from map 12112X15 and RMS fit error for small offsets
	in radial zero position
5.31	Fitted value of $G_{3,0}$ and $\left< B_T^2 \right>$ from map 12112X15 for small offsets in the
	z zero position
5.32	Fitted gradients $l = 0$ for B up, considering only maps taken in 2017 127
5.33	Fitted gradients $l = 1$ for B up, considering only maps taken in 2017 128
5.34	Fitted gradients $l = 2$ for B up, considering only maps taken in 2017 128
5.35	Fitted gradients $l = 3$ for B up, considering only maps taken in 2017 129
5.36	Fitted gradients $l = 4$ for B up, considering only maps taken in 2017 129
5.37	Fitted gradients $l = 5$ for B up, considering only maps taken in 2017 130
5.38	Fitted gradients $l = 6$ for B up, considering only maps taken in 2017 130
5.39	$< B_T^2 >$ computed for each B up, considering only maps taken in 2017 131
5.40	Fitted gradients $l=0$ for B down, considering only maps taken in 2017 131
5.41	Fitted gradients $l=1$ for B down, considering only maps taken in 2017 132
5.42	Fitted gradients $l=2$ for B down, considering only maps taken in 2017 132
5.43	Fitted gradients $l=3$ for B down, considering only maps taken in 2017 133
5.44	Fitted gradients $l=4$ for B down, considering only maps taken in 2017 133
5.45	Fitted gradients $l=5$ for B down, considering only maps taken in 2017 134
5.46	Fitted gradients $l=6$ for B down, considering only maps taken in 2017 134
5.47	$< B_T^2 >$ computed for each B down map, considering only maps taken in
	2017
5.48	Fitted gradients $l = 0$ for B up, comparing maps taken in 2013, 2014 and
	2017, including the 2017 maps taken in the 2014 setup
5.49	Fitted gradients $l = 1$ for B up, comparing maps taken in 2013, 2014 and
	2017, including the 2017 maps taken in the 2014 setup
5.50	Fitted gradients $l = 2$ for B up, comparing maps taken in 2013, 2014 and
	2017, including the 2017 maps taken in the 2014 setup

5.51	Fitted gradients $l = 3$ for B up, comparing maps taken in 2013, 2014 and	
	2017, including the 2017 maps taken in the 2014 setup	. 137
5.52	Fitted gradients $l = 4$ for B up, comparing maps taken in 2013, 2014 and	
	2017, including the 2017 maps taken in the 2014 setup	. 137
5.53	Fitted gradients $l = 5$ for B up, comparing maps taken in 2013, 2014 and	
	2017, including the 2017 maps taken in the 2014 setup	. 138
5.54	Fitted gradients $l = 6$ for B up, comparing maps taken in 2013, 2014 and	
	2017, including the 2017 maps taken in the 2014 setup	. 138
5.55	$< B_T^2 > {\rm computed}$ for each B up map, comparing maps taken in 2013, 2014	
	and 2017, including the 2017 maps taken in the 2014 setup	. 139
5.56	Fitted gradients $l = 0$ for B down, comparing maps taken in 2013, 2014 and	
	2017, including the 2017 maps taken in the 2014 setup	. 139
5.57	Fitted gradients $l = 1$ for B down, comparing maps taken in 2013, 2014 and	
	2017, including the 2017 maps taken in the 2014 setup	. 140
5.58	Fitted gradients $l = 2$ for B down, comparing maps taken in 2013, 2014 and	
	2017, including the 2017 maps taken in the 2014 setup	. 140
5.59	Fitted gradients $l = 3$ for B down, comparing maps taken in 2013, 2014 and	
	2017, including the 2017 maps taken in the 2014 setup	. 141
5.60	Fitted gradients $l = 4$ for B down, comparing maps taken in 2013, 2014 and	
	2017, including the 2017 maps taken in the 2014 setup	. 141
5.61	Fitted gradients $l = 5$ for B down, comparing maps taken in 2013, 2014 and	
	2017, including the 2017 maps taken in the 2014 setup	. 142
5.62	Fitted gradients $l = 6$ for B down, comparing maps taken in 2013, 2014 and	
	2017, including the 2017 maps taken in the 2014 setup	. 142
5.63	$< B_T^2 >$ computed for each B down map, comparing maps taken in 2013,	
	2014 and 2017, including the 2017 maps taken in the 2014 setup	. 143
5.64	Example fits of selected gradients versus currents, for maps of TTC. Two	
	sets of maps were measured, which are fitted with a common trimcoil field	
	(directly proportional to applied current) and independent remnant fields,	
	allowed to vary freely between degaussings	. 145
5.65	$G_{3,0}$ predicted from B_0 and trimcoil maps, versus measured $G_{3,0}$ from base	
	configuration maps	. 146
5.66	Histogram of $G_{3,0}$ residuals when comparing configuration maps to values	
	predicted from B_0 maps and trimcoil maps $\ldots \ldots \ldots \ldots \ldots \ldots \ldots \ldots$. 147

5.67	$< B_T^2 >$ predicted from B_0 and trimcoil maps, versus measured $< B_T^2 >$	
	from base configuration maps	. 147
5.68	Histogram of $\langle B_T^2 \rangle$ residuals when comparing configuration maps to	
	values predicted from B_0 maps and trimcoil maps	. 148
6.1	Existing laboratory and observational constraints on axion-photon coupling,	
	figure taken from [17], based on work cited within. Does not include latest	
	results from ADMX in [213].	. 158
6.2	The results of a detection analysis as a function fitted amplitude versus trial	
	frequency. The black line represents the actual measured periodogram of the	
	dataset, while green indicates the mean power measured in each frequency	
	across all generated Monte Carlo datasets. The orange line indicates the	
	computed power necessary for significance at levels of 1-5 σ , for the fitted	
	$N_{\rm eff} = 1026$ (the true number of frequencies tested was $N = 1334$). If the	
	black line crosses the <i>n</i> th orange line, it is significant at $n\sigma$. 164
6.3	Comparison of exclusion results obtained using the CL_S technique(below),	
	versus a naïve $P(\text{data} \text{signal})$ (above). The blue line represents the bound-	
	ary of the 95% exclusion region. The area in black is strongly excluded,	
	while the yellow regions are allowed. The simple method results in claims	
	of anomalously strong (stronger than the static EDM limit) exclusions in	
	regions where the power spectrum fluctuated low; claims of such sensitivity	
	must be unphysical.	. 166
6.4	Periodograms computed for the PSI dataset, with features as described in	
	text. The black line is the periodogram computed for each electric field	
	state, while the dark green shows the mean of all generated Monte Carlo	
	periodograms under the null hypothesis. The three progressively lighter	
	green bands represent the 1,2 and 3σ deviations extracted from the Monte	
	Carlo data. Finally, the five orange lines indicate the 1,2,3,4 and 5 σ false	
	alarm thresholds, as defined within the text. In each insert, the area around	
	the inverse cycle spacing $(1/(5\text{minutes}))$ is shown, where a large increase in	
	power is not unexpected. Apparent detections within this reason should be	
	treated sceptically.	. 169
6.5	Computed exclusion region for EDM amplitude as a function of frequency	
	using (left) the naïve $P(\text{data} \text{signal})$ approach and (right) the CL_S method.	

6.6	Combined 95% exclusion region using CL_S analysis technique on both data-		
	sets		
6.7	Final exclusion regions obtained for the axion-gluon coupling (above) and		
	axion-neutron wind type coupling below. The axion-gluon result represents		
	the first laboratory bound on this quantity, improving by up to a factor 1000		
	for cosmologically interesting 10^{-22} eV axions, while the axion-nucleon result		
	represents an improvement by up to a factor 40 over previous laboratory		
	limits over several orders of magnitudes of mass, but is still less restrictive		

Chapter 1

Introduction

This thesis details work done as part of the PSI Neutron Electric Dipole Moment collaboration.

Measurements of the neutron EDM have been said over the years to have disproved more theories of physics than any other experiment. The first measurement took place in 1951, and a series of measurements since then have improved the sensitivity to this parameter by six orders of magnitude. It is a sensitive probe of the violation of CPsymmetry (the same symmetry as between matter and antimatter). This is an important topic in contemporary physics, as the standard model (our current-best theory of particle physics) predicts far too little CP violation to be able to explain the observed matterantimatter asymmetry in the universe. In the standard model, the neutron EDM (from the CKM phase) is predicted to be unmeasurably small, however in a wide variety of theories of physics beyond the standard model, new sources of CP violation induce large EDMs, of the order of magnitude that current and next generation EDM experiments are sensitive to. Observation of a non-zero EDM would be a clear sign of physics beyond the standard model.

Paradoxically, the standard model is also unable to explain why the neutron EDM is so small. The so-called QCD theta term in the strong-interaction part of the standard model has the capacity to naturally produce huge neutron EDMs. The *CP* violating phase $\bar{\theta}$ has been constrained to be less than 10^{-10} , where it is naturally of order 1. This presents a substantial fine-tuning problem in the standard model. One solution to this is the Peccei-Quinn model. This model proposes a mechanism to naturally absorb this phase so it relaxes to zero, and also implies the existence of a new particle, the axion. This particle is also an exciting dark matter candidate. In this case, it has been proposed that the value of the neutron EDM would oscillate. In chapter 2, the motivation to search for a neutron electric dipole moment and the theory behind it is explored.

In chapter 3, the layout and construction of the nEDM experiment at the PSI is presented.

The next chapter, chapter 4 contains a description of the statistical analysis of the data taken with the apparatus, which is liable to produce the world's best limit on the nEDM to date.

nEDM measurements are difficult experiments requiring extraordinary control of systematic effects. In chapter 5, the measurement and analysis of data acquired in the 2017 magnetic field mapping campaign to constrain several important systematic effects is described.

Finally, in chapter 6, a novel analysis of the data measured at the previous generation Sussex-RAL-ILL nEDM experiment to produce the first laboratory measurements of the couplings of hypothetical ultralight dark matter axions to gluons via searching for an oscillating neutron EDM is presented.

Chapter 2

Theory and Motivation

The neutron EDM is a powerful observable, giving a sensitive window through which a large variety of theories of physics which promise to solve issues within the standard model can be tested. A nonzero neutron EDM would violate parity P and time reversal T, and therefore also CP, the symmetry between matter and antimatter. One powerful motivation to search for new physics is the unexplained origin of the observed matter-antimatter asymmetry in the universe, which requires new sources of CP violation. New physics theories readily provide such asymmetries in abundance, often predicting large EDMs, and must often be tuned so as to avoid limits on these from EDM measurements. As of yet, a non-zero EDM of a fundamental particle is yet to be discovered, and upper limits on their possible values provide strong constraints on CP violation in physics theories beyond the standard model.

2.1 C, P and T Invariance in the Standard Model

The standard model of particle physics is our current best description of the subatomic world, providing predictions which have been experimentally verified up to an incredibly high level. It describes a world made up of matter particles (quarks and leptons) interacting via three forces: electromagnetism, the strong interaction and the weak interaction, each of which are carried by gauge bosons. The electromagnetic and weak forces are further unified to form the electroweak interaction. The force of gravity is not described by the standard model, however it has negligible effects on particle physics scales. Efforts to unify gravity and the standard model have been plagued by difficult theoretical problems. C, P and T refer to the operations of reversing charge, parity or time

$$\hat{C}: q \to -q \tag{2.1}$$

$$\hat{P}: \vec{r} \to -\vec{r} \tag{2.2}$$

$$\hat{T}: t \to -t. \tag{2.3}$$

For example, to transform a wavefunction with the Parity operator \hat{P} , one must invert all spatial coordinates $\vec{r} \rightarrow -\vec{r}$, for example,

$$\psi(\vec{r},t) \to \hat{P}\psi(\vec{r},t) = \psi(-\vec{r},t). \tag{2.4}$$

If a system is said to be invariant under or to conserve a symmetry, the observables of that system must not change under the transformation.

Additionally, each of these operators can be shown to be Hermitian and therefore they correspond to some observable property of the system. Furthermore, if the interaction Hamiltonian commutes with these operators, then the observable will also be a conserved quantity. If so, then if the wavefunction is an eigenstate of the Hamiltonian, then it must also be an eigenstate of the operator. In the case of parity, it can be shown that the eigenvalues of \hat{P} are ± 1 . For Dirac particles, the particles and antiparticles must have opposite parity. By convention, particles are denoted to have parity +1 and antiparticles parity of -1. From quantum field theory, it can be shown that the vector bosons γ , g, W^{\pm} and Z have parity -1[1].

It is possible to construct combined transformations, the most common examples of which are CP and CPT. The transformation CP is the combination of the C and P operators, and also represents the transformation to convert a system of particles to a system comprised of antiparticles. Under the CPT theorem, the standard model (and almost all reasonable variants) must conserve the combined symmetry CPT [1]. As long as the CPT theorem holds, T must be equivalent to CP.

The first discovery of the violation of one of these fundamental symmetries was in an experiment observing β decay in ⁶⁰Co by C.S. Wu et al in 1957 [2]. Wu observed a clear angular bias in the direction of the emitted electron. This inspired a development of the Fermi theory of beta decay [3], to one proposed by Lee and Yang [4]. The Lee-Yang theory was originally proposed to explain the θ - τ puzzle - the θ^+ and τ^+ particles were experimentally identical, but had different primary decay modes, to $\pi^+\pi^0$ and to $\pi^+\pi^+\pi^$ respectively. As these final states had opposite parity, the θ and τ could not be the same particle if parity were conserved. Lee and Yang proposed that they were the same particle, but that the beta decay process violated parity. This model proposes that the weak interaction has both a vector and axial vector component. Vector quantities change sign under parity, but axial vectors do not, therefore a mixture of the two quantities will change under \hat{P} , leading to a \hat{P} violating decay rate. It was determined that the axial vector coupling g_A was -1 times the vector coupling g_V , leading to this model to be named the V-A theory [5, 6]. This model also proposes the existence of the W^{\pm} mediator bosons. This theory eventually evolved into the full electroweak theory within today's standard model [7].

To maintain a uniform coupling constant for weak interactions involving quarks and those involving leptons, and to explain the slower-than-expected decays for particles with strangeness, Cabbibo proposed a modification where the weak eigenstates of quarks were not the same as the mass eigenstates [8]. In order to correctly predict the branching ratio of $K_L^0 \to \mu^+ \mu^-$, the GIM mechanism was proposed [9], which added the charm quark.

Up to this point, it had been noted that the combined symmetry CP appeared to be conserved by this theory. However, in 1964 CP violation was detected in the decays of neutral kaons [10]. There are two neutral kaons, K^0 and $\overline{K^0}$, which are nominally comprised of $d\bar{s}$ and $s\bar{d}$ quarks respectively. They are allowed to mix via the weak interaction. Therefore, they propagate as a combined state. In practice two distinct states are observed, being very similar in mass but very different in lifetime, the K_S , with a lifetime of $0.9 \times$ 10^{-10} s, and the K_L , with a lifetime of 0.5×10^{-7} s. If CP were a symmetry of the weak interaction, then each of those states would be eigenstates of CP. In practice the K_S decays almost entirely to two-pion states $\pi^+\pi^-$ or $\pi^0\pi^0$ and the K_L to three-pion states $\pi^+\pi^-\pi^0$ or $\pi^0\pi^0\pi^0$, or semileptonic states. It can be shown that the two-pion states have CP of +1, while the three-pion states have a CP state of -1. If CP is conserved, then this means that $K_L \to \pi\pi$ is not allowed. However, this was observed, meaning that CPis violated somewhere. There are two solutions: either the mass eigenstates K_S and K_L are not equal to the CP eigenstates K_1 and K_2 (indirect CP violation), meaning the CPviolation occurs in the mixing process, or the decay itself $K_2 \rightarrow \pi\pi$ occurs, violating CP in the decay process itself (direct CP violation). In this example, indirect CP violation is the dominant process, but direct CP violation has also been observed, first by the NA31 experiment at CERN[11].

This was eventually reconciled with theory by Kobayashi and Maskawa, who proposed a third generation of quarks, which allowed the addition of a CP violating phase to the CKM matrix [12].

Since the original observation in the Kaon system, CP violation has been observed

in many systems, including in neutral B mesons in 2001 by the Babar [13] and Belle[14] experiments. The LHCb experiment has found a hint of CP violation in the D meson system, at a significance of 2.7σ [15]. Continuing to probe CP violating physics is an ongoing goal at the LHC experiments. So far, observations are fully consistent with CP violation being fully explained by the single CP violating phase in the Kobayashi-Maskawa mechanism¹ [17].

In the QCD sector of the standard model, there is an additional allowed CP violating term in the Lagrangian, often called the QCD-theta term[18, 19],

$$\mathcal{L} = \frac{g_s^2}{32\pi^2} \bar{\theta} G^a_{\mu\nu} \tilde{G}^{a\mu\nu}.$$
(2.5)

However, the *CP* violating phase $\bar{\theta}$ has been constrained to be extremely small, $|\bar{\theta}| \leq 10^{-10}$ by measurements of the neutron [20, 21] and mercury [22] electric dipole moments, where it is naturally of order 1. The smallness of this phase is considered a major fine-tuning problem in the standard model, and there are several extensions to the standard model, including the Peccei-Quinn model [23–27]. This topic is explored more thoroughly later in this thesis (chapter 6), where new limits on the couplings of axions (particles arising as a result of the Peccei-Quinn symmetry) are derived, in the scenario where they also constitute a major component of the local dark matter concentration.

2.2 Baryon Asymmetry and Baryogenesis

It can readily be observed that our universe contains a vast overabundance of matter over antimatter, both in our local vicinity and across space at large[28]. Sakharov proposed in 1967 three requirements for an interaction to produce a matter-antimatter asymmetry [29]. These are:

- Violation of baryon number B
- Violation of C and CP symmetry
- Interactions out of thermal equilibrium

At first sight the standard model satisfies these conditions, with B violation possible through sphaleron processes, C and CP violation well established in the weak interaction and the conditions far from thermal equilibrium available during phase transitions in the

¹It should be noted that there a hint of CP violation in the neutrino sector, with the T2K experiment having found a 2.0σ confidence interval for the CP violating phase in the neutrino mixing matrix not including the CP-conserving values of 0 or π [16]. This is distinct from the CKM phase.

early universe. The particular scenario where the baryon asymmetry was created during the out-of-equilibrium conditions at the time of the electroweak phase transition is referred to as electroweak baryogenesis [30]. However, the need for new physics to enable this soon became clear.

First, electroweak baryogenesis is only possible if the electroweak phase transition is a strong first order transition, which is only true for values of the Higgs mass less than around 80 GeV [31, 32], which was ruled out first by LEP limits[33] requiring a Higgs mass over 114 GeV, followed by the discovery of the Higgs by the LHC experiments ATLAS[34] and CMS[35], at 125 GeV[17]. This provides a powerful motivation for a modified Higgs sector, to allow for a first order phase transition at the measured Higgs mass. Ongoing measurements and analyses at the ATLAS and CMS detectors at the LHC and later at the HL-LHC searching for additional Higgs sector particles and precisely measuring and comparing the properties of the Higgs to those in the standard model aim to thoroughly probe this frontier[36–38]. Such searches are complementary to EDM searches and to experiments searching for charged lepton flavour violation (CLFV) such as MEG[39] and Mu2e[40].

A further deficiency of the standard model in explaining baryogenesis is that the CKM phase (the only source of CP violation in the standard model) does not provide nearly enough CP violation to explain the size of the observed matter-antimatter asymmetry, failing by orders of magnitude[41]. Thus, some additional source of CP violation is required to enable baryon-asymmetry-generating processes. A generic feature of many theories of beyond the standard model physics, including supersymmetric theories, is the addition of new many potential CP violating phases, many of which are suitable to drive the creation of the baryon asymmetry [42–45]. Many new CP violating phases produce EDMs at a one or two loop level, resulting in large predicted EDMs, meaning EDM measurements already constrain the building of BSM models, which need to produce enough CP violation to enable baryogenesis while simultaneously producing only small EDM contributions.

An alternative hypothesis is that the baryon asymmetry was produced thanks to CP violating physics in the lepton sector, so called leptogenesis [46–48]. There is a hint towards the existence of substantial CP violation in the neutrino sector from the T2K experiment[16]. The most common leptogenesis scenarios require that neutrinos are Majorana, which would be validated if neutrinoless double beta decay were to be discovered. These scenarios are being probed by a variety of ongoing experiments in the neutrino sector, and also have implications for EDMs[49, 50].

In summary, the need to explain the production of the baryon asymmetry provides a powerful motivation for a modified Higgs sector to produce a strong first order EWPT and new sources of CP violation to drive the creation of the baryon asymmetry when constructing new models of physics going beyond the standard model. EDM measurements (among many other experimental results) provide tight constraints on such BSM physics theories.

2.3 The Neutron Electric Dipole Moment

It is well known that the neutron, like many particles, has a finite magnetic moment [51– 53]. This is aligned precisely along the spin axis \vec{s} , as the only available axis in a Dirac particle. One might also expect that it may have an electric dipole moment. An electric dipole moment of a fundamental particle can be analogised to that of an asymmetric charge distribution $\rho(\vec{r})$ in classical physics, defined by

$$\vec{d} = \int_{V} \vec{r} \rho(\vec{r}) d^3 \vec{r}.$$
(2.6)

In the case of a Dirac spin 1/2 particle, this must be aligned with the spin, like the magnetic moment. The Hamiltonian for a particle with magnetic dipole moment μ and electric dipole moment d in electric and magnetic fields can be written

$$H = \mu \frac{\vec{s}}{|s|} \cdot \vec{B} + d \frac{\vec{s}}{|s|} \cdot \vec{E}.$$
(2.7)

Under the discrete \hat{C} , \hat{P} and \hat{T} symmetries, the components of this expression change as such:

	\hat{C}	\hat{P}	$\hat{T} \ (\equiv \hat{C}\hat{P})$
\vec{r}	\vec{r}	$-\vec{r}$	$-\vec{r}$
\vec{s}	$-\vec{s}$	\vec{s}	$-\vec{s}$
\vec{E}	$-\vec{E}$	$-\vec{E}$	$ec{E}$
\vec{B}	$-\vec{B}$	\vec{B}	$-\vec{B}$

Table 2.1: Transformation of some common quantities under \hat{C} , \hat{P} and \hat{T} .

As such, under parity, the Hamiltonian becomes

$$H = \mu \frac{\vec{s}}{|s|} \cdot \vec{B} - d\frac{\vec{s}}{|s|} \cdot \vec{E}, \qquad (2.8)$$

and under time-reversal it becomes

$$H = -\mu \frac{\vec{s}}{|s|} \cdot -\vec{B} - d\frac{\vec{s}}{|s|} \cdot \vec{E} = +\mu \frac{\vec{s}}{|s|} \cdot \vec{B} - d\frac{\vec{s}}{|s|} \cdot \vec{E}.$$
 (2.9)

Thus, a nonzero electric dipole moment violates both P and T, and by extension CP.

2.4 EDMs of other systems

The neutron is only one system among many attractive targets for EDM searches. While nEDM experiments search for an electric dipole moment of the free neutron, other experiments search for the EDMs using atoms, molecules and in storage rings. Atomic and molecular searches fall generally into two catagories: paramagnetic and diamagnetic. Experiments using paramagnetic atoms (such as ¹⁹⁹Hg) aim to measure the nuclear EDM, which includes contributions from the neutron and proton EDMs as well as CP violating interactions at the nuclear level, while experiments in diamagnetic systems such as ThO aim to make use of the high electric fields in atoms and molecules to gain enhanced sensitivity to the electron EDM. Experiments using storage rings are able to measure the EDMs of muons, the proton and of light bare nuclei directly[54–57].

Given the cost and difficulty of obtaining relatively low densities of ultracold neutrons even compared to many rare isotopes, and comparing the huge difference between the sensitivities of the most recent neutron and mercury EDM experiments $(|d_{199}H_g| < 7.4 \times 10^{-30} e$ cm compared to $|d_n| < 10^{-26}$ e cm) it may seem that measuring nuclear EDMs is clearly superior to the neutron. However, the most recent limits on the neutron and mercury EDMs actually provide comparable sensitivity to BSM physics. Measurements of the mercury EDM are performed on atoms. According to the Schiff theorem [58-60], EDMs of neutral atoms arising from a nuclear EDM are zero classically, and only relativistic effects [61] and the effects arising from the finite size of the nucleus [62] violate this. Experiments using 225-radium atoms 63] are an exciting prospect, as the large octopole deformation of the nucleus [64] leads to an enhancement of these effects [65-67]. One additional difficulty in EDM searches in nuclei, is that interpretation of the measured values is difficult theoretically. While the atomic physics calculations are well understood, relating the nuclear EDM to the underlying CP violating effective interactions still requires complicated many-body nuclear physics calculations [68]. The relation of the measured values of the neutron and proton EDMs to the underlying physics is still nontrivial, but substantially easier.

Contrastingly, the interpretation of leptonic EDMs is substantially easier. The electron EDM is measured typically as part of an atomic or molecular system[69, 70]. These systems are chosen such that a relatively small applied electric field results in a huge internal electric field, magnifying the effect of an electron EDM[71, 72]. Additionally, these experiments are sensitive to a CP violating quark-electron interaction[68], though the results are most

often quoted as the equivalent pure electron EDM, neglecting this possible interaction. It is possible for this interaction and an eEDM to offer some cancellations, allowing constraints on the smallness of the eEDM to be evaded [73, 74]. The muon EDM is most often measured directly in storage ring experiments. The most sensitive measurement to date [57] ran in parallel with the g - 2 measurement [75] in the storage ring at BNL, and a new, more sensitive, measurement is planned using an upgraded apparatus at Fermilab [56]. The absolute numerical sensitivity of the muon EDM is substantially less than for the electron. In minimally flavour violating scenarios, the muon EDM expected to follow $d_{\mu} = \frac{m_{\mu}^2}{m_e^2} d_e$, however new physics theories can predict large enhancements for the muon EDM over the electron EDM beyond this [76].

It is useful to measure the EDMs of a variety of systems for several reasons. First, EDM experiments are difficult precision experiments, which can have large systematic effects which are nontrivial to compensate. It is likely that the discovery of a non-zero EDM would need to be confirmed by a second experiment before it would be taken seriously by the wider physics community. More importantly however, the EDMs of different systems depend on different combinations of the underlying physical interactions. The discovery of a non-zero EDM would set the scale for new physics, but would not indicate the source without comparison to other measurements. It is also possible that in new physics theories, there may be more than one interaction which leads to EDM contributions. For example, [74] proposes that a CP-violating quark-electron coupling gives a contribution to the measured electron EDM, evading eEDM constraints on new physics theories. To this end, a global fit has been performed linking several EDM measurements to find limits on the strengths of the underlying effective interactions, allowing for more than one source to be present[73].

One might raise the immediate objection that many common molecules have extremely well measured electric dipole moments, for example ammonia (NH₃). Taking the molecule to be a rigid body, it is clear that the dipole must lie along the symmetry axis of the molecule. However, in these molecules, there are states with opposite rotational angular momentum, as measured by the projection onto the symmetry axis. These are at first sight degenerate, however atomic physics effects lift the degeneracy, meaning that in the nondegenerate states of fixed angular momentum both 'directions' of EDM state exist in equal amounts, meaning that the molecule has no net EDM [72, 77].

2.5 Standard Model predictions for the nEDM

2.5.1 Weak Sector

The first diagrams which could generate a neutron EDM from the CKM phase arise at 2-loop level. However, it can be shown that when summing over all diagrams, these contributions all cancel exactly[78]. Therefore, the first non-zero contributions arise only at three loop order, and are therefore extremely suppressed. At this point there are many possible contributions, and the calculation is complex. Estimates for the nEDM in the standard model have ranged from 10^{-29} e cm to 10^{-34} e cm over the years[68, 79]. The current best value is calculated in [80], at somewhere between $(1 - 6) \times 10^{-32}$ e cm, far below current experimental sensitivities and common predictions arising from many BSM theories. For this reason, the neutron EDM (among other EDMs) is often referred to as a 'background-free' observable.

2.5.2 Strong *CP*

The *CP* violating phase in the strong interaction, contrastingly, can generate extremely large EDMs. Measurements of the neutron and mercury electric dipoles have constrained the *CP* violating phase to be $|\bar{\theta}| \leq 10^{-10}$, where it would naturally be of order 1. This presents a fine-tuning problem in the standard model, referred to as the strong CP problem. Several mechanisms have been proposed to explain the smallness of this parameter, including the Peccei-Quinn theory, which implies the existence of axions. This term and the strong CP problem are explored more fully in chapter 6.

2.6 Relation of EDMs to Underlying Physics Through Effective Field Theories

While most interactions in particle physics could be expressed at the lowest levels of quarks, leptons and gauge bosons, in practice at lower energy scales a simplified model (known as an Effective Field Theory, or EFT) involving only the particles accessible at such energy levels is often more suitable. One example of an effective field theory is the Fermi theory of beta decay for the reaction $n^0 \rightarrow p^+ + e^- + \nu_e^0$ or $\mu^- \rightarrow e^- + \bar{\nu}_e + \nu_{\mu}$. In the full standard model description this is mediated by a (relatively very heavy) W⁺ boson, however this deeper structure was not known at the time. The Feynman diagrams for these two processes are compared in Figure 2.1 Therefore, in the Fermi theory the

reaction was modelled as a point interaction of the four participants. This served as a satisfactory model for many years, solving the apparent energy non-conservation problem by adding the neutrino.



Figure 2.1: Comparison of Feynman diagrams for the Fermi theory, and the underlying high-energy process. From [1].

Analogously, it is possible to determine the effects of fundamental interactions at the scale of TeV and PeV on nuclear scale (GeV) measurements such as EDMs via an effective field theory. At 1 GeV we can consider all fields other than the u, d and s quarks, the photon, muon and electron to be heavy and integrate them out. We can construct a finite set of all CP odd operators up to dimension 6 accessible at 1 GeV. We arrive at a series of effective operators at the scale of the neutron.

The only dimension 4 example is the QCD Theta term, mentioned in subsection 2.5.2, which will be further discussed later in chapter 6.

To extend the standard model and add new effective interactions in a relatively modelindependent way, a set of dimension 6 operators which violate CP and lead to EDMs can be added. These include the quark and electron EDMs (arising from diagrams which have been integrated out) and the quark chromo-EDMs (analogous to a normal EDM, but for the strong force). Several further interactions, such as the Weinberg 3-gluon operator, different 4-quark operators and quark-electron couplings can lead to an effective CP violating pion-nucleon coupling, which can lead to so-called 'long range' contributions to nucleon EDMs, and which are an important part of the calculation of nuclear EDMs. See Figure 2.2 for an explanatory schematic diagram illustrating the role of several effective operators as a general intermediary between model-independent fundamental physics and observable EDMs.

A wide variety of theoretical techniques are available to link the values of these operators with the observable EDM values of hadrons, atoms and nuclei. See [68] and references therein for an excellent quantitative summary and reference.



Figure 2.2: Schematic diagram illustrating the parameters coming into play to compute different EDMs at different scales from the underlying BSM physics. Dashed lines indicate a weaker dependence. In this notation, d_e is the electron EDM, C_{qe} and C_{qq} are various 4-fermion couplings (quark-electron or quark-quark), θ is the QCD $\bar{\theta}$, w is the Weinberg 3-gluon operator and d_q and \tilde{d}_q and the quark EDMs and chromo-EDMs. From these, operators at the nuclear scale $g_{\pi NN}$ (CP-violating pion-nucleon couplings) and $C_{S,P,T}$ (possible electron-nucleon couplings) can be derived and used to compute atomic EDMs. From [79].

2.7 Historic and Competing nEDM Measurements

Purcell and Ramsey first proposed to measure a neutron EDM in 1950[60], to test for parity violation. The first dedicated measurement of the neutron EDM was performed in 1951 by Smith, Purcell and Ramsey, but not published until 1957[81], spurred by the discovery of P violation in the beta decays of ⁶⁰Co nuclei by Wu et. al. [2]. This measurement was performed by applying the Ramsey technique[82] to a polarised neutron beam travelling through a uniform magnetic field, and searching for an electric-field correlated frequency shift. This gave a result of $d_n = (-0.1 \pm 2.4) \times 10^{-20}$ e cm. A series of experiments based on a similar beam-based Ramsey measurement concluded with the measurement of Dress *et al.* published in 1977, which gave a final measurement of $d_n = (0.4 \pm 1.5) \times 10^{-24}$ e cm[83], over a factor 10,000 increase in sensitivity. These beam-based measurements were ultimately limited by systematic errors. As neutrons travel through an electric field \vec{E} with velocity \vec{v} , relativistic effects cause them to experience a motional magnetic field in

their rest frame of $\vec{B}_m = -\frac{1}{c^2} \vec{v} \times \vec{E}$ [84]. If the electric field is not perfectly perpendicular to the magnetic field, this motional magnetic field will add with the main *B* field, causing an *E*-field correlated frequency shift, mimicking the effect of a true EDM. This systematic effect proved to be the limiting factor in the last beam experiment, motivating a shift to a new technology.

To avoid these systematic effects, the community shifted towards measurements using stored ultracold neutrons (UCN). The first experiment of this type achieved a result of $d_n = (4\pm7.5)\times10^{-24}$ e cm in 1980 by a group based at LNPI [85]. Since then, experiments by that group and by the Sussex-RAL-ILL collaboration have improved by over two orders of magnitude in sensitivity to reach the latest limit of $d_n < 3.0 \times 10^{-26}$ e cm (90% CL) first published in 2006, and revised in 2015 [20, 21, 86–88].

After the end of data taking, the apparatus used by the Sussex-RAL-ILL collaboration was transferred to a new experimental collaboration, intending to use it as a basis upon which to develop a novel room temperature stored UCN apparatus [89]. After several years of development with the apparatus installed at ILL, the apparatus was moved to the new UCN source at the PSI[90], with the aim to increase the statistical sensitivity by achieving greater UCN densities. During the upgrade process, almost every part of the apparatus, with the exception of the vacuum tank and the four layer μ -metal magnetic shield, was replaced and upgraded. While the increase in neutron statistics was less than hoped, improvements in all aspects of the apparatus have resulted in a substantial improvement in statistical sensitivity per day, as well as in the control of systematics. The apparatus took EDM data from August 2015 to December 2016, in which time a statistical sensitivity of approximately 1.1×10^{-26} e cm (1σ) was accumulated, representing an improvement on the Sussex-RAL-ILL experiment [91]. After unblinding, this result is expected to set the world's best limit. The experimental setup and analysis of the resulting data is the topic of this thesis.

The PSI nEDM collaboration has since begun construction of the next-generation n2EDM apparatus [92]. This apparatus is partially inspired by the dual chamber design of [85]. It aims to improve upon the statistical sensitivity of the previous generation apparatus in several aspects. Going to a double chamber design with chambers of a substantially increased radius will improve on the total UCN statistics, as will optimisation of all aspects of the UCN guide geometry to the PSI UCN source[93]. Improved control of magnetic fields will result in improved neutron polarisation. Finally, improved high voltage design will allow greater electric field strengths to be achieved. With the current

'baseline' design, using only known and tested techniques and technology, a final sensitivity approaching 1×10^{-27} e cm is envisaged. Further upgrades have the possibility to bring this sensitivity into the 10^{-28} e cm range.

Several other experiments also attempt to measure the neutron EDM experiment using a double chamber apparatus at room temperature. The Japanese-Canadian UCN collaboration aims to construct a double chamber apparatus at the TRIUMF superfluid ⁴He UCN source, with both a ¹⁹⁹Hg (as used in the Sussex-RAL-ILL and PSI experiments) and ¹²⁹Xe comagnetometer. This experiment proposes to reach a statistical sensitivity of 1×10^{-27} e cm after 100 days of beam time, based on simulation [94]. Another collaboration seeks to build an experiment at the upgraded LANL UCN source, which has the possibility to reach a statistical sensitivity of 2.1×10^{-27} e cm for five calendar years of running, based on demonstrated source performance [95]. The FRM-II/PanEDM experiment was originally developed at the FRM-II reactor, but following regulatory difficulties was moved to the ILL. It is connected to the SuperSun source [96] and in its first iteration intends to reach a sensitivity of 10^{-27} e cm, with a second iteration achieving an improvement of a factor 3-4 by use of the SuperSun magnetic reflector. There is also the possibility for a 50% improvement in electric field if the comagnetometer can be omitted. A new UCN source is under development at the WWR-M reactor [97] which promises to offer high UCN densities; it is planned to run the previous PNPI-ILL experiment[98] at this source to improve the statistical sensitivity of this apparatus.

An innovative technique was developed to measure the neutron EDM by taking advantage of the large electic fields available within certain crystal structures. Test measurements using this technique resulted in an nEDM measurement of $d_n = (2.4 \pm 6.5 (\text{stat.}) \pm 5.5 (\text{syst.})) \times 10^{-24} \text{e}$ cm [99], and is has been proposed that further measurements using this technique have the possibility to reach a sensitivity of the order a few $\times 10^{-27} \text{e}$ cm[100].

Following the success of the room temperature experiment published in 2006, the Sussex-RAL-ILL collaboration embarked on an ambitious effort to produce an nEDM experiment with the entire UCN production, NMR measurement and detection taking place submersed in superfluid liquid helium at 0.7 K [101]. This would allow the experiment to achieve an order of magnitude higher electric field, and extremely high neutron densities, with a statistical reach on the level of 10^{-28} e cm. Unfortunately, the experiment was plagued by technical difficulties and discontinued before taking any useful data.

Undeterred, a collaboration centred on the Spallation Neutron Source (SNS) at Oak Ridge National Laboratory (ORNL) [102] are developing an experiment based on a similar
concept of performing a nEDM search in superfluid helium [103]. This experiment will use an innovative detection technique relying on looking for a beating in the (spin-dependent) UCN capture on cohabiting ³He atoms, whose spins are dressed to almost exactly match the neutron precession frequency [104].

In recent years, there has also been a renewed interest in performing an nEDM experiment using a cold beam of neutrons [105]. It is proposed to counter the systematic effects using a time-of-flight technique on a chopped input beam, which is then extrapolated down to a zero-velocity, systematic free measurement. This technique has the possibility to reach high statistical sensitivities of 5×10^{-28} e cm by taking advantage of the extremely high cold neutron flux available at pulsed spallation sources, such as the ESS [106]. Several test experiments have been completed during beamtimes at cold neutron sources such as those at the ILL and at SINQ (PSI).

2.8 Ultracold Neutrons

The final section of this chapter will be devoted to the theory and properties of ultracold neutrons (UCN), the key experimental tool used in most modern neutron EDM experiments. Ultracold neutrons are extremely low energy neutrons, defined as having energies of up to 250 or 300 neV. This corresponds to velocities of up to a few metres per second, or a temperature of around 3.5 mK. Their most valued property is that they can be stored for long periods of time in macroscopic bottles made of relatively ordinary materials for several hundreds of seconds.

Their ability to be stored lies in the properties of their strong interaction with matter. The force between a free neutron and a free proton can be approximated by a square well with a depth around 40 MeV and a radius of around 2 fm [107], effectively a forest of delta functions for such low energy neutrons. However, as the spacing between nuclei is extremely large compared to the wavelength of the neutrons, the potential will be averaged out. It is therefore possible in many cases to instead model the interaction between a free neutron and a material through which it passes as a potential step. This potential may be positive or negative, corresponding to a repulsive or attractive interaction. Fermi potentials for typical materials used in UCN storage range between +335 neV for 58 Ni (which may be used as a coating on the inside of glass neutron guides) to +95 neV for the quartz used for the insulator ring of the Sussex-RAL-ILL nEDM experiment. The effective Fermi potential depends on the bound coherent scattering length a, number density N

and neutron mass $m_{\rm n}$ as [107]

$$V_f = \frac{2\pi\hbar^2}{m_{\rm n}} Na. \tag{2.10}$$

When a neutron reaches a boundary between two media, the interaction can be modelled using as transmission or reflection at a potential barrier. Neglecting capture and upscattering losses, reflection will occur when the neutron kinetic energy in the direction perpendicular to the surface is less than the difference between the Fermi potential of the two materials. In the case where the total neutron kinetic energy is less than the Fermi potential difference, the neutron will always be reflected. However, typical surfaces are generally rough. In this case, the reflection may not be totally specular (i.e. with momentum parallel to the surface remaining unchanged, and perpendicular being exactly reversed), rather it may be *diffuse*. In UCN guiding tubes, specular reflection is highly desirable to maximise transmission (thus minimising the number of neutrons reflected back towards the source, or which are reflected in such a way that their next collision occurs at greater than the critical angle, resulting in loss). However, since UCN storage chambers are typically filled from a single small port, the initial velocity distribution will be far from isotropic. Therefore, in order to ensure that the population quickly equilibrates such that the assumption of isotropic velocities holds when analysing the experimental data, some roughness is often desirable in UCN traps.

At each reflection, there is a finite chance for the neutron to be lost. This may be due to capture on nuclei in the wall, or inelastic scattering. In the quantum mechanical description of reflection at a finite potential, the neutron wavefunction penetrates a little way into the material. This allows for two potential loss mechanisms: capture on a nucleus, and inelastic upscattering from phonons in the material. Losses can be modelled by adding an imaginary component to the potential U = V - iW. The imaginary component W can be found as

$$W = \frac{\hbar}{2} \sum_{i} N_i \sigma_l^{(i)} v, \qquad (2.11)$$

where N_i is the number density of each nuclear species i, $\sigma_l^{(i)}$ is the loss cross section and v is the incident velocity. The depth of the penetration depends on the ratio of the incident neutron energy (considering only the direction perpendicular to the surface), and the Fermi potential of the material. Another useful quantity is the ratio $f = W/V_f$, given by

$$f = \frac{W}{V_f} = \frac{\sigma_l k}{4\pi a} = \frac{\sigma_l}{2a\lambda},\tag{2.12}$$

where k is the neutron wavenumber, or λ the wavelength. It can be computed that the loss probability per reflection for a neutron with total kinetic energy E incident at an angle θ is

$$\mu(E,\theta) = 2f\left(\frac{E\cos^2\theta}{V - E\cos^2\theta}\right).$$
(2.13)

Note that at higher UCN energies the loss probability per bounce is higher, and the rate of bounces is higher. This leads to a higher loss rate for higher energy UCN, meaning that during typical storage experiments the energy spectrum will soften (be shifted towards lower energies) over time. In addition, any regions with defects in the quality of the coating or other issues will typically show a reduced Fermi potential or higher loss probability per bounce, causing UCN losses which typically affect higher energy UCN preferentially, further softening the spectrum.

Ultracold neutrons are subject to beta decay via the weak interaction. They have a half life of 880.2±1.0 s as given by the PDG average [17]. However, this number is much debated with measurements differing from each other by over 5σ . There are two complementary methodologies: 'beam' and 'bottle'. Beam type experiments take a cold neutron beam of well-known flux, pass it through a detector, and count the protons or electrons resulting from the decay, while bottle type experiments store ultracold neutrons for different times, and count how many remain. Beam experiments typically produce results of around 888 s, while bottle experiments generally report lower numbers closer to 880 s. Both experiments are subject to large systematic effects, the incorrect assessment of which could explain the deficit. However, the discrepancy has also inspired alternative hypotheses, such as a decay to dark matter particles [108], or an oscillation into a hypothetical *mirror neutron*[109]. The nEDM apparatus at the PSI has been used for two sensitive searches for this phenomenon [110, 111].UCN are also the subject of several experiments to determine the neutron beta decay correlation coefficients [112].

While neutrons have no net charge, they do have a well measured magnetic moment [53], comparable in size to the proton magnetic moment. As a result of this, there is an interaction with a magnetic field of

$$V_{\text{mag}} = -\vec{\mu} \cdot \vec{B} \approx 60 \text{neV} / \text{T} \frac{\vec{s}}{|\vec{s}|} \cdot \vec{B}.$$
(2.14)

As a result of this, in an inhomogeneous magnetic field a force will be exerted upon the neutron

$$F_{\text{mag}} = -\nabla V_{\text{mag}} = \pm |\mu_{n}|\nabla|\vec{B}|.$$
(2.15)

This formula assumes the adiabatic case, where the gradients of the field are gradual and thus the neutron spins follow the same relationship with the direction of the field. This leads to two classes of UCN, depending on their spin state: 'high-field seekers' and 'low-field seekers', each being drawn to areas of maximal or minimal magnetic field respectively. This effect is used to polarise beams of ultracold neutrons using magnetic fields, either using strong magnetic fields created by a superconducting magnet solenoid [113], or more traditionally by using thin iron foils which greatly amplify an applied magnetic field to produce an equivalent barrier[107]. A superconducting magnet polariser with an aluminium foil of 100 µm is able to almost fully polarise a beam using a field strength of 5 T (= 300neV)[113].

Ultracold neutrons are unusual among systems often studied in particle physics in that gravity is a very significant influence on their motion. Writing the interaction with the Earth's gravitational field as

$$V_{\rm grav} = mgh, \tag{2.16}$$

and using the standard value for the neutron mass [17], one finds the rule of thumb that 1 cm of height gained equates to 1.03 neV of energy. Considering typical UCN energies of up to 250 neV, it is therefore clear that in experiments with dimensions measured in metres this is a non-negligible effect. In order to optimise the flow of UCN to, from and within experiments, in general Monte Carlo tools must be used [93]. Even when stored in chambers of limited size, such as those used in the PSI and previous Sussex-RAL-ILL experiment (of 12 cm height), the effect of gravity shifts the centre of mass of the stored neutrons approximately 3 mm below the centre of the chamber [114], causing a frequency shift from what might naïvely be expected in the presence of vertical gradients. The differing centres of masses of different UCN energy populations causes additional depolarisation effects [115–117]. The measurability of gravitational effects on UCN has spurred several experiments to investigate the quantum levels of UCN-gravity bound states such as GRANIT [118] and qBounce [119].

Chapter 3

The Search for the Neutron Electric Dipole Moment at the Paul Scherrer Institute

3.1 Introduction and Principle of Operation



Figure 3.1: Sketch showing the layout of key components of the nEDM experiment at the PSI, described in turn throughout this chapter. From [91].

The PSI nEDM experiment aimed to measure the neutron electric dipole moment by searching for a shift in the Larmor frequency of stored ultracold neutrons upon the application of an electric field. A simplified schematic of the major components of the experiment is presented in Figure 3.1, each of which will be in turn explained in more detail in the course of this chapter. The experiment is located at the ultracold neutron source at the Paul Scherrer Institute, Switzerland, which is depicted in an aerial photograph in Figure 3.2.

The Hamiltonian of a particle with electric and magnetic dipole moments in an electric and magnetic field is

$$\mathcal{H} = -d\frac{\vec{s}}{|s|} \cdot \vec{E} - \mu \frac{\vec{s}}{|s|} \cdot \vec{B}, \qquad (3.1)$$

where d is the EDM, μ is the magnetic dipole moment, and \vec{s} is the spin vector. Such a particle, of spin 1/2, will precess with a frequency of

$$f_{\rm n} = \frac{2}{h} |\mu_{\rm n} B \pm d_{\rm n} E|. \qquad (3.2)$$

By measuring the neutron Larmor frequency in configurations of parallel $(f_n^{\uparrow\uparrow})$ and antiparallel $f_n^{\uparrow\downarrow}$ magnetic and electric fields, one can extract the neutron electric dipole moment as

$$d_{\rm n} = \frac{1}{4E} \left(h(f_{\rm n}^{\uparrow\uparrow} - f_{\rm n}^{\uparrow\downarrow}) - \mu_{\rm n}(B^{\uparrow\uparrow} - B^{\uparrow\downarrow}) \right).$$
(3.3)

As such, it is critical to measure the neutron precession frequency as accurately as possible, and to measure any changes in the magnetic field as accurately as possible. The neutron precession frequency is measured using the Ramsey technique, while changes in the magnetic field are primarily compensated by the mercury magnetometer system.

The Ramsey technique of separated oscillatory fields [82] was first proposed in 1950 by Norman Ramsey for measuring the precession frequencies of atomic and molecular beams, for which he won the Nobel prize. It was developed from the Rabi technique. In the Rabi technique [120], a polarised beam passing through a static magnetic field is subject to a transverse oscillating (or rotating) field close to the resonant frequency. The oscillating field causes the spins to begin to be flipped. The closer to the resonant frequency the spin flip is, the more effective it is. As such, by varying the RF frequency and recording the spin states of the particles in the final beam, the resonant frequency can be measured. In the Ramsey technique, a pulse close to the resonant frequency is applied for the appropriate time and amplitude for a $\pi/2$ rotation into the transverse plane. The spins are then allowed to precess freely for a set time, before a second $\pi/2$ pulse is applied, at the same frequency and in phase with the first. If the frequency of the pulses matches the spinprecession frequency of the particles, then the spin flip will be successfully completed, and (assuming no depolarisation) the particles will finish the sequence with opposite spins to



Figure 3.2: Aerial view of the Paul Scherrer Institute straddling the river Aare, looking north. The round building in the foreground is the SLS. The HIPA accelerator is located in the large white building behind it, and the UCN source, cooling plant and nEDM experiment are located in the attached building on the left.

the starting state. However, if the second pulse arrives with a π phase difference to the neutron free precession, the initial spin flip will be fully reversed, and the particles will be returned to the initial state. The behaviour of the neutron spins in case where the pulses are applied at precisely the neutron resonant frequency is illustrated in Figure 3.3 The final state for pulses close to the resonant frequency (within the first fringe, where the Rabi window effect is negligible) can be approximated by

$$N_{\uparrow}(\Delta\nu) = \frac{N_0}{2} \left[1 - \alpha \cos\left(\frac{\pi\delta\nu}{\Delta\nu}\right) \right]$$
(3.4)

$$N_{\downarrow}(\Delta\nu) = \frac{N_0}{2} \left[1 + \alpha \cos\left(\frac{\pi\delta\nu}{\Delta\nu}\right) \right], \qquad (3.5)$$

where N_{\uparrow} and N_{\downarrow} are the approximate number of particles in each final spin state, N_0 is the total number of particles, α is the Ramsey visibility and $\delta \nu = \nu_{\rm RF} - \nu_{\rm n}$ is the detuning of the pulses $\nu_{\rm RF}$ from the resonant frequency of the neutrons $\nu_{\rm n}$. The Ramsey visibility α is defined to be the final polarisation of the neutron ensemble, convoluted with the efficiency and specificity of the spin-sensitive neutron detection system. The linewidth $\Delta \nu$ is given by

$$\delta\nu = \frac{1}{2T + \frac{8t}{\pi}},\tag{3.6}$$

where T is the free precession time between pulses, and t is the length of each $\pi/2$ pulse. The visibility α is related to the polarisation of the neutron ensemble, including additionally detector efficiencies. The maximum achievable α is 1 for a fully polarised ensemble with perfect separation of the two spin states upon detection.

In the original conception of the Ramsey technique, it was envisioned that a beam of particles would pass first through a region with the oscillating field continuously applied, then pass through a region with only the main static magnetic field, before passing through a second region with the same oscillating field applied in phase with the first. However, in modern experiments using stored particles rather than beams, the separated oscillatory fields are separated in time rather than space: the oscillator runs freely and continuously, however a gate is applied to the signal such that the oscillating field is applied only when desired, but the phase of the second and first pulse are guaranteed to match. The Ramsey technique is often compared to 'a double slit experiment in time' and compared to interferomety and other techniques because it acts like an interference experiment between the frequency of the pulses and the precession frequency of the particles.

For illustrative purposes, the results of a typical fit of the Ramsey curve a few hours of typical data are shown in Figure 3.4. In 2017, a special run was performed exploring a broader part of the Ramsey resonance to validate the form of the Rabi envelope, far beyond the central fringe where normal datataking is performed. The results of this measurement are shown in Figure 3.5.

In order to maximise the statistical sensitivity to the central frequency of the fringe (i.e., the Larmor frequency), it is most effective to measure at the regions of maximum slope, i.e. close to $\delta\nu \approx \pm \frac{\Delta\nu}{2}$. In practice, it is useful to add a small detuning to give sensitivity to nuisance parameters such as the Ramsey visibility α from Equation 3.4 and Equation 3.5 [122]. This is illustrated in Figure 3.4. It can be shown that the statistical sensitivity of a series of m Ramsey measurements taken at the optimum working point when $T \gg T_{\rm RF}$ with Poisson counting statistics on the counts is [72]

$$\sigma_{\nu_{\rm n}} = \frac{1}{2\pi T \alpha \sqrt{N_0 \cdot m}}.\tag{3.7}$$

Combining this with Equation 3.3, we then find that the best achievable statistical sensitivity of an nEDM experiment is [72]

$$\sigma_{d_{n}} = \frac{\hbar}{2ET\alpha\sqrt{N_{0}\cdot m}}.$$
(3.8)

As nEDM experiments have historically been statistics limited, the design philosophy was to prioritise the statistical sensitivity, while maintaining sufficient control of all systematic effects. Each of the statistical factors must be balanced against one another to optimise the final sensitivity accumulated. For example, reducing the height of the chamber may increase the attainable electric field E, at the cost of decreasing the neutron counts per cycle N_0 , both due to the reduced volume and decreased storage time constant due to an increased rate of wall collisions. Increasing the free precession time T would decrease the Ramsey visibility α and the neutron counts N_0 , as well as possibly decreasing the achievable number of cycles per day.

Although the PSI nEDM experiment is extremely similar in principle to its successor, the Sussex-RAL-ILL experiment[20, 21, 88], it was improved in several areas to increase the statistical sensitivity. First, the experimental setup was moved to run at the PSI UCN source[90] rather than the ILL turbine [123]. This resulted in an increase in the neutron counts. The PSI source was recently shown to be the most intense UCN source in the world, in terms of producing as many storable UCN as possible [124, 125]. The development of improved UCN coatings [126] and switching from a quartz to a Rexolite ring for the chamber also improved neutron storage times. Substantially improved magnetic field homogeneity due to a novel optimisation algorithm [127] using the caesium magnetometer array significantly improved the achievable T_2 relaxation time, improving on the final Ramsey visibility α . A photo of the neutron precession chamber and other key components within the vacuum chamber is displayed in Figure 3.6.

The remainder of this chapter will be devoted to documenting the experimental subsystems of the apparatus that made such improvements possible, and finally to exploring the outlook to and sensitivity goals of the next generation n2EDM spectrometer.



Figure 3.3: Sketch of the neutron precession during a Ramsey measurement sequence at the resonant frequency. Adapted from [121]



Figure 3.4: Ramsey fit of the Asymmetry $A = \frac{N_{\uparrow} - N_{\downarrow}}{N_{\uparrow} + N_{\downarrow}}$ to a typical section of EDM datataking. Explanation of fit parameters indicated will be included in chapter 4.



Figure 3.5: The Ramsey fit of the asymmetry $A = \frac{N_{\uparrow} - N_{\downarrow}}{N_{\uparrow} + N_{\downarrow}}$ from a special measurement covering a larger section of the Ramsey resonance curve beyond the central fringe. Shown above are the results over the full measurement range, while below is a zoomed view of the central few fringes. Note that at the extremes, the Rabi envelope becomes clearly visible over the tight Ramsey fringe pattern.



Figure 3.6: View inside the vacuum tank, showing the precession chamber used for the nEDM experiment. The castellated high voltage feed at the top connects to the upper electrode, which is separated from the lower electrode by the insulator ring. The UCN enter and leave the chamber from below. The mercury probe light travels across the chamber through two windows in the insulator ring.

3.2 The PSI UCN Source

The UCN source at PSI [90] is a spallation source driven by the 590 MeV HIPA accelerator at the PSI, shown in Figure 3.7. A fast kicker magnet [128] is able to divert the entire (typically 2.2 mA, but up to 2.4 mA) beam current of the accelerator to the UCN source for up to 8 seconds, for a maximum duty cycle of 3%. The maximum average beam current on the UCN spallation target is 60 μ A. There is also the option to continuously split 10 μ A of beam current to the UCN source for testing purposes. An upgrade programme is planned to increase the HIPA beam current to 3.0 mA [129]. The target is comprised of 760 lead-filled zirconium tubes[130]. Around 8 neutrons are released per incident proton. These neutrons are then thermalised in a moderator volume consisting of 3300 l of heavy water (D₂O) at approx 31 °C [131].

The neutrons are converted to ultracold neutrons in a 30 l solid ortho-deuterium (D₂) converter at 5 K. The incident neutrons are downscattered by phonon interactions to become ultracold neutrons [132]. Upon leaving the crystal to the vacuum above, they gain around 100 neV energy from the Fermi potential of the deuterium crystal, which is lost as they rise up the guide. Then, they enter a large storage vessel of around 2 m³ in volume. The walls are coated with diamond-like carbon (DLC) to maximise UCN storage properties. After the spallation pulse, a flap at the bottom of the storage vessel is closed. A total of three guide ports leave the storage vessel: South, West 1 and West 2. South and West 1 leave from the bottom of the storage vessel to experimental areas South and West, while West 2, connected to the top of the chamber, is mostly used for diagnostics and monitoring. The nEDM experiment was installed permanently in area South for the entire duration of the experiment.

An experimental comparison of worldwide UCN sources was recently completed [125], using a portable UCN chamber [124]. The achievable UCN densities in that chamber and their storage time constants (giving information about the UCN spectrum) were measured. It was found that the PSI UCN source achieved the highest UCN densities after two seconds of storage, while also recording a longer storage time constant than achieved at the next best competitor, the ILL PF2 beamline with the turbine source [123]. The result of the comparison is presented in Figure 3.8.



Figure 3.7: Diagram of the PSI UCN Source, from [131].



Figure 3.8: Comparison of the PSI UCN source performance [131] with several competitors worldwide, from [125].

3.3 UCN Systems

The systems for guiding, storing and counting UCNs form the core of the experiment. Free ultracold neutrons are produced at the source in several-second long pulses, at a repetition period of several minutes. These must be polarised, bought into the precession chamber, stored for hundreds of seconds for the duration of the Ramsey sequence, and then the neutrons in each final spin state must be counted.

In general, the neutrons are transported from place to place within UCN guides. In our experiment, these are typically formed of smooth glass tubes with the inner surfaces sputtered with a coating of material with a high Fermi potential, such as an alloy of nickel and molybdenum, or isotopically pure ⁵⁸Ni. The high smoothness maximises the specularity of the reflections and the high Fermi potential maximises the chance for the neutrons to be successfully reflected at each reflection. This contributes to maximising the proportion of neutrons successfully transported to and from the chamber, maximising the statistical sensitivity.[133]

The neutrons emerging from the source are unpolarised. To polarise the neutrons, they are passed through a 100 μ m aluminium foil as the UCN are guided through the bore of a 5 T superconducting magnet [113]. The magnetic moment $\vec{\mu}$ of the UCN in the 5 T field [52] imparts $\vec{\mu} \cdot \vec{B} \approx 300$ neV of kinetic energy to the neutrons with spins aligned with the field, and fully repels those with lower energies. The neutrons of the correct spin state are accelerated towards and pass through the foil in greatly increased numbers thanks to their high kinetic energy. Upon leaving the high field region, this kinetic energy is again lost, decelerating them to a storable energy. The neutrons emerging from this polariser are > 99 % polarised. The polarisation of the neutrons leaving the magnet is maintained by a series of guiding field coils, described in section 3.4.

Between the polariser and the experiment, an adiabatic spin flipper labelled SF1 is installed. It consists of a pair of coils above and below the incoming neutron guide, driven with RF. This allows the initial polarisation state of the neutrons to be flipped. This allows us to both perform specialised measurements of the efficiency of the detection of each spin state, and to control systematics related to an asymmetry in this. The spin flipper can be turned off or on at will, and during normal operation this was done every 112 cycles. The efficiency was determined from fits to several combinations of spin flipper settings, and was found to be close to perfect at (100.4 ± 0.3) %.

To control the path of neutrons in and out of the chamber, a mechanical neutron switch was devised. This consists of several bent UCN guides on a rotating platform such that the correct path for the UCN can be chosen. There are four positions: empty (connects storage chamber with USSA), fill (connects chamber with UCN source), monitor (connects source directly to USSA) and pump (allows UCN guides to be pumped down). The switch is shown in Figure 3.9.



Figure 3.9: Photo of the UCN switch. The inner platform with different guides mounted can be rotated in order to make different connections between the various ports for filling and emptying of UCNs.

The chamber itself consists of a Rexolite (polystyrene) ring of 47 cm inner diameter and 15 cm height, sandwiched between two aluminium electrodes, sitting in grooves of 15 mm depth, giving an effective chamber height of 12 cm. The inner surface of the ring is coated with deuterated polystyrene using a technique developed by the collaboration [126, 134], while the two electrodes are coated with diamond-like carbon [135]. These coatings were chosen for their high Fermi potentials and low neutron losses per reflection, maximising neutron storage times. Typical neutron storage characteristics during datataking are displayed in Figure 3.10. It is also important to verify that the chosen material cannot cause loss of neutron polarisation upon reflection (for example, due to strong local magnetic fields). Such an effect would cause severe reductions of the measured T_1 and T_2 times for the neutrons [141]. Measurements of the T_1 time of the neutron population demonstrating that this depolarisation mechanism has a negligible effect are shown in Figure 3.11. In the insulator ring, two circular sapphire windows were installed to allow the mercury magnetometer probe beam to cross the chamber. These are coated with deuterated polyethylene to maximise UCN performance. In the lower (ground) electrode, two shutters are installed. The UCN shutter is mounted in the centre, and allows neutrons to enter or leave the chamber. The mercury shutter is opened only at the beginning of the cycle to allow the polarised mercury vapour to enter the chamber. The UCN shutter is opened during the filling and emptying of the chamber. During the Ramsey sequence, both shutters are sealed tight to maximise the storage time.



Figure 3.10: Storage curve showing measured neutron counts after a given storage time, measured at the beginning of 2016 data taking. The storage curve cannot be well described by a simple exponential because slow and fast UCN will have different storage time constants, and the population of stored UCN will typically contain a range of energies. The curve is better approximated by the sum of two exponentials, with fast and slow decay constants.

At the end of each measurement cycle, the chamber is emptied and the neutrons in each spin state must be counted. The two spin states are separated by the USSA (U-Shaped Simultaneous Spin Analyser), shown in Figure 3.12, [136] before being separately counted by the two NANOSC [137] detectors. At the USSA, the guide from the chamber is split into two arms, each featuring a polarising iron foil, magnetised by a strong permanent magnet, permitting only one spin state to pass. The analysing power (using the formalism of [138], equivalent to the measured asymmetry for an incident hypothetical fully polarised beam) for the foil in arm A is (91.0 ± 4.4) % and in arm B is (89.7 ± 4.3) %. Before each foil, each arm features its own independent adiabatic spin flipper, referred to as SF2a and SF2b. In the case where the spin flipper is off, only neutrons with the correct spin state are permitted to pass through the foil in each arm. However, if the spin flipper is active, the spin states of all neutrons will be inverted on each pass, meaning that neutrons of



Figure 3.11: UCN longitudinal polarisation after a given storage time, from the same dataset as storage curve Figure 3.10. The longitudinal depolarisation rate commonly denoted as T_1 is fitted and shown to be in excess of 4000 s. T_1 is expected to be limited by interactions between neutrons and small magnetic impurities on the chamber walls. It is therefore shown that interactions with the chamber walls are not the limiting factor in the neutron polarisation time, since these would be expected to have similar depolarising effects for all neutron polarisation states.

the wrong spin state will be reflected by one arm (with their initial state restored) and may on a second attempt enter and be counted in the other arm. Typical efficiencies were (97.0 ± 1.2) % for SF2a and (97.1 ± 0.9) % for SF2b. The analysing power for the entire USSA assembly was measured to be (83.5 ± 4.0) %. Compared to the old sequential spin analysis system, as used by the Sussex-RAL-ILL experiment, the counting statistics are improved by (23.9 ± 1.0) % and the Ramsey visibility improved by (6.2 ± 4.9) %, resulting in an overall improvement in the statistical sensitivity of (18.2 ± 6.1) %.

To each arm of the USSA, an independent NANOSC UCN detector [137] is mounted, photographed in Figure 3.13. Each NANOSC consists of 9 independent PMT-scintillator units. The scintillator consists of two layers, the first consisting of 60 μ m of GS30 scintillating plastic (depleted in ⁶Li), and second 120 μ m of GS20 scintillating plastic (enriched in ⁶Li). The scintillator is bonded to a 80 mm PMMA light guide, which then leads to a R11187 Hamamatsu PMT. To detect a UCN, it must be first captured by the ⁶Li, which reacts to form an α particle and a triton (³H nucleus). The first ⁶Li depleted layer is thick enough that the capture occurs deep enough that both resulting particles are caught in the scintillator, giving a single energy peak for a neutron capture well distinguished from



Figure 3.12: Cut view of the USSA, as described in text. From [136].

ambient gamma rays or other noise. The relative efficiency was shown to be between 87 % and 130 % of a ³He Strelkov detector, depending on the precise experimental configuration and incoming UCN energy spectrum [137].

Although the background is very low and well distinguished from signal events, it is possible that a correlation between any remaining background and other experimental parameters may cause a systematic effect. This was investigated thoroughly and ruled out during dry runs of the apparatus without UCN during downtime [111].



Figure 3.13: Photo of the NANOSC detector. On the front can be seen the nine scintillator stacks, connected by the blue light guide assembly to nine PMTs mounted at the rear. From [137].

3.4 Magnetic Field Environment

In order to successfully measure the Larmor frequency of neutrons using the Ramsey technique, great care must be taken in the preparation of the magnetic field environment. The neutron spin polarisation must be maintained for their transport to and from the chamber, and the field within the precession chamber must be extremely uniform to maximise the sensitivity of the Ramsey cycle. The experiment takes place within a large experimental facility in close proximity to the PSI synchrotron, and other experiments in the neighbouring hall using large magnets, which are periodically ramped and create noise. It is also located next to a road where (ferromagnetic) vehicles are used for deliveries. Across this road lies a major superconducting magnet test facility, the stray fields from which can be strong enough to reverse the direction of the Earth's ambient field as it is ramped up and down each day. These external fields are passively monitored by several remote magnetometer modules spread around the experimental hall to allow the sources of magnetic disturbances to be identified, but nEDM experimenters cannot influence significantly the running patterns of other experiments. Significant effort must therefore be invested to achieve a uniform and stable field within the experiment.

The design of the superconducting magnet solenoid used to polarise the UCN means that the guiding field must smoothly vary from the horizontal field coaxial with the UCN guide from the source to the vertical field used inside the precession chamber. A series of solenoidal coils colinear with the horizontal guides maintains polarisation up to the switch, where the field smoothly rotates to vertical. This is then maintained by the vertical guide coil assembly, a set of three solenoids mounted around the vertical UCN guide. This meets the uniform B_0 field within the vacuum tank smoothly. The direction of the field in the vertical guide and switch must match the direction of the intended B_0 field in the tank. The field in the USSA upon emptying the tank must also be optimised to ensure there is no additional loss of polarisation during transport, and to ensure proper operation of SF2a and SF2b.

In order to stabilise the field within the precession chamber against external interruptions, both dynamic active shielding using large coils ('Surrounding Field Coils', or SFC for short) [139] and passive shielding using a four-layer μ -metal shield is employed. The SFC system consists of six large rectangular coils (X+, X-, Y+, Y-, Z+, Z-) surrounding the experimental area, with sides 6-8 m in length, depicted in Figure 3.14. The field close to the experiment is monitored by nine three-axis fluxgates, mounted within the coils but outside the μ -metal shield. The current through each coil can be independently adjusted between $-10 \rightarrow +10$ A or $-20 \rightarrow +20$ A. A live-feedback algorithm dynamically adjusts the current through each coil in order to stabilise the field as measured by the fluxgates. This system is successful in shielding the experiment by a factor of 5–50 at a bandwidth from 10^{-3} Hz up to 0.5Hz.

The four-layer μ -metal shield is responsible for shielding higher frequency signals, and for providing additional protection against slower field changes. This shield was originally used in the previous-generation experiment by the Sussex-RAL-ILL collaboration, which is responsible for setting the current best limit on the nEDM. It consists of four concentric cylindrical shields, with push-fit end caps. There are a number of holes through the shield to provide for vacuum pumping, the probe beam for the Hg magnetometer, degaussing coils, UCN guides and the high-voltage feedthrough. All components within the innermost μ -metal shield are screened for magnetic contamination using the BMSR-2 facility at PTB Berlin, to avoid creating large local gradients which may cause systematic effects[140, 141].

To allow for the shields to be degaussed, a cable containing several windings is threaded



Figure 3.14: Photo of the experimental installation, with the μ -metal shield lifted by a crane during installation. The six SFC coils are marked.

axially through each layer of the shield, passing through the centre of each endcap. These coils are driven by a large power amplifier with up to 10A current, via a 1:1 transformer to remove any DC offset, and a special isolation box, which allows the coils to be shorted, isolated and gradually reopened to allow them to be isolated without introducing large transients which may magnetise the shield. This system is driven by a 10 Hz sine wave, decaying exponentially to zero over several minutes. The initial field is sufficiently strong to drive the μ -metal into saturation to ensure full degaussing, and as the field is reduced the μ -metal domains are randomised and the shield is demagnetised.

The shielding factor of the shield was determined by applying known fields externally using the SFC coils, and measuring the field within using a fluxgate. The shielding factor in the X direction was determined to be 13300 ± 600 , in the Y direction to be 1600 ± 20 , and in the Z direction to be 8600 ± 300 .

The nominal 1μ T vertical field within the shield, named B_0 , is produced by a $\cos\theta$ coil wound on the outside of the vacuum tank. It is driven by a highly stable fixed current supply, at a current of approximately 17 mA. The current is continuously monitored and recorded by a Keithley 6487 picoammeter. The field created is substantially larger than would be predicted for the same coil in free space: a substantial portion is due to magnetisation of the innermost layer of the μ -metal shield. The resulting magnetic field is already impressively uniform, with gradients of a few tens of pT/cm close to the centre, however in order to maximise neutron depolarisation times and control systematic effects, it is necessary to further trim the field.



Figure 3.15: Measured Ramsey visibility as a function of time for an optimised trimcoil current configuration using the homogenisation algorithm [127], from special datataking runs during the 2016 ramp up period. The measured T_2 of (1568 ± 19) s shows a vast improvement over that obtained in previous experiments, which were of the order 600s [88]. This directly increases the statistical sensitivity of the apparatus.

To further homogenise the field, 33 trimcoils are wound on the vacuum tank, or on chimneys protruding through the shields. Each of these is driven by an independent computer-controlled current supply, adjustable in two ranges between $-20 \rightarrow +20$ mA or $-200 \rightarrow +200$ mA, in 1 μ A steps. The precise magnetic state of the μ -metal shields is not completely reproducible between degaussings, which are necessary whenever the direction of the field is inverted or the shields must be opened. Therefore, attempts to optimise the trimcoil currents offline based on information obtained from field mapping campaigns were only partially successful. A technique to measure the field state after each degaussing using the caesium magnetometer array was developed [127], which was then used to optimise the trimcoil settings, minimising the gradients of the vertical component of the field across the chamber to maximise the neutron T_2 . An experimentally measured T_2 decay curve is shown in Figure 3.15. To validate the collaboration's understanding of the effect of field gradients on neutron polarisation characteristics, a dedicated investigation of the effect of known gradients on final state neutron polarisation for different pulse sequences was performed, detailed in Figure 3.16.

For the Ramsey technique, it is necessary to apply a rotating field to apply the $\pi/2$ pulses at around the Larmor frequency of the neutrons (≈ 30 Hz) at the beginning and end of the free precession period. This is approximated by a single coil AS1, which provides an oscillating 30Hz field driven by the custom DAQ board [142]. This oscillating field is equivalent to two counter-rotating fields at the same frequency; the additional counter-rotating field is adjusted such that a 2 second pulse corresponds to a $\pi/2$ flip, as it is difficult to predict a priori due to the large conducting masses in the electrodes and tank which significantly disrupt the field shape and magnitude. For the mercury magnetometer (described later in section 3.5), it is also necessary to apply a $\pi/2$ pulse at the beginning of each cycle close to the mercury Larmor frequency (≈ 8 Hz) to start the free spin precession. This is achieved with a pair of perpendicular RF coils, driven with sinusoidal currents, with a 90° phase difference, in order to create a rotating field. For these coils, the amplitude and phase relation is also empirically adjusted so as to maximise the mercury sensitivity.

While this scheme resulted in an extremely uniform and stable field, it was still necessary to monitor and correct for variations from cycle to cycle. For this reason, the experiment was equipped with online magnetometry systems, described in section 3.5. In addition, to empirically determine the exact response of each trimcoil and constrain higher order gradients contributing to several systematic effects, offline campaigns to map the field using a dedicated mapper robot took place in 2013, 2014 and 2017. These measurements are documented in chapter 5.



Figure 3.16: Measured Ramsey visibility after 180s compared to Spin-Echo measurement[117] and longitudinal polarisation measurements, versus additional vertical gradient applied on top of optimised configuration. The longitudinal polarisation measurements probe the baseline efficiency of the initial polarisation and of the specificity and efficiency of the USSA spin sensitive neutron detection system, as well as ' T_1 ' (or spin-lattice) type depolarisation such as that due to small scale magnetic contamination of the chamber walls. The Spin-Echo measurements probe ' T_2 ' (spin-spin) type depolarisation, such as that due to large scale *B*-field inhomogeneities such as a $\frac{\partial B_x}{\partial x}$. The Ramsey measurement also includes the effect of 'gravitationally enhanced depolarisation', a T_2^* effect. This type of depolarisation occurs when UCN are stored in a vertical gradient $\frac{\partial B_x}{\partial z}$, as UCN of different energies have a different height distribution throughout the chamber. These results provide clear validation of the measurements presented in [116], which are used to inform the data taking strategy.

In order to precisely monitor the magnetic environment within the chamber, the experiment is equipped with two additional magnetometry systems. The first, the cohabiting mercury magnetometer, is designed to measure the volume average magnetic field within the neutron chamber, while the array of 16 external caesium magnetometers is designed to capture information about the gradients of the field.

The mercury system was originally developed by the Sussex-RAL-ILL collaboration [20, 21, 88, 143]. It differs from previous magnetometers used in nEDM experiments because it measures the free precession of atoms within the neutron storage volume, in contrast to previous efforts where a limited number of magnetometers were placed around the cell, limiting the accuracy in cases where the field was not perfectly uniform. First, ¹⁹⁹Hg vapour is produced in the source by heating isotopically pure powdered HgO. This diffuses into the polarising cell where it is optically pumped by circularly polarised light from a 204 Hg discharge lamp. This causes the spins of the mercury atoms to be aligned (anti) parallel to the vertical B_0 field. At the beginning of each measurement cycle, a shutter is opened to allow the polarised vapour to enter the precession chamber. Before the neutron Ramsey sequence begins, a $\pi/2$ pulse is applied at the Larmor frequency of the mercury to flip the spin into the XY plane (transverse to the magnetic field), in which they begin to precess. A probe beam of circularly polarised light from a second ²⁰⁴Hg discharge lamp is shone across the chamber, and the transmission of the beam is modulated as the atoms precess. The light enters and leaves the chamber through sapphire windows on either side. The light leaving the chamber is detected by a photomultiplier tube. The sensitivity of the magnetometer is limited by the signal to noise ratio of the readout, which depends on a combination of several factors: the initial polarisation and density of the mercury, the depolarisation time of the mercury, and the noise in the probe light amplitude. Upon reversal of the electric field, the depolarisation rate initially decreases but recovers over several cycles; the mechanism for this is however not well understood [143–145]. To improve the quality of the probe light beam, a laser-based readout system was developed [144, 146, 147], however this was not used for the bulk of nEDM datataking in 2015 and 2016. The typical sensitivity of the mercury magnetometer was sufficient that it was estimated to contribute no more than 2.88~% to the statistical uncertainty of the nEDM experiment [144].

While the mercury magnetometer effectively compensates for drifts in the vertical average magnetic field B_0 , it also comes with several major drawbacks. First, the ultracold

neutrons are slow, with velocities of a few m/s, meaning that the centre of mass of the stored ultracold neutrons is around 3-5 mm below the centre of the chamber[114, 117], while the thermal mercury atoms fill the chamber more evenly. The result of this is that a vertical gradient will impact the ratio of neutron and mercury frequencies. In the case of small vertical gradients (up to ± 30 pT/cm or so), the effect is linear, however for larger gradients, it must be noted that the Ramsey technique is sensitive to the phase accumulated during the free precession modulo 2π . So-called Ramsey wrapping causes a deviation from linearity in this case [21, 115, 117]. Second, as the mercury atoms are fast, they will sample the chamber nonadiabatically, effectively experiencing an average field of $| < \vec{B} > |$, while the slow neutrons will sample the magnetic field adiabatically, effectively experiencing a field of

$$\left\langle \left| \vec{B} \right| \right\rangle \approx \left| \left\langle \vec{B} \right\rangle \right| + \frac{\left\langle B_T^2 \right\rangle}{2 \left| \left\langle \vec{B} \right\rangle \right|}.$$
 (3.9)

The effect of this is that transverse inhomogeneities in the field will serve to increase the neutron precession frequency, while leaving the mercury frequency unchanged. Finally a conspiracy between the radial components of the magnetic field and motional magnetic fields (arising from relativistic transformation of the electric field into the rest frame of the mercury atoms, $\vec{B}_m = \frac{1}{c^2}\vec{v} \times \vec{E}$ where c is the speed of light, \vec{v} is the mercury velocity and \vec{E} is the electric field) results in a frequency shift in the mercury atoms proportional to the applied electric field, thus mimicking the effect of a true EDM [140, 141, 148–152]. Each of these effects and how they were compensated will be described in more detail in chapter 4.

To compensate for drifts in the vertical gradient within a run, and to optimise the magnetic field before each run, an array of 16 optically pumped caesium magnetometers is employed [153–156]. The cesium array consists of two different types of caesium magnetometer units: 'High Voltage' and 'Ground' magnetometers, with the HV magnetometers mounted above the precession volume on the HV electrode, and the ground magnetometer consists of a paraffin-coated glass spherical cell of 25-30 mm internal diameter, filled with caesium vapour. At 45° to the magnetic field, a circularly polarised beam of 894.6 nm laser light crosses the cell, pumping the caesium D1 transition. This effectively builds up a net magnetisation along the beam axis. A coil applies an RF field at around the Larmor frequency of the caesium atoms transverse to the main magnetic field, which begins to drive the precession around the main magnetic field much like a classical driven oscillator. The transmitted light is modulated as the atoms precess. By observing the phase rela-

tion between the driving field and the modulation of the transmitted light, the difference between the frequency of the driving field and the resonant frequency of the caesium atoms can be observed. The driving frequency is tuned to match the resonant frequency using a phase-locked loop. A sensitivity to field changes of less than 0.1 pT is readily achieved in the experiment, however the absolute accuracy is substantially worse than this.

3.6 Data Acquisition

The experiment was largely automated, and the majority of operations were controlled by a computerised data acquisition system. The core of the system was a custom-built data acquisition board [142], which was responsible for generating the RF pulses for the Ramsey sequence and the mercury magnetometer, recording the mercury PMT signal, and controlling the timing of other experimental components such as valves and switches. This was interfaced with a computer program, AcqEDM, which provided an interface for programming of the board and was the primary user interface for control of the experiment. The UCN detector PMTs were read by the FASTER DAQ system, which interfaced with AcqEDM (via the blinding package) to provide information on neutron counts in each detector each cycle, which was needed to choose the correct operating points on the Ramsey fringe. The caesium magnetometer array was read out by a pair of custom FPGA-based boards, which read out the precession signals, and provided the PLL-controlled driving RF signal. This was interfaced with the slow-control system, which was also responsible for controlling and passively monitoring of many of the components of the system, such as the vacuum system, temperature, pressure and humidity gauges and the trimcoil supplies. All systems were synchronised to a 10 MHz rubidium atomic clock, which was itself synchronised to GPS time.

A data blinding system was implemented in order to disguise the true data such that psychological biases towards a particular expected experimental result of those analysing the data cannot influence the running of the experiment or of the analysis. An experimenter may, consciously or subconsciously, make decisions about, for example, which data to cut and which to include in the analysis based on such biases. Such biases have been suspected to distort measurements of several physical quantities, which are often seen to fall closer to previous results and further from the true modern value than statistics might suggest. The use of blind techniques when analysing and running experiments has long been standard practice in medical trials, and is increasingly becoming standard practice across particle physics measurements.

Several techniques were investigated to provide the required disguise. As all of the parameters which could be potentially hidden to remove any possibility of computing the EDM before unblinding (such as neutron counts or magnetometry signals) were needed to be known for diagnostic and monitoring purposes during the running of the experiment, the only option was to somehow subtly alter the results in a way that would not affect the running of the experiment. Early options considered were to record the neutron spin flip frequencies used in the Ramsey sequencing with an electric field correlated offset, or to add an electric-field-correlated offset to the recorded mercury frequency, however the implementation of these options would either impact the data taking scheme in a way that may bias the results, or be easily reverse engineered to reveal the offset. It was decided that the best option would be to add an electric field correlated offset to the neutron counts, thus mimicking a real neutron EDM signal. While an online fit of the Ramsey fringe is performed in order to ensure that the working points on the Ramsey fringe are correctly set, only cycles taken at E = 0 are included in this fit, so datataking is unaffected. The size of shift required is small enough so as not to affect the manual diagnostic checks: typically 1-2 neutrons must be shifted out of 10,000 to 20,000 counts.

The implementation of the blinding is in two parts: online and offline. In order to accurately shift the correct number of neutrons each cycle in order to produce the desired offset, it is necessary to know several parameters that can only be known well by fitting a large dataset: in particular the Ramsey visibility α and the central position of the Ramsey fringe. However, AcqEDM requires knowledge of the neutron counts after each cycle in order to compute the correct working points on the Ramsey fringe. Therefore, the online blinding program is responsible for providing the cycle-by-cycle live neutron counts to AcqEDM, while the offline blinding program is responsible for providing realistic, statistically well-behaved blinded data to be used for further analysis. The offset used for online blinding is large and set randomly each run, such that the statistical noise induced is enough that blinding offsets cannot be determined by comparison of the online and offline blinded counts, while the offline blinding uses a smaller offset computed to give a false EDM of the order $1 \times 10^{-25} e$ cm. The actual number is randomly computer generated, and encrypted such that a human could never accidentally see it. The offline blinding was performed in two identical stages: primary and secondary. The primary blinding stage is applied once to all raw data. The secondary blinding stage is performed twice on the already-blinded data to produce two separate datasets, one to be analysed by each analysis groups, differing only by the offset applied. The unblinded and primary-blinded data is stored in an encrypted form and not accessible to the analysis groups. Once the analysis groups each believe they have a final analysis setup, they will submit an unblinding request to the collaboration board. If the collaboration board is satisfied, the secondary level of blinding will be removed, and it will be checked that the two analyses are consistent. If this holds, each analysis will be performed on the original dataset to give the final EDM result.

3.7 High Voltage

In order to measure a neutron EDM, it is necessary to apply an electric field across the precession chamber. As sensitivity to an EDM is linearly proportional to the applied electric field, it is critical to apply the largest field possible across the chamber. The high voltage system consists of a commercial power supply manufactured by FUG, which is rated to supply up to ± 200 kV into a commercial high voltage cable. The polarity can be electronically switched. The HV line is connected to the experiment via a large resistor, to limit the potential damage caused by a breakdown. A custom-design feedthrough with the bulk of the construction made of ceramic is used to bring the high voltage into the vacuum tank.

The geometry of the chamber was optimised to maximise the achievable high voltage. The height of 12cm was chosen in order to balance the achievable electric field with building a suitably large chamber to store as many neutrons as possible in order to maximise the statistical sensitivity. The precise profile of the grooves in the electrode in which the insulator sits was also optimised to maximise the achievable high voltage. Corona rings were mounted to each electrode in order to avoid sharp edges (which create regions of high electric field that seed high voltage breakdowns) and to ensure that the high voltage and ground caesium magnetometers remain in low field regions to avoid any possible high-voltage effects on those.

In order to further optimise the attainable electric field, several steps were taken. During normal EDM running, the pressure inside the apparatus is maintained at approx 1.5×10^{-3} mbar of helium in order to suppress high-voltage discharges. In the end, the limiting factor in high-voltage operation was discharges occuring due to the high-voltage caesium magnetometer fibres. During operation, voltages of ± 132 kV were consistently achieved with these fibres installed, while the full 200kV could be consistently achieved without. On one occasion, a voltage of 200kV could be applied with the fibres in place. The normal running voltage of ± 132 kV corresponded to a field of ± 132 kV/12 cm = 11 kV/cm.

Stray currents flowing within the magnetically shielded region can create high-voltage correlated magnetic field effects, which can mimic a neutron EDM. This can either be due to the field produced by the current flow itself, or the current flow causing magnetisation of, for example, the μ -metal shielding. In order to constrain this, the return current feed from the ground electrode is fed through a picoammeter with special protective circuits in order to measure any current flow. This current feed leaves the ground electrode, flows to the picoammeter, then the return line is routed back through a screened cable to the vacuum tank ground. The vacuum tank itself is electrically isolated from other parts of the experiment, and is only grounded via the high voltage feed connection. This absence of current loops for the high voltage charging current is designed such that it is impossible for the electrode charging currents and associated systematic effects are detailed fully in [111, 157].

3.8 Conclusion: Expected Sensitivity and Outlook to n2EDM

In section 3.1 the factors that contribute to the ultimate statistical sensitivity of nEDM experiments were introduced. These can be summarised by Equation 3.8

$$\sigma d_{\rm n} = \frac{\hbar}{2ET\alpha\sqrt{N_0\cdot m}}$$

where E is the achieved electric field (in each polarity), α is the Ramsey visibility, T is the free spin precession time between the $\pi/2$ pulses in the Ramsey sequence, N_0 is the average number of neutrons per cycle and m is the total number of measurement cycles. As outlined throughout this chapter, the recently completed nEDM experiment implements many advances over the previous generation Sussex-RAL-ILL experiment, improving on both the statistical precision and the control of systematic errors. The apparatus was used for sensitive EDM datataking from August 2015 to December 2016, during which time enough data was collected to reach a raw sensitivity (1σ) of $1 \times 10^{-26} ecm$. The analysis of this data is presented in chapter 4. The apparatus was also used for a number of auxiliary measurements, both to constrain new physics phenomena, and to better understand systematic effects affecting nEDM experiments in preparation for n2EDM. These include measurements of neutron 'disappearance' to constrain mirror neutron oscillations [110, 111, 158], an improved measurement of the neutron to ¹⁹⁹Hg magnetic moment ratio, constraints on axionlike particle mediated forces [159], investigations of Lorentz violating parameters [160] and most recently constraints on ultralight axion dark matter [161], detailed as part of this thesis in chapter 6. Measurements were taken to validate calculations of systematic effects due to linear vertical magnetic field gradients [144, 151] and higher order magnetic field inhomogeneities [141, 144]. An investigation[116] of gravitationally enhanced depolarisation[115] and Ramsey wrapping was conducted. Finally, a novel spinecho technique to explore the energy spectrum and depolarisation mechanisms of ultracold neutrons was developed [117], using a π pulse between the initial and final $\pi/2$ pulses to disentangle gravitationally enhanced depolarisation effects and intrinsic depolarisation effects under the application of known vertical gradients.

This intense development work stands the collaboration in good stead to construct and operate the n2EDM experiment. The guiding principle of the design of this experiment is to maximise the statistical sensitivity while maintaining sufficient control of systematics. The most striking change is the change to a double chamber apparatus, with a single shared high voltage electrode mounted in the centre of two ground electrodes, separated by plastic insulator rings as before. The overall experimental layout is displayed in Figure 3.17, with a detail view of the precession chamber at the heart of the experiment is shown in Figure 3.18. The double chamber design has the effect of substantially increasing the volume (increasing UCN counts) and, as a key benefit, provides substantial protection from systematic effects, many of which will cancel as both directions of electric field are probed simultaneously in each chamber. Also, as there will no longer be any caesium magnetometers mounted at high voltage, the limiting effect of the caesium fibres on the attainable electric field will be lifted. The guide geometry is also well optimised for the UCN characteristics and energy spectrum available at the PSI UCN source, which further improves UCN statistics. Furthermore, it was chosen to use a larger precession chamber, with an internal diameter of 80 cm compared to the 47 cm of the previous experiment. This serves to improve UCN statistics, but at the expense of greater vulnerability to systematic effects related to the homogeneity of the magnetic field, and having stricter homogeneity requirements to maintain the neutron polarisation. This will be achieved using a larger magnetically shielded room (with more layers of μ -metal) and B_0 coil, and monitored by a large array of less sensitive but highly accurate caesium magnetometers. The new generation experiment promises to achieve a sensitivity of 1.2×10^{-27} e cm for 500 days of running, based on realistic UCN Monte Carlo simulations and models validated by real world measurements of the PSI UCN source, assuming only demonstrated performance of the UCN source and using only demonstrated technologies. The key figures leading

to this forecast are summarised in Table 3.1. Research and development is ongoing to develop upgrades for the experiment, such as coated non-metallic electrodes to reduce the impact of Johnson noise on the statistical sensitivity. Work on better understanding and optimising the UCN source is also underway, which will further improve the statistical sensitivity of the experiment directly.



Figure 3.17: Sketch of entire n2EDM apparatus in area South. Grey blocks around area are concrete blocks for radiation protection. Magnetically shielded room (MSR) to contain vacuum tank (5) and precession chambers (4), mounted on large granite blocks (7) for stability. Superconducting magnet polariser (1) provides a polarised UCN beam from beam port (2). This feeds a Y-piece and switch (3) to feed each precession chamber. Upon emptying the chambers, the neutrons of each spin state are counted by a pair of USSAs (6) (as used in the previous generation experiment, described in section 3.3).



Figure 3.18: Close up of n2EDM precession chamber conceptual design: dual chambers surrounding a single high voltage electrode, with independent mercury magnetometry systems for each chamber. Ground electrodes on top and bottom to support vast caesium magnetometer array for online determination of field configuration.

	Sussex-RAL-ILL	nEDM	n2EDM Baseline
Diameter (cm)	47	47	80
Height (cm)	12	12	12 (× 2 - Double
			Chamber)
Coatings	Quartz, DLC	dPS, DLC	dPS, DLC
E (kV/cm)	10	11	15
T (s)	130	180	180
α	0.6	0.75	0.8
N (typical, per	14,000	15,000	Simulated
cycle)			121,000
Typical Sensitiv-	26	11	2.6
ity Per Day $\sigma d_{\rm n}$			
$(\times 10^{-25} ecm)$			

Table 3.1: Forecast sensitivity parameters of n2EDM experiment in comparison with previous experiments. Calculations for n2EDM are from the n2EDM Project Overview Report (January 2018) [162], considering improvements achievable using only presently proven technology - the 'Baseline', and assuming no improvements in UCN source performance. Performance figures of ILL experiment from [20].
Chapter 4

Data Analysis

The nEDM experiment at the PSI took sensitive data from August 2015 to December 2016, with a break during the accelerator maintainance window taking place January 2016 to May 2016. Analysis of the data to obtain a final nEDM result is currently underway, with the 'statistical' part of the analysis described in this chapter complete. Following the finalisation of the analysis of systematic effects, the data will be unblinded, to reveal what is likely the most statistically sensitive measurement of the nEDM to date.

4.1 Structure of the Analysis

In order to provide a cross-check, the data arising from the experiment is analysed by two distinct 'analysis groups'. Each group is responsible for completing a full analysis of the data. The data is blinded as it is produced, by adjusting the neutron counts in each detector to simulate a real EDM (see section 3.6). Before publishing a result, and as part of the unblinding process, it must be ensured that the results arising from each analysis are consistent. Each analysis group has access to the full dataset, however the data is blinded differently to prevent comparison of preliminary results biasing the development process. First, a 'primary' blinding is applied to the entire dataset, common to both groups. Then, a second 'secondary' blinding is performed on top on this, with different randomly generated offsets for each group. This chapter describes the analysis developed by the eastern analysis group, of which the author is a part.

The eastern analysis is a large and complex software project with many contributors over a period of years. It is divided into a series of 'steps': import, cycle-level analysis, runlevel analysis and crossing-point analysis. Each discrete step, in addition to the mercury fit program[163], the map analysis (see chapter 5) and some shared routines for file access, is stored in its own git repository for version control [164]. Each step of the analysis can be run independently and records in its output all versions of software used to produce its results, allowing traceability of all results.

The majority of the code is written using MATLAB[165] and its associated toolboxes. The mercury fit analysis and the map analysis is done using Python3[166] with the NumPy[167], SciPy[168] and Pandas[169] libraries.

4.2 Importing the Data

The data arising from the apparatus is recorded by several programs split across multiple computers. The majority of this is stored in a text based format, with the exception of data from the UCN detectors, which is stored in a binary format. The purpose of the 'import' stage of the analysis is to collect these many independent files together, synchronise the timestamps, and to store them in a fast, compressed binary file format for fast and efficient access. The HDF5 format [170] was found to meet these requirements, and is also used to store intermediate results at each step of the analysis. For each data run, a single HDF5 database is produced. All data recorded by the apparatus is stored in the database, with the exception of the individual neutron detector pulse data. The import step is responsible for reading the full data produced by the FASTER detector DAQ system [171], and applying predetermined thresholds to count the number of neutrons of each spin state detected each cycle as a function of time.

4.3 Cycle Level Analysis

The cycle-level analysis is responsible for analysing everything that is independent cycleto-cycle. This consists of the analysis of the mercury free induction decay (FID) signal, the counting of neutrons arriving during the emptying period of the apparatus, and the computation and reporting of 'performance parameters', identifying bad cycles where the apparatus did not function properly, which are used to decide cuts at the run level.

The mercury magnetometer fit is performed by a specialised algorithm described in [163]. The signal recorded by the mercury magnetometer is an FID signal, a decaying sine wave at approximately 8 Hz (the Larmor frequency of ¹⁹⁹Hg in a 1 μ T field), and a decay time constant between approximately 30 s and 200 s. The signal is passed through a filter to allow only a restricted range of frequency components to pass, and digitised[142]. The frequency computation is most precise when the signal is fitted considering the entire

neutron precession period. However, in the case of a rapidly decaying signal in the presence of significant field drifts there is a risk of inducing a bias towards measuring only the field at the beginning of the cycle rather than a true average through the free precession time. In this case, it is possible to sacrifice some precision in the name of accuracy by using a 'two window fit', as used in the previous Sussex-RAL-ILL data analysis[88, 143, 172]. Two sections of the beginning and of the end of the FID signal are fitted, and by comparing the relative phase of each signal, it is possible to deduce the mean precession frequency over the whole period. In the choice of window sizes, there is a further bias-variance tradeoff, with shorter window lengths giving less bias but poorer precision. A Monte Carlo simulation was used to optimise the fit type and window length chosen as a function of the noise level and estimated magnitude of magnetic field drift, in order to maximise accuracy. The results of this optimised fit were used for all subsequent stages.

4.4 Run Level Analysis

The next level is the run level analysis, responsible for all analysis which can be done with just one run, including the determination of the neutron frequencies by fitting the Ramsey curve, and the determination of the EDM measured from each run.

The run level analysis is performed on 'sets' of data, which may not correspond precisely to exact run numbers used in datataking. A set consists of a consecutive set of runs performed in the same nominal magnetic field configuration (same B_0 direction and trimcoil settings), with no degaussing in between. Sometimes the run had to be stopped temporarily for technical reasons, and restarting without any changes resulted in the run number being incremented. Working with sets of merged runs allows cycles to be most effectively divided into 'subsets'.

4.4.1 Division of Sets into Subsets

An entire 'set' consisted of typically two days of datataking at one magnetic field configuration. Within this set, the electric field direction and the states of each spin flipper were regularly reversed according to the scheme shown in Figure 4.1. However, during one set, it is possible that many parameters, such as field gradients, may drift. Therefore, it is useful to further divide the data into shorter *subsets*. The choice of subsets is informed by the electric field reversal interval, as cycles in at least two different electric fields are required to extract an EDM. In particular, it was found necessary to cut the data into 'ABBA' subsets, combining two high voltage reversals between both electric field orientations and the intervening zero-field cycles, in order to be able to distinguish between linear time drifts in the vertical field gradient and a real EDM-like effect. This process is mostly automated, with a few manually defined exceptions where the algorithm failed: for example where large parts of the run had to be cut.



Figure 4.1: Time structure of datataking scheme. Upper plot: yellow = USSA SF2a/b state (inverted every four cycles, see section 3.3), blue = SF1 state (inverted every 112 cycles). Lower plot: green = E field polarity (8-48-8-48 cycle of 8 cycles datataking at zero field followed by 48 cycles at alternating HV states), alternating light/dark grey = indicative division of single long run into 'ABBA' sections, combining both HV states. From [91].

4.4.2 Caesium Gradients Fit

The PSI nEDM experiment includes an array of 16 caesium magnetometers mounted above and below the precession chamber, though for much of the datataking period only 15 of these were functional. These sensors have relatively poor absolute accuracy, making them unsuitable as the sole measure of the vertical gradient; however they provide excellent precision (of the order 0.1pT, when averaged over one cycle). Thus, they can be used to compensate drifts in the vertical gradient over the course of one run.

To obtain an estimate of the gradient drifts across a run, first a reference field (typically the average measured field across the entire subset) is subtracted from each sensor, so that only relative field changes remain. To this, a field model is fitted, typically consisting of the terms $G_{0,0}$, $G_{1,0}$, $G_{1,1}$, $G_{1,-1}$, $G_{2,0}$, $G_{2,1}$, $G_{2,2}$, $G_{2,-1}$ and $G_{2,-2}$ from the standardised adequate basis detailed in [141] ¹ (i.e. all longitudinally visible terms up or order 2; terms with no z component such as $G_{1,2}$ will produce only a minimal effect in the caesium sensors which measure the magnitude of the field).

4.4.3 Ramsey Fit

In order to determine the neutron frequency for each cycle, it is necessary to fit a Ramsey resonance to a longer set of data, in order to determine the central value of the fringe and nuisance parameters such as the visibility and detector asymmetries. To compensate for variation in the total number of neutron counts, the asymmetry is used in place of the raw number of neutron counts in each spin state [173]

$$A = \frac{N_{\uparrow} - N_{\downarrow}}{N_{\uparrow} + N_{\downarrow}}.\tag{4.1}$$

In the basis of neutron counts, the equations defining the Ramsey curve are

$$N_{\uparrow}(\delta\nu) = \frac{N_0}{2} \left[1 - \alpha \cos\left(\frac{\pi\delta\nu}{\Delta\nu}\right) \right]$$
(4.2)

$$N_{\downarrow}(\delta\nu) = \frac{N_0}{2} \left[1 + \alpha \cos\left(\frac{\pi\delta\nu}{\Delta\nu}\right) \right],\tag{4.3}$$

with $\delta\nu$ the detuning from resonance $\delta\nu = \nu_{\rm RF} - \nu_{\rm n}$, α the Ramsey visibility and $\Delta\nu$ the linewidth $\Delta\nu = 1/(2T - 8t/\pi)$, where T is the free precession time and t the length of each spin flip. These are combined, such that in the case of perfectly symmetric detection efficiency and visibility for each state [173]

$$A = -\alpha \cos\left(\frac{\pi\delta\nu}{\Delta\nu}\right). \tag{4.4}$$

In the case where small asymmetries in detector efficiency are present, this is slightly modified. If the spin-up detector is slightly more efficient than the spin-down detector, so that $N_0^{\uparrow} = N_0 + D$ and $N_0^{\downarrow} = N_0 - D$, and the visibility as measured in each spin state is modified by d so that $\alpha^{\uparrow} = \alpha + d$ and $\alpha^{\downarrow} = \alpha - d$, then one finds

$$A = \frac{D}{N_0} - \alpha \cos\left(\phi\right) - \left(\frac{\alpha^2 D}{N_0} + d\alpha\right) \cos^2(\phi) \tag{4.5}$$

$$= A_{\rm av} - \alpha \cos\left(\phi\right) - \delta \cos^2(\phi), \tag{4.6}$$

where $\phi = \frac{\pi \delta \nu}{\Delta \nu}$. Close to the working points $\phi \approx \pi/2$, the term $\delta \cos^2(\phi)$ is small, so it can be neglected.

¹A brief summary of the basis used throughout this work is included in Appendix A, any gradient expressed using the format $G_{l,m}$ is as described in this basis.

During datataking, the states of the three spinflippers SF1 and SF2a/b are inverted every 112 and 4 cycles respectively (see Figure 4.1). Each combination of spin flipper states may have different asymmetries, as each spin flipper and analysing foil may have a different efficiency, and also which detector counts which spin state is also changed. All cycles within one 112 cycle block with the same SF1 state are combined into one joint fit. An additional parameter is added to absorb any E field dependence of the phase (i.e. an EDM-like effect, such as that added during the blinding process). The overall equation fitted is

$$A = A_{\rm av}({\rm SF2}) - \alpha \cos\left(\frac{\pi\Delta f}{\nu} - \Phi - \frac{d\Phi}{dE}E\right),\tag{4.7}$$

where A is the measured asymmetry of each cycle (with error arising from Poisson statistics on N_{\uparrow} and N_{\downarrow}), $A_{\rm av}$ the fitted mean asymmetry for each SF2 state, α the fitted Ramsey visibility, Δf the difference between the neutron RF pulse frequency and nominal Ramsey frequency (corrected by Hg, in the case without using caesium for gradient drift correction $\Delta f = f_{\rm RF} - \gamma_{\rm n}/\gamma_{\rm Hg}f_{\rm Hg}$), ν the (fixed) linewidth, Φ a free phase to absorb any non-*E*dependent frequency shifts, and $\frac{d\Phi}{dE}$ also a free parameter to absorb any *E*-dependent frequency shifts. This first Ramsey fit is referred to as the 'HVCombi' fit. Typical results are shown in Figure 4.2.

In order to minimise the chance of any as-of-yet unknown effects causing some bias and systematic effect in the Ramsey fit, the dataset is then divided again by *E*-field state. A new 'chunk' is begun at each HV ramp. Another fit is performed, this time only to find the central phase Φ as a free parameter, the rest fixed from the first 'HVCombi' fit.

$$A = A_{\rm av}({\rm SF2}) - \alpha \cos\left(\frac{\pi\Delta f}{\nu} - \Phi\right).$$
(4.8)

Typical results from this second state fit are displayed in Figure 4.3. This two-stage fit is needed to ensure that the nuisance parameters α and $A_{av}(SF2)$ can be sufficiently well determined, even when the chunk consists of relatively few cycles (for example for E = 0chunks, which nominally consist of eight cycles, but can consist of fewer if some must be cut away).



Figure 4.2: HVCombi Ramsey fit to a typical subset, with no correction for gradient drifts.



Figure 4.3: Second stage Ramsey fit of only phase to one HV chunk of a typical subset. The blue and black points correspond to different states of the USSA spin flippers SF2a/b.

Caesium Gradient Drift Correction

In doing the Ramsey fit, drifts in the average magnetic field are compensated using the mercury comagnetometer. Vertical gradients (including but not limited to $G_{1,0}$) however cause shifts in the ratio of the neutron and mercury precession frequencies, as the neutron population sags by a few mm under the effect of gravity[114]. This shift is normally absorbed into Φ , however, random drifts in the vertical gradients can cause a worsening of the quality of the fit, and therefore the overall sensitivity.



Figure 4.4: Fitted relative vertical gradient drift over a typical run.

Relative drifts in the vertical gradient can be effectively measured to high precision using the caesium magnetometer array. Typical results from this measurement are shown in Figure 4.4. For small drifts in the vertical gradient $G_{1,0}$ (resulting in a total gradient less than around $\pm 30 \text{pT/cm}$, avoiding the 'gravitational depolarisation' effects[116]), the frequency shift induced will follow $\delta R = \pm \frac{\gamma_n}{\gamma_{\text{Hg}}} (\langle z \rangle \, \delta G_{1,0}/B_0)$, where $\langle z \rangle$ is the mean height difference between the neutrons and mercury.

In order to correct for gradient drifts before doing the Ramsey fit, the term Δf in Equation 4.7 and Equation 4.8 can be substituted with

$$\Delta f = \left(f_{\rm RF} - \frac{\gamma_{\rm n}}{\gamma_{\rm Hg}} f_{\rm Hg} - \gamma_{\rm n} \langle z \rangle \, \delta G_{1,0} \right). \tag{4.9}$$

This requires prior knowledge of the correct value of $\langle z \rangle$. The value depends on the precise form of the UCN energy spectrum in the chamber. This may be predicted by

Monte Carlo simulation[93] or obtained empirically via a spin-echo technique[117] or by measuring the shift in the measured neutron-mercury frequency ratio when applying a known gradient[91, 116, 174].

Correction with this method results in a substantial improvement in the goodness of fit. The Ramsey fit shown previously (in Figure 4.2 and Figure 4.3) was repeated after this correction, the results of which are shown in Figure 4.5 and Figure 4.6.



Figure 4.5: HVCombi Ramsey fit to typical subset, with CsM gradient compensation. The colour of each point corresponds to the high voltage state, with red corresponding to an applied voltage of +132 kV, blue corresponding to -132 kV, and black to cycles where no high voltage was applied.



Figure 4.6: Second stage Ramsey fit of only phase to one HV chunk of typical subset, with CsM gradient compensation.

Neutron Frequency Fit

While the fitted phase Φ contains information about the average frequency of the neutrons during that chunk, it is informative and useful to invert Equation 4.8 to infer Δf and therefore the neutron frequency during that cycle in order to be able to test for, for example, time structures and any further uncompensated drifts. Rather than inverting the equation analytically, a least-squares fit is used to invert the equation. This is required because otherwise points with a measured asymmetry greater or less than $A_{\rm av} \pm \alpha$ will cause the operation to fail.

From the fitted neutron frequency for each cycle, the ratio of neutron-to-mercury frequencies $R = \frac{f_{\rm n}}{f_{\rm Hg}} \approx 3.84$ is computed. If the caesium magnetometer array is used to correct for gradient drifts, then the ratio R can also be substituted for a gradient-corrected $R = \frac{f_{\rm n} + \langle z \rangle \delta G_{1,0}}{f_{\rm Hg}}$. The error on R is computed considering the counting statistics contributing to the measured asymmetry, the error in the fit parameters, and the error in the mercury and (if used) caesium magnetometer readings and fits.

4.4.4 'RvsE' and 'RvsT' fit

Rather than using the central phase offset Φ arising from each Ramsey fit to measure the average frequency across each E state, the R values computed for each cycle are used. In order to compute the measured EDM during each subset, one must look for a correlation between the applied electric field, and the measured neutron frequency. The ratio R is used as a proxy, to compensate for drifts in the applied magnetic field B_0 .

In the simplest case, a linear relationship $R = p_1 E + R_0$ (with E = -U/d, with U the applied voltage and d = 12 cm the separation of the two electrodes) is fitted. This fit is referred to as the 'RvsE' fit. p_1 is related to the measured EDM by $d_{n,meas} = -\text{sign}(B_0) \cdot p_1 \cdot h \cdot \bar{f}_{\text{Hg}/2}$, where h is Plank's constant and \bar{f}_{Hg} is the mean mercury frequency over the subset. R_0 is then the neutron-mercury frequency ratio, which is used later in the analysis effectively as a gradiometer, taking advantage of the height offset between the neutrons and the mercury atoms. Typical results of such a fit are shown in Figure 4.7, or in the case when magnetic field gradient drifts are corrected using the cesium array in Figure 4.8.



Figure 4.7: Typical RvsE fit, without CsM gradient drift correction.

Particularly in the case where the caesium magnetometers are not used to compensate drifts in the magnetic field gradients, it is possible to additionally fit a time drift component. This is named the 'RvsT' fit. The equation fit then becomes

$$R = R_0 + p_1 \cdot E + R_1(t - \bar{t}) + R_2(t - \bar{t})^2.$$
(4.10)

The fit is often insensitive to R_2 , in these cases it is excluded from the fit. Including the time drift in the fit increases the goodness of fit even in the case where the caesium drift correction is included, as there will still be residual higher-order field drifts not measured



Figure 4.8: Typical RvsE fit, with CsM gradient drift correction.

by the second-order decomposition. Typical results of this fit with and without gradient drift correction using the cesium array are displayed in Figure 4.9 and Figure 4.10.



Figure 4.9: Typical RvsT fit, without CsM gradient drift correction. The red points correspond to cycles taken at positive high voltage, blue to cycles taken at negative high voltage, and grey to cycles taken at zero high voltage.



Figure 4.10: Typical RvsT fit, with CsM gradient drift correction.

4.5 Crossing Point Analysis: A First Exposition

The final stage of the analysis is driven by the need to correct the leading systematic, the mercury 'false EDM'[140, 141, 148–152]. This systematic produces a frequency shift proportional to the applied electric field, exactly the signature of an EDM. It arises due to a conspiracy between radial field components from magnetic field gradients $G_{l,0}$ with odd l, and the Lorentz transformation of the static electric field into the rest frame of the mercury atoms, giving $B_m = -\frac{1}{c^2}v \times E$. This effect applies to both the mercury atoms and to the neutrons, however the effect on the mercury atoms is much larger. When using the frequency ratio $R = f_n/f_{Hg}$ to correct for B_0 drifts, any frequency shifts affecting the mercury atoms are transferred to the measurement. In the case of a pure linear gradient $G_{1,0}$, the frequency shift is directly proportional to $G_{1,0}$. Converting this shift into an apparent nEDM, one computes

$$d_{\mathrm{n\leftarrow Hg}}^{\mathrm{false}} = \frac{\hbar |\gamma_{\mathrm{n}}\gamma_{\mathrm{Hg}}|}{8c^2} r^2 G_{1,0}$$

$$(4.11)$$

$$\approx \frac{G_{1,0}}{1 \mathrm{pT/cm}} \times 4.42 \times 10^{-27} \mathrm{e} \mathrm{~cm},$$
 (4.12)

where r is the radius of the precession chamber. The absolute accuracy of the cesium magnetometer array is too poor to compensate for this effect. However, as previously mentioned, the height difference between the neutrons and the mercury comagnetometer also induces a shift in the ratios of their precession frequencies proportional to $G_{1,0}$. This effect is sufficiently large to enable this quantity to act as an effective gradiometer. The shift in R follows

$$R = \frac{\gamma_{\rm n}}{\gamma_{\rm Hg}} \left(1 + \frac{G_{1,0} \langle z \rangle}{B_0} \right). \tag{4.13}$$

The shift inverts in sign with the direction of B. This immediately suggests the 'crossing lines' technique, first used in the analysis of the previous Sussex-RAL-ILL experiment[20, 21, 148]. Combining the two expressions, one obtains

$$d_{\rm n}^{\rm meas} = d_{\rm n}^{\rm true} + B_0 \frac{\hbar \gamma_{\rm Hg}^2}{8c^2 \langle z \rangle} r^2 \left(R - \left| \frac{\gamma_{\rm n}}{\gamma_{\rm Hg}} \right| \right)$$
(4.14)

The slope of this relationship is equal and opposite for each direction of B. At the crossing point of the two lines, the $G_{1,0}$ must be zero, the measured R will be equal to the independent measurements of $\frac{\gamma_n}{\gamma_{\text{Hg}}}$, and the geometric phase systematic will be removed from d_n^{meas} . The preliminary results of this fit are displayed in Figure 4.11.



Figure 4.11: Result of the crossing point analysis for all data of 2015 and 16. A correction is applied for the Earth's rotation, but corrections for $G_{3,0}$ and $\langle B_T^2 \rangle$ are not applied. In the upper plot, the result from each run taken with B up is marked by a red point, and in the lower plot the result taken with B down is marked by a blue point. The slopes for each of the two lines are equal in magnitude but of opposite sign, and fixed.

4.6 Corrections to the Crossing Point Analysis

In the previous section, the method used to correct for the leading systematic effect is described in the simplest case. However, a number of other systematic effects must be taken account of and be properly compensated. Both direct systematic effects (those inducing a frequency shift correlated with E, mimicking an EDM) and indirect systematic effects (those causing a frequency shift correlated with B, which may induce an error in the final reported EDM through the crossing point procedure) are considered. The measured R can be represented as a function of a sum of several corrections

$$R = \frac{\gamma_{\rm n}}{\gamma_{\rm Hg}} \left(1 + \delta_{\rm EDM} + \delta_{\rm falseEDM} + \delta_{\rm grav} + \delta_{\rm T} + \delta_{\rm LS} + \delta_{\rm Earth} + \delta_{\rm others} \right).$$
(4.15)

These can be divided into terms which will vary according to the applied electric field, namely the effect of a true EDM δ_{EDM} and false EDMs δ_{falseEDM} (which cause direct systematic effects), and those which depend only on the magnetic field such as the gravitational shift δ_{grav} and the effect of transverse inhomogeneities δ_{T} (which may cause indirect systematic effects). There are a number of additional effects such as the scalar mercury light shift δ_{LS} (which depends strongly on the intensity of the mercury probe beam, measured in [144]), and the effect of taking the measurement in a rotating reference frame due to the rotation of the Earth. In this chapter only those relating to magnetic field homogeneity and the Earth's rotation will be discussed.

4.6.1 False EDMs and Gravitational Shifts

In the first exposition, only a mercury false EDM arising from the linear vertical gradient $G_{1,0}$ was considered. However, any vertical gradient $G_{l,0}$ can cause gravitational shifts, and those with odd-l will induce false EDMs.

A more general set of formulae for the false EDM was derived in [152]

$$d_{\rm n}^{\rm false} = -\frac{\hbar v_{\rm RMS}}{2c^2 B_0^2} \left\langle \frac{\partial B_z}{\partial z} \right\rangle \tag{4.16}$$

$$d_{\rm Hg}^{\rm false} = -\frac{\hbar\gamma_{\rm Hg}^2}{8c^2} \left\langle xB_x + yB_y \right\rangle \tag{4.17}$$

$$d_{\mathrm{n\leftarrow Hg}}^{\mathrm{false}} = -\frac{\hbar |\gamma_{\mathrm{n}}\gamma_{\mathrm{Hg}}|}{8c^2} \langle xB_x + yB_y \rangle , \qquad (4.18)$$

which allows for the false EDM due to higher order gradients to be computed

$$d_{n\leftarrow Hg}^{\text{false}} = \frac{\hbar |\gamma_n \gamma_{Hg}|}{8c^2} r^2 \left[G_{1,0} - G_{3,0} \left(\frac{r^2}{2} - \frac{H^2}{4} \right) + G_{5,0} \left(\frac{5r^4}{16} - \frac{5r^2H^2}{12} + \frac{H^4}{16} \right) + \cdots \right].$$
(4.19)

Such higher order gradients also induce a gravitational shift. It can be shown that a stored population of UCN of kinetic energy ϵ (measured at the bottom of the chamber) will have a density of [114]

$$\rho_{\epsilon}(z) = \rho_{\epsilon}(z=0)\sqrt{1 - \frac{mgh}{\epsilon}}, \qquad (4.20)$$

with $mg \approx 1.03$ neV/cm, expressed in convenient units. Each independent, non-mixing, neutron energy group will therefore experience a different average field. In real experiements, the full range of storable UCN will be measured together. In larger vertical gradients, this can lead to 'gravitationally enhanced depolarisation' as the different neutron energy groups drift apart[115, 116]. This can be measured and used to estimate the UCN spectrum though a spin-echo technique[117]. Due to the phenomenon of 'Ramsey wrapping', this dephasing can cause non-linear spectrum-dependant frequency shifts, as the Ramsey technique is only sensitive to the accumulated phase module 2π - neutron populations which drift too far from the central phase can 'wrap around'. In order to avoid this effect, the experiment was always operated at a vertical gradient equivalent to $|G_{1,0}| < 30 \text{pT/cm}$, where the response is expected and observed to be linear[91, 115, 116].

Neglecting these effects, the average field experienced by the mercury atoms or neutrons, considering only terms up to l = 3, can be written as

$$\langle B_z \rangle = G_{0,0} + G_{1,0} \langle z \rangle + G_{2,0} \langle -\rho^2/2 + z^2 \rangle + G_{3,0} \langle z^3 - \frac{3}{2}\rho^2 z \rangle.$$
(4.21)

For both species, $\langle \rho^2 \rangle = r^2/2$. As the mercury atoms fill the volume isotropically, $\langle z \rangle = 0$, $\langle z^3 \rangle = 0$ by symmetry, and $\langle z^2 \rangle = H^2/12$. The magnetic field as measured by mercury is therefore

$$\langle B_z \rangle_{\rm Hg} = G_{0,0} + G_{2,0} \left(\frac{H^2}{12} - \frac{r^2}{4} \right).$$
 (4.22)

Neutrons however clearly do not fill the chamber uniformly. The density $\rho(z)$ is approximated to be a linear function of z

$$\rho(z) \propto 1 + kz. \tag{4.23}$$

Fixing k to produce a given $\langle z \rangle$, one finds

$$\langle z^2 \rangle \approx \frac{H^2}{12}$$
 (4.24)

$$\langle z^3 \rangle \approx \frac{3H^2}{20} \langle z \rangle \,.$$
 (4.25)

These expressions have been validated numerically for a variety of realistic UCN energy spectra, and have shown to produce results accurate to better than 5%, more than sufficient at the current level of sensitivity.

Combining these expressions, one finds

$$\langle B_z \rangle_{\rm n} = G_{0,0} + G_{1,0} \langle z \rangle + G_{2,0} \left(\frac{H^2}{12} - \frac{r^2}{4} \right) + G_{3,0} \left(\frac{3H^2}{20} - \frac{3r^2}{4} \right) \langle z \rangle \,. \tag{4.26}$$

The gravitational shift can be expressed as an equivalent linear gradient G_{grav}

$$\delta_{\text{grav}} = \frac{\langle B_z \rangle_{\text{n}}}{\langle B_z \rangle_{\text{Hg}}} - 1 = \pm \frac{G_{\text{grav}}}{|B_0|}, \qquad (4.27)$$

where

$$G_{\rm grav} = G_{1,0} + G_{3,0} \left(\frac{3H^2}{20} - \frac{3r^2}{4}\right).$$
(4.28)

The gravitational shift effect can therefore, in the case of higher order gradients beyond $G_{1,0}$, be used as a powerful measurement of G_{grav} . Rewriting Equation 4.19 in terms of G_{grav} , one sees that the crossing point analysis first presented already compensates a substantial part of the systematic effect from $G_{3,0}$

$$d_{\rm n \leftarrow Hg}^{\rm false} = \frac{\hbar |\gamma_{\rm n} \gamma_{\rm Hg}|}{8c^2} r^2 \left[G_{\rm grav} + G_{3,0} \left(\frac{r^2}{4} + \frac{H^2}{10} \right) \right] .$$
(4.29)

In the analysis, the effect of G_{grav} is compensated through the crossing lines fit, but the remaining false EDM from $G_{3,0}$ must be compensated on a run-by-run basis using values measured in the offline field mapping measurements as described in chapter 5.

4.6.2 The Earth's Rotation

The experiment takes place on Earth, with the B_0 field oriented approximately vertically up or down, and so naturally the measurement is made in a rotating reference frame. This effect induces a shift in the measured frequency, which inverts with the sign of B [86, 175, 176]. The correction can be computed as

$$\delta_{\text{Earth}} = \mp \left(\frac{f_{\text{Earth}}}{f_{\text{n}}} + \frac{f_{\text{Earth}}}{f_{\text{Hg}}} \right) \cos \theta, \qquad (4.30)$$

where $f_{\text{Earth}} \approx 11.6 \text{mHz}$ and θ is the angle between the earth's axis of rotation and the quantisation axis of the system, 42.4833° at PSI. As this effect causes an opposite sign of shift for each direction of B, the crossing point of the two lines is shifted. If uncorrected, this would result in a shift in the measured EDM of $\Delta_{d_n} \approx -2.6 \times 10^{-26}$ e cm. If the quantisation axis (i.e. the direction of the magnetic field) is correlated with E or B, then this may induce an additional direct or indirect systematic effect.

4.6.3 $< B_T^2 >$ Correction

Inhomogeneities in the transverse magnetic field can also cause shifts in the neutronmercury frequency ratio R. Ultracold neutrons have velocities of the order < 5 m/s, and precession frequencies ≈ 30 Hz, meaning that they effectively follow inhomogeneities of the field as they cross the chamber, putting them firmly in the adiabatic regime, with $f_n \sim \left\langle \left| \vec{B} \right| \right\rangle$. Meanwhile, the thermal mercury atoms cross the chamber many times during one precession period², so do not track local inhomogeneities in the field, and therefore precess according to $f_{\rm Hg} \sim \left| \left\langle \vec{B} \right\rangle \right|$. This effect causes the neutron precession frequency to always be increased relative to the mercury precession frequency. If the size of the frequency shift is larger for one sign of *B* than the other, this can result in an indirect systematic effect. Additionally, as different trimcoil currents were applied to optimise the field after each degaussing, the size of this shift can vary from run to run. The size of the shift can be computed as

$$\delta_T = \frac{\langle B_T^2 \rangle}{2B_0^2},\tag{4.31}$$

where $\langle B_T^2 \rangle$ is the RMS transverse inhomogeneity, given by

$$\langle B_T^2 \rangle = \langle (B_x - \langle B_x \rangle)^2 + (B_y - \langle B_y \rangle)^2 \rangle.$$
 (4.32)

This can be expressed as a function of the gradients $G_{l,m}$, or computed numerically. Considering only terms up to l = 1, the expression reads

$$\left\langle B_T^2 \right\rangle_{\rm LO} = \frac{r^2}{2} \left(G_{1,-2}^2 + G_{1,2}^2 + \frac{1}{4} G_{1,0}^2 \right) + \frac{H^2}{12} \left(G_{1,-1}^2 + G_{1,1}^2 \right).$$
 (4.33)

Similar expressions can be computed to arbitrarily high order using a computer algebra program, and in the full analysis terms up to and including l = 6 are considered.

The gradients $G_{l,m}$ are obtained from offline field mapping measurements, as detailed in chapter 5.

4.7 The Crossing Point Analysis: Recapitulation

The crossing point analysis described in section 4.5 can be modified to correct for the effects described in section 4.6. A series of substitutions are made

$$d_{\rm n}^{\rm corr} = d_{\rm n}^{\rm meas} - \frac{\hbar \left|\gamma_{\rm n}\gamma_{\rm Hg}\right|}{8c^2} r^2 \left(\frac{r^2}{4} + \frac{H^2}{10}\right) G_{3,0} \tag{4.34}$$

$$R^{\rm corr} = R^{\rm meas} / (1 + \delta_T + \delta_{\rm Earth} + \delta_{\rm other}). \tag{4.35}$$

With these new definitions, one finds

$$d_{\rm n}^{\rm corr} = d_{\rm n}^{\rm true} + B_0 \frac{\hbar \gamma_{\rm Hg}^2}{8c^2 \langle z \rangle} \left(R^{\rm corr} + \left| \frac{\gamma_{\rm n}}{\gamma_{\rm Hg}} \right| \right).$$
(4.36)

 $^{^{2}}$ The chamber is operated at a sufficiently low pressure that the mercury atoms follow a ballistic trajectory with a mean free path substantially larger than the chamber diameter, however if a buffer gas were used, the mean free path could be reduced to bring them into the adiabatic regime

This may be fitted to the measured EDM and R value from each subset to find the EDM value, corrected for the aforementioned systematics. Other possible frequency shifts can be incorporated into δ_{other} if they may vary from run to run, otherwise they can be corrected after the crossing point fit.

4.8 Sensitivity and Result

The theoretical maximum sensitivity of an EDM type experiment is given by

$$\sigma_{d_{n}} = \frac{\hbar}{2ET\alpha\sqrt{N_{0}\cdot m}},\tag{4.37}$$

where *m* is the total number of measurement cycles. Using typical values for the Ramsey visability α , electric field $E = \pm 132 \text{ kV}/12 \text{ cm} = \mp 11 \text{ kV/cm}$ and precession time T = 180 s, we compute a raw statistical sensitivity of 0.94×10^{-26} e cm (1 σ). However, this cannot be reached in real world conditions. The experiment is not operated at the most optimal configuration for a number of reasons. First, the working points chosen to measure the Ramsey resonance are slightly offset from the steepest (most sensitive) part of the curve, in order to gain some knowledge of nuisance parameters such as the visibility α and detector efficiency asymmetries. Second, a portion of the datataking must be undertaken at zero electric field, in order to enable the working points to be chosen in a way not biased by *E*-correlated frequency shifts such as those introduced by the blinding scheme. Another reason that the theoretical sensitivity cannot be reached is the uncertainty on the reading of the mercury magnetometer, which is used to correct drifts of the magnetic field.

After the crossing lines procedure, the statistical sensitivity of the experiment is $\sigma_{d_n} = 1.14 \times 10^{-26}$ e cm (1 σ). Compared to the current best result of $d_n = (-0.21 \pm 1.53 (\text{stat.}) \pm 0.99 (\text{syst.}) \times 10^{-26}$ e cm, this represents an improvement. The collaboration also targets a substantial improvement in the final systematic uncertainty, with a total error budget target of 5×10^{-27} e cm. The combination of these two factors will lead to an improvement in the first improvement in the measurement of the nEDM for ten years, and this limit will likely stand for several years until it is succeeded by the n2EDM experiment from the same collaboration.

Chapter 5

Mapping of the Magnetic Field within the nEDM apparatus



Figure 5.1: Photo of the mapper as installed inside the nEDM vacuum chamber, with the radial \hat{r} , tangential $\hat{\phi}$ and vertical \hat{z} directions annotated.

5.1 Motivation

In the PSI nEDM experiment, the largest systematic errors arise from the gradients and nonuniformities in the magnetic field within the precession chamber. While the PSI nEDM experiment possesses both a ¹⁹⁹Hg magnetometer to measure the average field within

the chamber and an array of 16 optically pumped scalar ¹³³Cs magnetometers in the vicinity of the chamber, these do not provide a complete picture of the magnetic field environment within the experiment. Firstly, the finite number of caesium magnetometers allows the field to be expanded up to only quadratic order, while it has been shown that the $G_{3,0}$ magnetic field component (see [141] for original definition of basis functions, summarised in Appendix A) can produce a substantial systematic effect. Additionally, as scalar magnetometers measuring only the modulus of the field $|\vec{B}|$, they offer only poor access to purely transverse field modes such as $G_{1,\pm 2}$ or $G_{2,\pm 3}$, as they add only in quadrature to the large 1 µT nominal field ¹. Finally, the caesium magnetometers each suffer from unexplained offsets in their readings, which make them unsuitable to measure the absolute gradient. However, they are still useful to correct for relative drifts in the vertical gradient (see chapter 4).

To allow for the field inside the experiment to be directly measured, the mapper was conceived. This robot (photographed in Figure 5.1) consists of a cart mounted on an arm, itself mounted on a rotational axle within the vacuum tank, allowing independent movement in the radial and vertical directions, and rotation about the vertical axle. On one end of the cart, a fluxgate magnetometer is mounted inside a tube, allowing rotation nominally about the radial axis to allow the calibration of the sensor offsets. On the opposing end, a free alignment precession (FAP) fully optical scalar cesium magnetometer is mounted. This kind of magnetometer was chosen because it is free of many of the kinds of systematic offsets and frequency shifts present in other magnetometers [177]. In previous mapping campaigns, the mapper was used with just the fluxgate sensor or a vector cesium magnetometer module [178–180].

The mapper was designed and produced by the laboratories and workshops of the Laboratorie de Physique Corpsuculaire de Caen.

5.2 Accuracy Requirements

The measurements of the mapping campaigns were used to constrain systematic effects arising from two sources: the false EDM effect due to the cubic $G_{3,0}$ gradient, and the indirect systematic introduced by transverse inhomogeneities across the chamber, quantified by $\langle B_T^2 \rangle$. The accuracy requirements for the mapping campaign can be derived working backwards from the overall systematic error budget of 5×10^{-27} e cm.

¹These modes can however be accessed by applying known transverse fields and observing the parabolic dependence of the measured field on the applied transverse field [127]

As previously discussed in section 4.6, the gradient $G_{3,0}$ induces a false EDM of the mercury atoms, which is injected into the neutron EDM measurement when the mercury comagnetometer is used to compensate drifts in the B_0 field. The size of this effect is

$$\delta_{d_{n}^{\text{False,Cubic}}} = \frac{\hbar |\gamma_{n}\gamma_{\text{Hg}}|}{8c^{2}} r^{2} \left(\frac{r^{2}}{2} - \frac{H^{2}}{4}\right) G_{3,0}$$
(5.1)

$$= 1.10 \times 10^{-24} \text{ecm} \frac{G_{3,0}}{1 \text{pT/cm}^3},$$
(5.2)

where r is the chamber radius and H the total chamber height.

However, as explained in section 4.6, this gradient also induces a gravitational shift, which allows the effect to be partially compensated by the crossing lines procedure used to correct the linear false EDM effect due to $G_{1,0}$

$$\delta_{d_{n}^{corr}} = \frac{\hbar |\gamma_{n}\gamma_{Hg}|}{8c^{2}} r^{2} \left(\frac{r^{2}}{4} + \frac{H^{2}}{10}\right) G_{3,0}$$
(5.3)

$$= 7.01 \times 10^{-25} \text{ecm} \frac{G_{3,0}}{1 \text{pT/cm}^3}.$$
 (5.4)

Thus, the requirements are somewhat relaxed.

Transverse inhomogeneities cause a shift in the precession frequency in the adiabatic neutrons relative to the nonadiabatic mercury comagnetometer atoms. If $\langle B_T^2 \rangle$ is different for the two directions of B_0 , then the crossing point d_n will be affected. The fractional shift in R is given by

$$\delta_T = \frac{\langle B_T^2 \rangle}{2B_0}.\tag{5.5}$$

Such a shift will result in an systematic effect after the crossing lines correction procedure on the nEDM measured each run of

$$\delta_{d_{n}} = \pm \frac{\partial d_{n}}{\partial R} \cdot \delta_{T} \cdot R_{0} \tag{5.6}$$

$$= \pm \frac{B_0 \hbar \gamma_{\rm Hg}^2}{8c^2 \langle z \rangle} \cdot \frac{\langle B_T^2 \rangle}{2B_0} \cdot \left| \frac{\gamma_{\rm n}}{\gamma_{\rm Hg}} \right|$$
(5.7)

$$= \pm \frac{\hbar \left| \gamma_{\rm n} \gamma_{\rm Hg} \right|}{16c^2 B_0 \left\langle z \right\rangle} \left\langle B_T^2 \right\rangle \tag{5.8}$$

$$\approx \pm 6.28 \times 10^{-27} \text{ecm} \cdot \frac{\langle B_T^2 \rangle}{\ln T^2},\tag{5.9}$$

for $\langle z \rangle = 3.5$ mm, where $\langle z \rangle$ is the centre of mass shift between the mercury atoms and the neutrons within the chamber (explained more thoroughly in chapter 4). As such a shift will always increase R, once the crossing point analysis is applies, if the effect is equal for both signs of B, the effect will cancel. The relevant quantity is actually $\langle B_T^2 \rangle_{\rm up} - \langle B_T^2 \rangle_{\rm up}$, which will shift the crossing point of the two lines up or down.

In order that systematic effects related to the mapping contribute less than 5×10^{-27} ecm to the overall systematic uncertainty, strict requirements are placed on the accuracy of the field maps. For this requirement to be reached, $G_{3,0}$ must be measured with an accuracy of the order 0.005 pT/cm³, and $\langle B_T^2 \rangle$ with an accuracy of the order 0.4 nT². For $G_{3,0}$ this corresponds to a vertical field of the order 140pT at the furthest point probed by the mapper during normal running. For $\langle B_T^2 \rangle$ the requirements are at first sight less strict, with an RMS field across the precession chamber of $\sqrt{0.4nT^2} \approx 600pT$, however this field is also challenging to measure because small angular misalignments of radial and transverse sensor into the vertical direction (with a large vertical field of nominally 1μ T) can cause large fake fields in these sensors.

5.3 Mapper Hardware

The mapper itself consists of a radial arm, mounted on a vertical axle, which can rotate and slide freely within a pair of bearings mounted on the upper and lower flanges of the experiment vacuum tank. These flanges are normally used to mount the HV feedthrough from above and for the UCN guide from below when the experiment is installed. The radial motion is driven by a rack-and-pinion system, with the pinion gear being driven by a vertical shaft which runs inside the lower part of the vertical axle. This shaft is directly driven by a stepper motor, itself mounted at the lowest end of the shaft to rotate with the main vertical axle. The rack is mounted to the underside of the cart. The rotation of the entire setup around the vertical axis is performed by a second stepper motor. This is mounted such that it moves vertically with the mapper. A belt drive connects the motor spindle with the mapper axle. The entire mapper including the radial and rotational driving mechanisms can be raised or lowered by means of a vertical drive system. This is a commercial unit, originally intended to raise or lower a dentists' chair. A stepper motor drives a threaded rod, which meshes with a thread on the lower part of the assembly to give a vertical motion.

In the original design, each axis position was read by a potentiometer, to give an independent feedback of the absolute position of each axis and to prevent errors from allowing the motors to drive too far and destroy the apparatus. The radial potentiometer was a string potentiometer type unit, with the string attached to a monofilament Dyneema cord mounted to the cart via a system of pulleys. This cord however proved to be difficult to manipulate and regularly became tangled, and did not give an accurate enough reading. Additionally, this cord fouled the mounting for the vertical guide coil. The decision was therefore made to remove the potentiometer from the system and instead add a small counterweight in order to replace the tension this placed on the cart. The ϕ potentiometer was a string type potentiometer, and the end of the string was simply fixed to the rotating axle such that it wrapped around when the mapper rotated. The vertical potentiometer was more sturdy in construction, and consisted of a large linear potentiometer unit mounted to the vertical lifting mechanism.

To provide for an accurate calibration of the zero offsets of each fluxgate sensor, the calibration mechaism allows for a 180° rotation of the fluxgate about the radial axis. This is achieved by the fluxgate being mounted in a tube with a bearing at either end and two notches cut to give two stable positions. A helix is cut into the body of the tube, and the mechanism is actuated by a pneumatic drive driving a long finger penetrating from outside of the system from above running through the helical groove.

The B_0 coil and trimcoils are wound on the outside of the vacuum tank, with the exception of the coils UCK, HVK ('Kamin' coils, which are wound around 'chimneys' penetrating the layers of the μ -metal) and the vertical guide coils VCGT, VGCC and VGCB, which are wound on a separate support structure formed of epoxy resin. During nEDM running, this structure is mounted onto the vertical glass UCN guide, and penetrates all four layers of μ -metal, meeting the edge of the vacuum tank. One change for the 2017 campaign was the design of a special mounting for the vertical guiding coil mounting during mapping. The mapper axle runs through the centre of the vertical guiding coil former, while the coil is supported from below. The coil was positioned before the dismounting of the appartus: first the field produced by the coil as mounted onto the vertical guide was measured using the transverse field technique [127], then the vertical guide was removed, the new structure mounted, and the position was trimmed until the field measured by the cesium magnetometer array was matched. This was achieved to within 100pT.

5.4 Datataking Summary

Using the apparatus described in this chapter, three distinct measurement campaigns were performed, in 2013 and 2014 before the main nEDM datataking, and additionally in 2017 after datataking. The aims of the 2013 and 2014 campaigns were to charactarise the magnetic properties of the μ -metal shields and to measure the effect of each trimcoil. In these campaigns several maps were taken of the unmodified B_0 field, plus maps of each trimcoil. From these two ingredients, it is possible to predict the magnetic field for any given field configuration inside the apparatus. In 2017, there were four primary aims: first, to assess the reproducibility of the degaussing process; second, to assess if there had been any changes in the magnetic configuration since 2014; third to assess the effect of the vertical guiding field coils on the internal gradients (which were not included in previous campaigns) and finally to validate the assumption of linear superposition of each of the fields, allowing us to find the gradients of the field for arbitrary trimcoil configurations. In order to achieve this, a 24-hour measurement cycle was adopted for much of the datataking. Each day, the magnetic field direction would be reversed, the shield degaussed, and then 3-4 configurations were mapped, including always an unmodified B_0 up or down map, plus several of the field configurations used during datataking in 2015 and 2016. The repeated B_0 maps allow for assessment of the reproducibility of the degaussing process, while the configuration maps allow a more direct determination of the gradients used during datataking and assessment of the validity of the assumption of linearity.

A single configuration measurement cycle consisted of a series of maps. First, a calibration map was performed, a special cycle involving three full rotations of the system, with each of the three magnetometers located at the centre of the tank. The fluxgate was then flipped using the calibration mechinism and the three rings were repeated, before returning the system to the original state. With this process, it is possible to fix the offsets of each sensor to correct for a linear drift over the measurement period. The most important map was a map of the main volume. Full measurement rings were taken with a spacing of $\Delta \phi = 10^{\circ}$ every $\Delta R = 20$ mm from 0mm to 340mm. This was repeated at several planes from z = -160mm to z = +160mm. Afterwards, another calibration map was taken. In 2017, these were often supplemented with so-called z200 maps, which consisted of a series of measurements with the vertical fluxgate at the central axis of the tank, moved up and down along the z-axis. These were intended to give a measurement of the gradients $G_{l,0}$ independent of some of the systematics. However, the statistical precision of these maps is inferior to the full maps, and it was possible to constrain the systematics well enough to render these measurements redundant.

5.5 Known Imperfections of the Mapper

Although the mapper robot is of excellent mechanical construction, there are a few imperfections in the action. These arise partially due to the strict requirements in order to ensure an accurate measurement of the magnetic field. In order to measure small nT transverse components on a relatively large 1uT vertical field, in addition to accurate positioning, the angular alignment of the system must be ensured. As a rule of thumb, a 1mrad angular misalignment will cause a false field of 1nT in a transverse sensor. Therefore, angular misalignments must be controlled to better than 0.1mrad for a 100pT accuracy. Typical gradients in the apparatus are of order 100 pT/cm, so positioning precision is relatively unimportant, requiring a reproducibility of 2 mm or better to match the 20pT noise level of the fluxgate. However, systematic errors in the measured gradients may arise if there is a systematic positioning error. The mapper contains no metallic parts inside the μ -metal shields other than those required for the operation of the magnetometry system (which were separately verified to not affect the measurement). Apart from the obvious need to avoid magnetic contamination, it was previously found that fluxgate magnetometers do not perform correctly close to large conductive masses due to induced eddy currents. This requirement meant that standard high-precision construction from a non-ferrous metal was not suitable, and instead it was chosen to use PEEK as a primary construction material. This plastic has been verified to be non-magnetic to a very high level. PEEK is however difficult to machine precisely. As a result, the precision of the radial and rotational motions is limited. Additionally, the linear guiding assembly in the radial direction is substantially warped, which cause angular misalignments of several mrad.

The vertical motion system also presents mechanical challenges: while the system is actuated by a precise commercial drive unit, the axle is not perfectly rigid, and is fixed at several points: the connection to the drive unit, and bearings at the top and bottom of the tank. If the three points are not perfectly aligned, or the motion of the vertical drive is not collinear with this axis, the axle will be bent to a 'banana' shape, which will result in the headings of the three sensors and the axis of rotation to depend on the height, causing possible issues.

The fluxgate sensor also has several intrinsic limitations. The model used is a special model FL3-2 of three-axis fluxgate produced by Stefan Mayer Instruments, adapted for a measurement range of up to $\pm 2\mu$ T. It consists of three independant, nominally orthogonal sensors, labelled X, Y and Z, with Y pointing in the $-\hat{z}$ direction, Z pointing radially, and positioned 20mm offset from the first sensor, and X pointing to $-\hat{\phi}$, and 40mm offset from the Y sensor. The orthogonality of these three axes is specified to be accurate to better than 0.5°, corresponding to an 8.7 mrad uncertainty. This is large, so the precise orientation of each sensor must be determined. Additionally, each sensor has a zero offset specified to be < 5 nT, however in practice this is often larger, as the quoted specification is for a new sensor with a recently-degaussed ferrite. This offset value can drift substantially over time, so it must be measured in situ frequently.

In order to accurately analyse our data, each of these effects must be charactarised. A model of how each of these effects will manifest in the measurements is presented in section 5.6, and a series of techniques to measure each of these is described in section 5.7.

5.6 General Model of Fluxgate Mapper Misalignments, Nonorthogonalities and Offsets

In order to relate the measurements taken using the field mapper and the true external field, it is necessary to formulate a model of the misalignment angles which can be measured or extracted from our data. The ultimate goal is to be able to express the orientations of the three vector fluxgate sensors, denoted by \hat{X} , \hat{Y} and \hat{Z} in terms of the external coordinate system, with basis vectors \hat{r} , $\hat{\phi}$, \hat{z} . This will eventually depend on the orientation of the fluxgate rotation tube, and on the radial position R. Nominally, in the normal orientation, the \hat{X} sensor points in approximately the $-\hat{\phi}$ direction, the \hat{Y} sensor in the $-\hat{z}$ direction and the \hat{Z} sensor in the radial \hat{r} direction. In the inverted position, the \hat{X} sensor points in approximately the $+\hat{\phi}$ direction. The axis of the fluxgate calibration rotation is roughly coincident with \hat{r} , and the axis of the mapper arm rotation is roughly coincident with $-\hat{z}$.

First, the three sensor alignments are expressed in an intermediate coordinate system, with orthonormal basis vectors $\hat{r'}$, $\hat{\phi'}$, $\hat{z'}$. This is aligned so that at a nominal R = 0position, $\hat{\phi'} = \hat{\phi}$ and so that $\hat{r'}$ and $\hat{z'}$ lie in the same plane as \hat{r} and \hat{z} , with $\hat{r'}$ coincident with the axis of rotation of the fluxgate tube. Each of the vectors \hat{X} , \hat{Y} , \hat{Z} are expressed in terms of their azimuth and polar angles $\phi_{X,Y,Z}$ and $\theta_{X,Y,Z}$ as measured in this intermediate frame.

In the fluxgate tube coordinate system,

$$\hat{X}, \hat{Y}, \hat{Z}_{\text{local}} = \begin{pmatrix} \cos \phi_{X,Y,Z} \sin \theta_{X,Y,Z} \\ \sin \phi_{X,Y,Z} \sin \theta_{X,Y,Z} \\ \cos \theta_{X,Y,Z} \end{pmatrix}.$$
(5.10)

Then, a rotation by an angle ψ around the \hat{r}' vector is added, in order to allow for the calibration rotation action. This is the same for all three sensors as they must rotate together. In the normal position, ψ is nominally zero, and in the inverted position, ψ is nominally π .

$$\hat{X}, \hat{Y}, \hat{Z}_{\text{local}} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & \sin \psi & -\sin \psi \\ 0 & \sin \psi & \cos \psi \end{pmatrix} \begin{pmatrix} \cos \phi_{X,Y,Z} \sin \theta_{X,Y,Z} \\ \sin \phi_{X,Y,Z} \sin \theta_{X,Y,Z} \\ \cos \theta_{X,Y,Z} \\ \sin \phi_{X,Y,Z} \sin \theta_{X,Y,Z} \\ \sin \phi_{X,Y,Z} \sin \theta_{X,Y,Z} \\ \cos \theta_{X,Y,Z} \sin \phi_{X,Y,Z} \sin \theta_{X,Y,Z} \\ \sin \phi_{X,Y,Z} \cos \psi + \sin \phi_{X,Y,Z} \sin \theta_{X,Y,Z} \\ \sin \phi_{X,Y,Z} \sin \psi \\ \cos \theta_{X,Y,Z} \cos \psi + \sin \phi_{X,Y,Z} \sin \theta_{X,Y,Z} \\ \sin \phi_{X,Y,Z} \sin \psi \end{pmatrix}$$
(5.11)

The next step is to determine the relation between the intermediate coordinate system on the cart and the global coordinate system. This depends on the position of the cart, as the rails are substantially warped. The conversion between coordinate systems is performed using a series of rotations by angles $\alpha(R)$ ('roll' around \hat{r}), $\beta(R)$ ('pitch', or 'tilt', around $\hat{\phi}$) and $\gamma(R)$ ('yaw' around \hat{z}). This R is defined by the position of the cart (calibrated so that at R = 0mm the vertical Y sensor is at the centre, with the radial Z at a radius of +20mm and the tangential X sensor at a radius of +40mm), rather than the position of each sensor. The conventions are chosen such that α and β make the upwards facing angle, with positive γ having the same sense as positive ϕ .

The heading of each fluxgate component in the global coordinate system can be computed,

$$\hat{X}_{\text{global}} = R_r(\alpha)^{-1} R_{\phi}(\beta)^{-1} R_z(\gamma)^{-1} \hat{X}_{\text{local}}.$$
 (5.12)

Each fluxgate sensor reads the component of the \vec{B} field along its axis \hat{X} , \hat{Y} or \hat{Z} , plus an arbitrary offset due to the magnetisation of the ferrite core,

$$\vec{B}_{\rm FG,meas} = \begin{pmatrix} \hat{X} \cdot \vec{B} + \delta B_X \\ \hat{Y} \cdot \vec{B} + \delta B_Y \\ \hat{Z} \cdot \vec{B} + \delta B_Z \end{pmatrix}.$$
(5.13)

The 3-axis fluxgate magnetometer used for the mapping consists of three independent, nominally orthogonal, single axis fluxgates mounted inside a single device, offset from each other by 20mm. The fluxgate is mounted such that at the position R = 0, z = 0, the Y fluxgate is at the centre of the vacuum tank, pointing approximately along the $-\hat{z}$ axis. The Z fluxgate is located 20mm behind the Y fluxgate, pointing in approximately the radial direction, and the X fluxgate is 40mm behind the Y fluxgate, pointing in approximately the $-\hat{\phi}$ direction. Thus, when the mapper is positioned at R = 0, the X and Z sensors are effectively at a negative radius. The effect of this is to invert the \hat{r} and $\hat{\phi}$ components of the headings as measured in the global cylindrical system when these sensors are at negative radius. Additionally, a negative radius is outside the range normally allowed in a cylindrical coordinate system. To convert this to a valid position, one must transform $r \to |r|$ and $\phi \to \phi + 180^{\circ}$. This is done implicitly through the rest of the chapter.

In summary, a model to describe the largest observed imperfections in the construction of the mapper device has been constructed. It is formulated in terms of:

- Three offsets δB_X , δB_Y and δB_Z to describe the internal zero-offset in each fluxgate sensor
- Six angles $\phi_{X,Y,Z}$ and $\theta_{X,Y,Z}$ to describe the alignment of each fluxgate sensor with respect to the axis of rotation of the fluxgate calibration rotation mechanism
- One angle ψ to describe the orientation of the calibration mechanism (with two stable positions, nominally $\psi_{\text{norm}} = 0$ and $\psi_{\text{inv}} = \pi$)
- Three angles, α , β and γ to describe the orientation of the axis of rotation of the fluxgate calibration mechanism relative to the global coordinate system. These are allowed to depend on the radius of the mapper cart to account for imperfections in the rails.

Many of these parameters are indistinguishable in effect, or distinguishable only in certain magnetic field configurations. In the next section, the measurements to characterise and measure these angles is described.

5.7 Determination of Mapper Misalignments, Nonorthogonalities and Offsets

In the previous section, a model describing several possible angular imperfections of the mapper mathematically was described. In order to make use of it, the free parameters must be determined: the six angles $\phi_{X,Y,Z}$ and $\theta_{X,Y,Z}$, the value of ψ in both the normal and inverted fluxgate position, the values of α , β and γ at all radii measured. Any effects not accounted for in the model must be determined to be sufficiently small. As an initial validation step, the profile of the mapper rails was measured directly, to determine if curvature of the rails really can explain the observed radius-dependant misalignments. A set of dedicated measurements using an inclinometer were undertaken during the installation of the mapper. To characterise the behaviour with the fluxgate magnetometer

installed, a combined fit of the magnetic field and the misalignments of the mapper was performed. The fluxgate offsets were regularly measured throughout the measurement period through the calibration proceedure described below.

5.7.1 Direct Measurement of Mapper Arm Profile

In order to verify the previously constructed model further and confirm that the observed misalignment angles really do arise from a deformation of the rails, the profile of these rails was measured directly using a Mitutoyo Legex 776 measuring machine, with a measurement accuracy of up to $0.2 \ \mu\text{m}$. The surface of each of the lower two surfaces contacted by the mapper cart wheels during a normal map ($R = 0 \ \text{mm}$ to $R = 340 \ \text{mm}$) was mapped at a resolution of 1 cm. The results of this measurement are shown in Figure 5.2. Additionally, the two side rails were mapped at a resolution of 2 cm. The surfaces are sufficiently smooth that linear interpolation can be used between these points.



Figure 5.2: Uncorrected measured height profile of the mapper rails. This shows a clear deviation from straightness, on the order 1mm.

The cart sits upon 4 wheels, which rest upon the rails. Except in the case of perfect alignment, only three will make contact at any point. Which three wheels will make contact depends on the weight distribution, the distribution of other forces on the cart and the precise profile of the rails, this cannot trivially be computed a priori. The height differences of these three points will define the slope. The tilt (β) and roll (α) of the cart at a specified measurement radius R will be given by:

$$\begin{aligned} \alpha(R) &= \frac{h_{\rm RH}(R-R_0) - h_{\rm LH}(R-R_0)}{W} & \text{Both rear wheels touching} \\ \alpha(R) &= \frac{h_{\rm RH}(R-R_0+L) - h_{\rm LH}(R-R_0+L)}{W} & \text{Both front wheels touching} \\ \beta(R) &= \frac{h_{\rm LH}(R-R_0+L) - h_{\rm LH}(R-R_0)}{L} & \text{Both left wheels touching} \\ \beta(R) &= \frac{h_{\rm RH}(R-R_0+L) - h_{\rm RH}(R-R_0)}{L} & \text{Both right wheels touching} \end{aligned}$$

where R_0 is the offset position on the rail of the rear wheels from the centre when R is nominally 0, L is the distance between the centres of the front and rear wheels, W is the distance between the centres of the left and right pairs of wheels and $h_{\text{LH,RH}}(R)$ are the profiles of the right and left rails respectively. From the technical drawings of the mapper one finds L = 336 mm, $R_0 = 363$ mm and W = 80 mm.

The cart itself is known to be warped. This was not directly measured, so it is not possible to ascertain the absolute values of α and β . However, it is known that the cart was initially set up so that when the cart was located centrally $R = -R_0 + L/2$ all four wheels made contact. This is subtracted from the height for each wheel.

Using

$$\Delta h_{\rm FR} = h_{\rm RH}(+L/2)$$
Front right wheel offset
$$\Delta h_{\rm RR} = h_{\rm RH}(-L/2)$$
Rear right wheel offset
$$\Delta h_{\rm FL} = h_{\rm LH}(+L/2)$$
Front left wheel offset
$$\Delta h_{\rm RL} = h_{\rm LH}(-L/2)$$
Rear left wheel offset

The previous expressions for α and β are recast to account for this

$$\alpha (R) = \frac{(h_{\rm RH} (R - R_0) - \Delta h_{\rm RR}) - (h_{\rm LH} (R - R_0) - \Delta h_{\rm RL})}{W}$$

$$\alpha (R) = \frac{(h_{\rm RH} (R - R_0 + L) - \Delta h_{\rm FR}) - (h_{\rm LH} (R - R_0 + L) - \Delta h_{\rm FL})}{W}$$

$$\beta (R) = \frac{(h_{\rm LH} (R - R_0 + L) - \Delta h_{\rm FL}) - (h_{\rm LH} (R - R_0) - \Delta h_{\rm RL})}{L}$$

$$\beta (R) = \frac{(h_{\rm RH} (R - R_0 + L) - \Delta h_{\rm FR}) - (h_{\rm RH} (R - R_0) - \Delta h_{\rm RR})}{L}$$

Both rear wheels touching Both front wheels touching Both left wheels touching Both right wheels touching

Using this method, the values of $\alpha(R)$ and $\beta(R)$ are found in both extreme cases of one wheel being fully raised and three wheels making contact. These results are presented in Figure 5.3 and Figure 5.4.



Figure 5.3: Misalignment α versus cart radius, computed in the cases where the slope is defined by the two front wheels, and the two rear wheels.



Figure 5.4: Misalignment β versus cart radius, computed in the cases where the slope is defined by the two left wheels, and the two right wheels.

5.7.2 Analysis of Inclinometer Measurements

To ensure proper installation of the mapper and verify the mechanical performance, a pair of 2-axis inclinometers were mounted instead of the usual measurement fluxgate. One was mounted inside the calibration tube mechanism in place of the fluxgate, while the other was mounted directly on the base of the mapper cart. Both inclinometers were model Kelag KAS901-52, with a resolution of 0.001° (0.017 mrad), measurement repeatability of 0.014° (0.24 mrad) and long term stability of $0.036^{\circ}(0.63 \text{ mrad})$. To verify that the alignment had remained constant and good, at the end of datataking the inclinometers were once again installed. Each time several maps were taken. In general, these consisted of complete 'rings', consisting themselves of several measurements at the same radius and height around a full rotation of the mapper. At each point the tilt and roll as measured by each inclinometer was recorded.

To analyse the data, each ring the motion of the mapper was modelled as a pure 360° rotation of a rigid arm about the vertical axis, which may be tilted with respect to gravity. This gives

$$\theta_{\text{tilt}} = \theta_{\text{grav},x} \cos \phi + \theta_{\text{grav},y} \sin \phi + \beta$$
(5.14)

$$\theta_{\rm roll} = \theta_{{\rm grav},y} \cos \phi - \theta_{{\rm grav},x} \sin \phi + \alpha \tag{5.15}$$

 α and β are identified with the tilt and roll angles α and β described in section 5.6. $\theta_{\text{grav},x}$ and $\theta_{\text{grav},y}$ are the angles between the rotation axis and the true vertical axis (as defined by gravity), projected onto the xz and yz planes respectively. Thus, it is possible to determine the alignment of the axis of rotation with respect to gravity as well as the alignment of the cart with respect to this axis. This can be computed separately for each inclinometer. α and β may depend on the radius, however $\theta_{\text{grav},x}$ and $\theta_{\text{grav},y}$ should not depend on the radius.

Upon installation, the first task was to ensure that the vertical axis of the mapper was fully aligned. As described previously, as the mapper vertical axis (formed of solid PEEK, which may flex under load) is fixed in three places, if these are not fully aligned the axis will be bent into a parabola, which will manifest as a global tilt ($\theta_{\text{grav},x}$ and $\theta_{\text{grav},y}$) which will depend on the vertical position. Initially, the two tank mountings were fixed such that the mapper axis was vertical when mounted between them, then the axis was connected with the motor box. The global tilt as a function of the height was measured using the inclinometers, and the position and alignment of the motor box was trimmed until no further improvement in the straightness of the axis was possible. The result at the end of trimming is displayed in Figure 5.5, showing a final global tilt of less than 1 mrad (which arises because the two tank fixings were installed with the tank open, installing the tank door weighing approx 100kg causes some tilt.), with a 'banana' deformation of less than 0.1 mrad.

Next, the goal was to validate the performance of the radial motion. To do this, 'star' type maps were analysed, consisting of high resolution scan over the radial axis, with measurements taken every 0.5mm along the measurement range. This was repeated at 45° steps in ϕ , at z = 0mm. The raw data from one such measurement is displayed in Figure 5.6, with the results of the fit of α and β to this data shown in Figure 5.7. 'Oscillations' of approx 0.5 mrad in size can be observed in the roll measurement, and of 0.2 mrad in size in the tilt measurement. The believed cause of these is some imperfection in the construction of the ball bearings. This would correspond to a deviation from roundness or radial run-out of around 50 microns. As can be seen, the exact form is not reproducible, though it does seem to repeat with a period of around 60mm, corresponding to the circumference of the bearing. Additionally, there was no sign of such an 'oscillation' when the profile of the rails was measured in subsection 5.7.1. The result of this will be a fundamental limitation on the achievable measurement accuracy of the transverse field components in the case where a large 1μ T vertical field to around 200pT in the radial direction and 500pT in the tangential direction is mapped.

In the analysis of the inclinometer maps of previous mapping campaigns, it was noted that the extracted forms of $\alpha(R)$ and $\beta(R)$ from the base and tube inclinometers were incompatible by 2-3 mrad (see Figure 5.8). This was determined to be a result of the mounting of the cart: it was suspended between both the upper and lower parts of the railings, with four pairs of wheels in total used. This meant that the cart was overconstrained, and as a result of the deformations in the railings was itself forced to bend. For the 2017 measurement campaign this was mitigated by removing the upper set of wheels. The results of this are shown in Figure 5.9.



Figure 5.5: Projection of global tilt on xz and yz axes, presented as a function of height z, as measured at the start of the 2017 mapping campaign. The lack of z-dependence of the global tilt indicates that the vertical axle of the mapper is properly aligned through the apparatus.


Figure 5.6: Raw tilt and roll recorded during a 'finestar' map taken at the start of the 2017 mapping campaign. The apparant 'oscillations' arise from imperfections in the bearings which function as wheels for the mapper cart. Their apparant amplitudes of 0.5 mrad in the roll axis and 0.2 mrad in the tilt axis correspond to an imperfection (deviation from roundness of the outer housing, or radial run-out) of order 50 microns.



Figure 5.7: Computed α and β in 2017, from a fit to each ring of Figure 5.6.



Figure 5.8: Difference between α and β as measured by inclinometers mounted in the base and tube of the mapper cart. A constant offset is expected as no special effort was made to calibrate the absolute readings. In the 2014 mapping campaign the mapper cart was overconstrained between two distorted rails, leading to significant bending and twisting of the card varying by up to 2-3 mrad across the radial motion range as shown above.



Figure 5.9: Difference between α and β in 2017 as measured by inclinometers mounted in the base and tube of the mapper cart. A constant offset is expected as no special effort was made to calibrate the absolute readings. The bend of the cart is found to be constant to better than 0.1 mrad, while the twist shows some structure of approx 0.5 mrad in size.

5.7.3 In Situ Determination of Fluxgate Nonorthogonalities and Offsets

To allow for the calibration of the fluxgate sensor offsets, a calibration proceedure ('calibration map') was devised. This consists of:

- One ring at R = 0 mm (Z sensor at centre of tank), $\delta \phi = 10^{\circ}$, normal fluxgate orientation
- One ring at R = 20 mm (Y sensor at centre), $\delta \phi = 10^{\circ}$, normal fluxgate orientation
- One ring at R = 40 mm (X sensor at centre), $\delta \phi = 10^{\circ}$, normal fluxgate orientation
- One ring at $R = 0 \text{ mm}, \delta \phi = 10^{\circ}$, inverted fluxgate orientation
- One ring at $R = 20 \text{ mm}, \delta \phi = 10^{\circ}$, inverted fluxgate orientation
- One ring at $R = 40 \text{ mm}, \delta \phi = 10^{\circ}$, inverted fluxgate orientation

It is possible to parameterise the field at the point at the centre of the tank as

$$\vec{B}_{\text{true}} = \begin{pmatrix} B_r \\ B_\phi \\ B_z \end{pmatrix} = \begin{pmatrix} B_x \cos \phi + B_y \sin \phi \\ B_y \cos \phi - B_x \sin \phi \\ B_z \end{pmatrix}.$$
 (5.16)

For the normal fluxgate orientation, the reading of each sensor at the centre of the tank can be written as the sum of a sine and a cosine of the ϕ angle and a constant component, their amplitudes depending on the external field and the orientation of each sensor,

$$\vec{B}_{\rm FG,norm} = \begin{pmatrix} A_{\rm 1cosx} \\ A_{\rm 1cosy} \\ A_{\rm 1cosz} \end{pmatrix} \cos \phi + \begin{pmatrix} A_{\rm 1sinx} \\ A_{\rm 1siny} \\ A_{\rm 1sinz} \end{pmatrix} \sin \phi + \begin{pmatrix} A_{\rm 1cx} \\ A_{\rm 1cy} \\ A_{\rm 1cz} \end{pmatrix}, \qquad (5.17)$$

and similarly for the inverted orientation,

$$\vec{B}_{\rm FG,inv} = \begin{pmatrix} A_{\rm 2cosx} \\ A_{\rm 2cosy} \\ A_{\rm 2cosz} \end{pmatrix} \cos \phi + \begin{pmatrix} A_{\rm 2sinx} \\ A_{\rm 2siny} \\ A_{\rm 2sinz} \end{pmatrix} \sin \phi + \begin{pmatrix} A_{\rm 2cx} \\ A_{\rm 2cy} \\ A_{\rm 2cz} \end{pmatrix}.$$
 (5.18)

Each of the coefficients A is extracted from the data by a least-squares fit.

Fixing $\alpha = 0$ and $\gamma = 0$ at R = 0 (which is possible, as these angles can be absorbed into the ψ fluxgate tube angle and a rotation of the *B* field about \hat{z} respectively), Equation 5.12 can be simplified to express the heading of each sensor as

$$\hat{X}_{global} = R_{\phi}(\beta)^{-1} \begin{pmatrix} \cos \phi_X \sin \theta_X \\ \sin \phi_X \sin \theta_X \cos \psi - \cos \theta_X \sin \psi \\ \cos \theta_X \cos \psi + \sin \phi_X \sin \theta_X \sin \psi \end{pmatrix} \\
= \begin{pmatrix} \cos \beta & 0 & -\sin \beta \\ 0 & 1 & 0 \\ \sin \beta & 0 & \cos \beta \end{pmatrix} \begin{pmatrix} \cos \phi_X \sin \theta_X \\ \sin \phi_X \sin \theta_X \cos \psi - \cos \theta_X \sin \psi \\ \cos \theta_X \cos \psi + \sin \phi_X \sin \theta_X \sin \psi \end{pmatrix} (5.19) \\
= \begin{pmatrix} \cos \beta \cos \phi_X \sin \theta_X - \sin \beta \cos \theta_X \cos \psi - \sin \beta \sin \phi_X \sin \theta_X \sin \psi \\ \sin \phi_X \sin \theta_X \cos \psi - \cos \theta_X \sin \psi \\ \sin \beta \cos \phi_X \sin \theta_X + \cos \beta \cos \theta_X \cos \psi + \cos \beta \sin \phi_X \sin \theta_X \sin \psi \end{pmatrix} \\
= \begin{pmatrix} \hat{X}_r \\ \hat{X}_{\phi} \\ \hat{X}_z \end{pmatrix}.$$

Combining Equation 5.13 and Equation 5.16, for the X fluxgate sensor in the normal orientation, one can write

$$B_{X,\text{meas}} = \hat{X} \cdot \vec{B}_{\text{true}} + \delta B_X$$

$$= \begin{pmatrix} \hat{X}_r \\ \hat{X}_{\phi} \\ \hat{X}_z \end{pmatrix} \cdot \begin{pmatrix} B_x \cos \phi + B_y \sin \phi \\ B_y \cos \phi - B_x \sin \phi \\ B_z \end{pmatrix} + \delta B_X , \quad (5.20)$$

$$= \left(\hat{X}_r B_x + \hat{X}_{\phi} B_Y \right) \cos \phi + \left(\hat{X}_r B_y - \hat{X}_{\phi} B_y \right) \sin \phi + \hat{X}_z B_z + \delta B_X$$

$$= A_{1\text{cosx}} \cos \phi + A_{1\text{sinx}} \sin \phi + A_{1\text{cx}}$$

thus relating the Fourier coefficients extracted for one sensor in one orientation to the selection of angles. Equivalent expressions can be written for the X, Y and Z fluxgate sensors in the normal and inverted orientations. Thus, from each calibration map, 12 coefficients are extracted for the next step. The results of a typical measurement are shown in Figure 5.10.

In order to gain sensitivity to the different angles, a combined fit of several calibration maps taken in different field configurations is performed. It is known that the fluxgate sensor offsets can drift substantially over time, so these are fitted individually for each map in addition to the field. Several fits were performed of various combinations of maps in different field configurations, of approximately $1 \,\mu\text{T}$ in magnitude directed vertically up and down and in a transverse direction. When fitting, different combinations of parameters were fixed, allowed to vary per-map or allowed to vary globally in different situations:

Calibration Map 012005X07



Figure 5.10: Typical first-stage fit result for a calibration map taken in the remnant field.

- α , γ and ψ_{norm} fixed to zero. All other angles allowed to vary globally. Offsets and magnetic field at centre allowed to vary for each map. α , γ and ψ_{norm} can be all absorbed entirely into other angles. This configuration allows the determination of all other angles and offsets from a set of calibration maps taken in vertical and horizontal fields.
- All angles except ψ_{norm} and ψ_{inv} fixed from previous results. Offsets, ψ_{norm} and ψ_{inv} allowed to vary for each map. This configuration allows the stability of ψ_{norm} and ψ_{inv} to be evaluated over time, however with the disadvantage that these angles are hard to distinguish from the fluxgate offset δB_X in the case of a calibration map taken in a vertical field.
- Only offsets allowed to vary, all angles fixed from previous results: this configuration provides a precise determination of the three offsets from any single calibration map, assuming good values for each of the angles can be provided.

As a strategy, all of the angles are characterised once at the beginning of each datataking period after the apparatus is closed. This is because they can typically be altered by any mechanical disturbances or manipulations. The results from each measurement set used are shown in Table 5.1. During datataking during 2017, a calibration map was taken

Angle	Nominal	2013	2014	2017-1	2017-2
ϕ_X	$-\pi/2$	-1.58683(8)	-1.57254(8)	-1.57311 (11)	-1.57306(3)
θ_X	$\pi/2$	1.56583(10)	1.58058(10)	1.57214(13)	1.57211 (4)
ϕ_Y	0	-0.75238 (422)	-1.14278 (826)	-0.15549(3035)	-0.18174 (856)
θ_Y	π	3.16736(10)	3.15294(11)	$3.14695\ (11)$	3.14735(3)
ϕ_Z	0	0.00741(8)	-0.00147 (8)	-0.00194 (11)	-0.00197 (3)
θ_Z	$\pi/2$	1.56164(5)	1.57168(5)	1.57223(7)	1.57264(2)
α_0	0	0.00000 (fixed)	0.00000 (fixed)	0.00000 (fixed)	0.00000 (fixed)
β_0	0	-0.00570 (8)	-0.00920 (8)	-0.00427(11)	-0.00399(3)
γ_0	0	0.00000 (fixed)	0.00000 (fixed)	0.00000 (fixed)	0.00000 (fixed)
$\Psi_{ m norm}$	0	0.00000 (fixed)	0.00000 (fixed)	0.00000 (fixed)	0.00000 (fixed)
$\Psi_{ m inv}$	π	3.13936(17)	3.12946(17)	3.13340(23)	3.13811 (7)

Table 5.1: Results of fit to sets of calibration maps in different fields to determine angles described in text.

before and after each map taken. This allows precise determination of the sensor offsets at several-hour intervals. A linear drift in the offset is fitted, and the computed offset values are subtracted from each measurement taken as the very first step in the analysis.

5.7.4 In Situ Mapper Misalignments Measurement

The above fit can characterise the mapper correctly only when the cart is at R = 0. To extend the validity to cover the entire mapped region, it is necessary to find the parameters $\alpha(R)$, $\beta(R)$ and $\gamma(R)$ throughout the entire mapped region. This is possible to do by making a combined fit of both the magnetic field within the volume and the misalignment angles themselves.

First, the magnetic field within the volume must be parametrised. The approach used is described fully in [141], but briefly summarised here for completeness. In a region free of currents or magnetisation, in the magnetostatic regime, the \vec{B} field is divergenceless and curlless

$$\vec{\nabla} \cdot \vec{B} = 0 \qquad \qquad \vec{\nabla} \times \vec{B} = 0. \tag{5.21}$$

As such, the magnetic field can be represented as the gradient of a scalar potential Σ satisfying the Laplace equation,

$$\vec{B} = \vec{\nabla}\Sigma \qquad \qquad \nabla^2\Sigma = 0. \tag{5.22}$$

Solutions of the Laplace equation are harmonic functions. One possible basis is the spherical harmonics Y_l^m , which form a complete basis of homogeneous polynomials with l the degree of the polynomial and m being an integer from -l to l. It is therefore possible to express any possible field within the volume in terms of the gradients of all spherical harmonics up to a given order. The choice of when to truncate the series is arbitrary, as the field produced at the edge of the vacuum tank quickly becomes unphysically large for high order components. One particularly useful convenient choice, which is used for this analysis, is that presented in [91]. This consists of a complete basis of harmonic polynomials which can be conveniently expressed in both Cartesian and cylindrical coordinate systems, which automatically satisfy the Maxwell equations.

The values of $\alpha(R)$, $\beta(R)$ and $\gamma(R)$ are characterised simply by independent free parameters for each measured radius R. They are considered to be independent of the height z, in contrast to the approach of [178], as no feasible physical mechanism exists which can produce non-independent contributions. In the case of a 'banana' vertical axis, it is possible to have a separate contribution like $\alpha(R, z) = \alpha_1(R) + \alpha_2(z)$, however this contribution has been constrained to be at most 0.1 mrad by inclinometer measurements at the beginning and end of datataking.

Similarly to above, the field in each fluxgate channel is considered to be the scalar product of its heading in the global cylindrical coordinate system, and the actual magnetic field at that position, plus some constant offset δB

$$B_{X,\text{meas}} = \hat{X}(R, z) \cdot \vec{B}_{\text{true}}(R, \phi, z)$$

$$= \begin{pmatrix} \hat{X}_r(R, z) \\ \hat{X}_{\phi}(R, z) \\ \hat{X}_z(R, z) \end{pmatrix} \cdot \begin{pmatrix} B_r(R, \phi, z) \\ B_{\phi}(R, \phi, z) \\ B_z(R, \phi, z) \end{pmatrix} + \delta \vec{B}.$$
(5.23)

Each field component $G_{l,m}$ with $m \neq 0$ can only produce fields which are proportional to $\sin(m\phi)$ or $\cos(m\phi)$. The measured data therefore can be transformed to separate the effects of each component. A linear least squares fit of

$$\vec{B}_{FG}(r,\phi,z) = \vec{A}_0(r,z) + \sum_{m=1}^N \vec{A}_{m,\cos}(r,z) \cos m\phi + \vec{A}_{m,\sin}(r,z) \sin m\phi, \qquad (5.24)$$

is performed to determine the factors $\vec{A_0}$, $\vec{A_{m,cos}}$ and $\vec{A_{m,sin}}$, applied at each ring of data taken at constant R and z. This first step is described in more detail in subsection 5.8.1. Only the components $\vec{A_0}$ are considered, as these will be the largest and most visible in the case of a 1µT vertical field. As all of the polynomials $\Pi_{\phi,l,0}$ are independent of ϕ , and all polynomials $\Pi_{\phi,l,0}$ are zero,

$$A_{X,0}(R,z) = \begin{pmatrix} \hat{X}_r(R,z) \\ \hat{X}_{\phi}(R,z) \\ \hat{X}_z(R,z) \end{pmatrix} \cdot \Sigma_{l=0}^{l_{\max}} G_{l,0} \begin{pmatrix} \Pi_{r,l,0}(R,z) \\ 0 \\ \Pi_{z,l,0}(R,z) \end{pmatrix} + \delta B_X,$$
(5.25)

and similarly for the Y and Z sensors.

In the case of a vertical 1μ T field, the leading components contributing to A_0 are as follows:

- $A_{0,X}$: Roll misalignment α (around \hat{r}), bringing a component of the vertical field $G_{0,0}$ into the X sensor
- $A_{0,Y}$: Vertical magnetic field components $\sum_{l=0}^{l_{\max}} G_{z,l,0} \prod_{l,0}$, plus a small contribution from tilt misalignment β bringing a component of the radial field into the Y sensor.
- $A_{0,Z}$: Radial magnetic field component $\sum_{l=1}^{l_{\max}} G_{l,0} \Pi_{r,l,0}$, and tilt misalignment β (around $\hat{\phi}$), bringing a component of the vertical field $G_{0,0}$ into the Z sensor.

The next step is to substitute the fluxgate headings computed from the model Equation 5.12 into Equation 5.25 for each sensor X, Y and Z. The values of ψ , $\phi_{X,Y,Z}$ and $\theta_{X,Y,Z}$ (as described in section 5.6) are fixed to the values obtained by our in situ calibration map measurement described in subsection 5.7.3. This gives an expression for the measured coefficients $A_{0,X}, A_{0,Y}$ and $A_{0,Z}$ in terms of the magnetic field components $G_{l,0}$ and the misalignments α and β at each radius. A nonlinear least squares fit is performed using the least_squares routine provided by the SciPy library [168], an implementation of the Levenberg-Marquardt algorithm.

For a normal map, when terminating the field parametrisation at $l_{\text{max}} = 6$, there are then 7 parameters $G_{l,0}$ to describe the magnetic field and 34 parameters each to describe α and β at each radius. There are 170 rings measured with each sensor in a normal map, giving 510 total measurements(170 × 3). This gives 435 degrees of freedom.

5.8 Fit of Field Coefficients

The goal of this stage of the analysis is to extract all field gradients $G_{l,m}$ from a single measured map. In order to do this, the structure of the Standardised Adequate Basis (see [141]) is once again exploited to reduce the complexity. As previously mentioned, field components $G_{l,m}$ with $m \neq 0$ can only produce fields depending on $\sin(|m|\phi)$ and $\cos(|m|\phi)$, while components m = 0 can only produce fields with no ϕ dependence. Therefore, a two-stage fit is once again used: first fitting a series of sines and cosines separately for each sensor to each ring measured by the mapper, then fitting the obtained coefficients to the gradients $G_{l,m}$.

5.8.1 Decomposition into Rings

The first stage of the fit is to determine the coefficients $A_{0,X}(r,z)$, $A_{m \cos X}(r,z)$ and $A_{m \sin X}(r,z)$ for each ring of the map. The equation

$$B_{X,Y,Z}(r,\phi,z) = A_{0,X}(r,z) + \sum_{m=1}^{N} A_{m\cos X}(r,z) \cos m\phi + A_{m\sin X}(r,z) \sin m\phi \quad (5.26)$$

is fitted by linear least squares to each ring of constant R, z. These coeffecients are then identified with the coefficients of $\sin(m\phi)$ and $\cos(m\phi)$ for each polynomial $\Pi_{l,m}$ in the Standardised Adequate Basis ([141]) as expressed in cylindrical coordinates. Results of a typical fit to a single ring towards the edge of the chamber, where non-uniformities are largest, are shown in Figure 5.11, comparing results when the expansion is terminated at m = 2 and m = 6.



Figure 5.11: Example ring measurement and fit recorded towards the edge of mapped volume, where gradients are typically larger, including plots of the fitted field, truncated at order 2 or order 6.

5.8.2 Choice of Fit Order and Goodness of Fit

The noise in the fluxgate can be charactarised by the Allan standard deviation [181, 182]. This is a tool originally used to quantify the stability of atomic clocks over different timescales, however it can also be used to charactarise the stability of other measurements as a function of the integration time. It is defined, as a function of the integration time τ , as

$$\sigma_A(\tau) = \sqrt{\frac{1}{2(N-1)} \sum_{j=1}^{N-1} \left(f_{j+1}^{(\tau)} - f_j^{(\tau)} \right)^2}$$
(5.27)

and the related Allan variance is defined as

$$\sigma_A^2(\tau) = \frac{1}{2(N-1)} \sum_{j=1}^{N-1} \left(f_{j+1}^{(\tau)} - f_j^{(\tau)} \right)^2$$
(5.28)

Where $f_j^{(\tau)}$ is the *j*'th 'chunk' of the quantity of interest, averaged from $t = (j-1)\tau$ to $t = j\tau$. The variance corresponds to 1/2 of the mean square difference between two consecutive measurements.

During the characterisation phases of datataking, on two occasions a long measurement was taken lasting several hours with the fluxgate held still at the centre of the tank, with the field nominally constant. The Allan deviation was computed for these datasets to characterise the noise of the fluxgate. The results of a typical measurement are shown in Figure 5.12.

One ring taken in the mapper consisted of 38 measurement points, each taken for 0.5-3 seconds once the position was stabilised, and took approximately 180s to acquire. While several formulae exist (namely, the 'Bias Functions' B1 and B2 [183]) which can be used to predict the 'M-Sample Deviation' in a series of measurements - which would give a rigorous prediction of the noise on the measurement - these are complex and rely on knowledge of the spectral form of the noise. Instead, use will be made of the fact that the Allan Deviation is very flat at around 20pT for all three sensors for integration times between 0.1s and hundreds of seconds - meaning that one can intuitively fix an estimated error on all measurements of 20pT, with no excess correlation between consecutive measurements.

Comparison with the fit results obtained from the rings fit in a real map displayed in Figure 5.13 will reveal if this is accurate, or if assumptions start to break down. The RMS error considering all rings within a map displayed as a function of the order of field fitted is displayed in Figure 5.14. For the order 6 fit of the Y sensor data, the RMS error is 19.1pT. Given (in the case of a fit up to m=7) 15 parameters and 38 datapoints per sensor per ring (for 23 degrees of freedom), this is equivalent to measuring $\chi^2/\nu = 34.7/23$, well within

the acceptable range considering the rough nature of the estimation of σ . Alternatively, by fixing $\chi^2/\nu = 1$, this implies an error bar of 26.5pT, consistent with expectations.



Figure 5.12: Typical Allan standard deviation measured during data taking, with the fluxgate magnetometer stationary at the centre of the shield.



Figure 5.13: Colourmap of the RMS error for each ring fit as a function of position, computed for fits to up to orders 2,4 and 6. Higher order terms grow in size very strongly with radius and height, corresponding to a poorer goodness of fit at larger radius when truncating the fit too soon.



Figure 5.14: RMS fit error from a ring decomposition fit of all points within all rings of one B_0 map. For comparison, the measured Allan standard deviation at this timescale is approximately 20pT.

103

5.8.3 Fit of Magnetic Field Gradients to Ring Coefficients: An Oversimplified First Exposition

As an example, the fit of the gradients with $m = \pm 1$ up to $l_{\text{max}} = 1$ will be considered. The relevant polynomials are

$$\begin{split} \Pi_{\rho,0,-1} &= \sin \phi & \Pi_{\phi,0,-1} &= \cos \phi & \Pi_{z,0,-1} &= 0 \\ \Pi_{\rho,0,1} &= \cos \phi & \Pi_{\phi,0,1} &= -\sin \phi & \Pi_{z,0,1} &= 0 \\ \Pi_{\rho,1,-1} &= z \sin \phi & \Pi_{\phi,1,-1} &= z \cos \phi & \Pi_{z,1,-1} &= \sin \phi \\ \Pi_{\rho,1,1} &= z \cos \phi & \Pi_{\phi,1,1} &= -z \sin \phi & \Pi_{z,1,1} &= \cos \phi. \end{split}$$

The field due to these four components is therefore

$$\vec{B} = \Sigma_{l,m} G_{l,m} \begin{pmatrix} \Pi_{\rho,l,m} \\ \Pi_{\phi,l,m} \\ \Pi_{z,l,m} \end{pmatrix}.$$

For the purpose of the demonstration, it will be assumed that each sensor is exactly aligned with the relevant axis: $\hat{X} = -\hat{\phi}$, $\hat{Y} = -\hat{z}$ and $\hat{Z} = \hat{r}$, so that $B_X = -B_{\phi}$, $B_Y = -B_z$ and $B_Z = B_r$. This assumption will be relaxed later. Substituting the polynomials $\vec{\Pi}_{l,m}$ and selecting the components visible in each sensor,

$$B_X = (-G_{0,-1} - zG_{1,-1})\cos\phi + (G_{0,1} + zG_{1,1})\sin\phi$$
$$B_Y = -G_{1,1}\cos\phi + -G_{1,-1}\sin\phi$$
$$B_Z = (G_{0,1} + zG_{1,1})\cos\phi + (G_{0,-1} + zG_{1,-1})\sin\phi.$$

These may be identified with $\vec{A}_{1,\cos}$ and $\vec{A}_{1,\sin}$

$$\begin{aligned} A_{1,\cos,X}(R,z) &= -G_{0,-1} - zG_{1,-1} & A_{1,\sin,X}(R,z) = G_{0,1} + zG_{1,1} \\ A_{1,\cos,Y}(R,z) &= -G_{1,1} & A_{1,\sin,Y}(R,z) = -G_{1,-1} \\ A_{1,\cos,Z}(R,z) &= -G_{0,1} + zG_{1,1} & A_{1,\sin,Z}(R,z) = G_{0,-1} + zG_{1,-1}. \end{aligned}$$

This yields a set of equations which can be solved by linear least squares to yield the coefficients $G_{0,-1}$, $G_{0,1}$, $G_{1,-1}$ and $G_{1,1}$. This can trivially be extended to higher order fits.

This two-stage fit brings several advantages. First, it reduces the complexity of the fit from a single large fit of 62 parameters (for a fit of up to l = 6) to almost 20000 readings

² to a series of smaller fits. In the first step, $34 \times 5 \times 3$ fits of $2l_{\max} + 3$ parameters (15 for $l_{\max=6}$) are performed to 38 points each. In the second stage, for each m from 0 to $l_{\max} + 1$ a separate fit of the coefficients $G_{l,\pm m}$ for all permitted l from $\max(|m| - 1, 0)$ to l_{\max} to the $A_{m,\cos}$ and Am, sin from each ring is performed. Each fit takes as input 340 coefficients for each sensor X, Y, Z, for up to 16 parameters. In all stages the system is significantly over determined, with many more measurement values than parameters fitted.

The field components m = 0 behave slightly differently. These components are independent of the angle ϕ . Additionally, the B_{ϕ} component is always 0. A similar approach is used, where the polynomials $\Pi_{l,0}$ are fitted to the constant coefficients of the ring fits $A_{0,X,Y,Z}$.

As a simple example, the fit of the m = 0 components up to $l_{max} = 1$ is described. The relevant polynomials included in this fit are:

$$\begin{split} \Pi_{\rho,0,0} &= 0 & \Pi_{\phi,0,0} = 0 & \Pi_{z,0,0} = 1 \\ \Pi_{\rho,1,0} &= -\rho/2 & \Pi_{\phi,1,0} = 0 & \Pi_{z,1,0} = z. \end{split}$$

Fixing $\hat{X} = -\hat{\phi}$, $\hat{Y} = -\hat{z}$ and $\hat{Z} = \hat{r}$, so that $B_X = -B_{\phi}$, $B_Y = -B_z$ and $B_Z = B_r$, one finds

$$A_{0,X}(r,z) = 0$$

$$A_{0,Y}(r,z) = -G_{0,0} - zG_{1,0}$$

$$A_{0,Z}(r,z) = rG_{1,0}.$$

These equations can readily be fitted to the data to find $G_{0,0}$ and $G_{1,0}$. This technique can easily be extended to arbitrarily high order.

5.8.4 Compensating for Misalignments and Nonorthogonalities

In the first exposition of the fitting method, the simplifying assumptions that each of the sensors was colocated at precisely the same coordinate, and that each sensor was aligned precisely with one of the axes \hat{r} , $\hat{\phi}$, or \hat{z} were made. However, this is not precisely realised, as explained in section 5.6. The angular misalignments are measured in situ using

²A typical map consists of a series of rings taken with $\Delta \phi = 10^{\circ}$ from 0° to 360°, and a repeat of 0° for 38 points per ring. At each height 17 radii were sampled from R = 0cm to R = 34cm, with a spacing of $\Delta R = 2$ cm. Most often 5 heights were measured in full, from z = -16cm to z = +16cm.

the techniques described in subsection 5.7.3 and subsection 5.7.4. Using these results, the precise orientation of each sensor can be predicted at any given position. The consideration that each of the three mapper sensors are separated by 20mm will also be added.

The vector magnetic field at the position of each sensor when the mapper is positioned at a coordinate (R, ϕ, z) will be

$$\vec{B_X} = \Sigma_{l,m} G_{l,m} \begin{pmatrix} \Pi_{\rho,l,m} (r + \Delta r_X, \phi, z) \\ \Pi_{\phi,l,m} (r + \Delta r_X, \phi, z) \\ \Pi_{z,l,m} (r + \Delta r_X, \phi, z) \end{pmatrix},$$

and similar for the Y and Z sensors.

 Δr_X , Δr_Y and Δr_Z are the radial offsets of each of the fluxgate sensors as defined in section 5.6. These are fixed at $\Delta r_X = -40$ mm, $\Delta r_Y = 0$ mm, and $r_Z = -20$ mm. As radii between r = 0mm and r = +340mm are measured, what happens when $r - \Delta r$ is negative must be considered, as negative radii are not usually allowed in cylindrical coordinate systems. Working at a negative radius is equivalent to working at the modulus of the radius, and adding a ϕ offset of 180°. The effect of this 180° phase shift depends on the nature of the field. In the Standardized Adequate Basis, all field components $\Pi_{l,\pm m}$ are proportional to $\sin(m\phi)$ or $\cos(m\phi)$. In the case where m is even, a shift of ϕ of 180° will lead to an offset of a multiple of 360° being added, which leaves both the sine and cosine unaffected. However, when m is odd, the sign of the field will be inverted. Additionally, for all m, going to a negative radius will cause the \hat{r} and $\hat{\phi}$ components of the heading of each sensor to invert. These corrections will be applied automatically for the remainder of this derivation. The field actually measured by each sensor will be the dot product of this field with the sensor with its heading \hat{X} .

$$B_X(r,\phi,z) = \vec{B}_X(r+\Delta r_X,\phi,z) \cdot \hat{X}(r,z)$$
$$= \Sigma_{l,m}G_{l,m} \begin{pmatrix} \Pi_{\rho,l,m}(r+\Delta r_X,\phi,z) \\ \Pi_{\phi,l,m}(r+\Delta r_X,\phi,z) \\ \Pi_{z,l,m}(r+\Delta r_X,\phi,z) \end{pmatrix} \cdot \begin{pmatrix} \hat{X}_r(r,z) \\ \hat{X}_{\phi}(r,z) \\ \hat{X}_z(r,z) \end{pmatrix}$$

As a shorthand, $\vec{\Pi}_{l,m,\cos}(r,z)$ and $\vec{\Pi}_{l,m,\sin}(r,z)$ are defined to be the components of $\Pi_{l,m}(r,\phi,z)$ proportional to $\sin(|m|\phi)$ and $\cos(|m|\phi)$ respectively. Taking $\sin(m\phi)$ and $\cos(m\phi)$ outside of the summation over l,

$$B_X(r,\phi,z) = \Sigma_m \left(\Sigma_l G_{l,m} \hat{X}(r,z) \cdot \vec{\Pi}_{l,m,\cos}(r+\Delta r_X,z) \right) \cos(|m|\phi) + \Sigma_m \left(\Sigma_l G_{l,m} \hat{X}(r,z) \cdot \vec{\Pi}_{l,m,\sin}(r+\Delta r_X,z) \right) \sin(|m|\phi).$$

The parts of the expressions in brackets are equated to the coefficients $A_{m \cos X,Y,Z}(r,z)$ and $A_{m \sin X,Y,Z}(r,z)$.

$$A_{m,X,\cos}(r,z) = \sum_{l} G_{l,m} \hat{X}(r,z) \cdot \vec{\Pi}_{l,m,\cos}(r + \Delta r_X, z)$$
$$+ \sum_{l} G_{l,-m} \hat{X}(r,z) \cdot \vec{\Pi}_{l,-m,\cos}(r + \Delta r_X, z)$$
$$A_{m,X,\sin}(r,z) = \sum_{l} G_{l,m} \hat{X}(r,z) \cdot \vec{\Pi}_{l,m,\sin}(r + \Delta r_X, z)$$
$$+ \sum_{l} G_{l,-m} \hat{X}(r,z) \cdot \vec{\Pi}_{l,-m,\sin}(r + \Delta r_X, z)$$

This system of expressions can be fitted to the data to obtain the field coefficients $G_{l,m}$ by reforming this into a matrix computation, combining this and the similar expressions for the Y and Z sensors. This yields a system $\mathbf{M}x = A$, where \mathbf{M} is a large matrix containing the coefficients $\hat{X}, \hat{Y}, \hat{Z}(r_i, z_i) \cdot \Pi_{l,\pm m,\cos,\sin}(r_i, z_i)$, A is a vector containing the measured coefficients from each ring $A_{\pm m \cos X,Y,Z}$ and x a vector containing the gradients $G_{l_{\min},\pm m}$ to $G_{l_{\max},\pm m}$, which can be fitted by linear least squares. At each $\pm m$, there are 6 coefficients per map ring measured (sine and cosine-like components from 3 sensors), giving 510 values total. To this, all components $G_{l\min,\pm m}$, to $G_{l\max,\pm m}$, where $l_{\min} = \max(0, m-1)$ are fitted: up to 14 components for $l_{\max} = 6$. This system is therefore suitably over-determined.

As was discovered in the analysis of measurements taken with the inclinometers in subsection 5.7.2, there is a measurable nonreproducibility in the precise angular alignment of the fluxgate, which is attributed to imperfections in the bearings. This manifests as a noise of around 0.5mrad in the roll misalignment α , and 0.2mrad in the tilt misalignment β , measured 'peak-to-peak'. In the case of mapping a 1 μ T vertical field with only small transverse components, this misalignment would cause additional components of up to 200pT in the radial sensor and 500pT in the transverse sensor. As the vertical sensor points approximately along the direction of the field, any misalignments will cause fake fields proportional to the cosine of a small angle, giving significantly greater immunity to these fake fields. In order to test this effect, several fits were performed excluding the sensors affected.

In the case of the m = 0 fit, only the vertical (Y) and radial (Z) sensors can meaningfully probe the field. The radial and vertical sensor are both sensitive to the field, however the radial sensor is affected by false fields arising from the misalignment of the order $B_{r,\text{false}} = B_z \sin(\beta) \approx 5 \text{ nT}$. All components $G_{l,m}$ produce components in both the radial and vertical directions, therefore it is possible to access these components using only the vertical sensor. Misalignment effects will produce false fields in the vertical sensor only

	Fit of c	Fit of only Y		Fit of Y and Z		
	Value	Uncert	Value	Uncert	Diff	Diff σ
$G_{0,0} \ \mathrm{pT}$	1.03e+06	4.02	1.03e+06	10.7	12.2	1.07
$G_{1,0} \text{ pT/cm}$	22	0.409	18.1	0.874	3.83	3.97
$G_{2,0} \mathrm{ pT/cm^2}$	-8.02	0.0235	-8.1	0.0526	0.0844	1.46
$G_{3,0} \mathrm{ pT/cm^3}$	-0.0654	0.00103	-0.0669	0.00217	0.00141	0.587
$G_{4,0} \ \mathrm{pT/cm^4}$	-0.00393	5.01e-05	-0.00397	9.66e-05	4.53e-05	0.416
$G_{5,0} \mathrm{\ pT/cm^5}$	-2.05e-05	1.27e-06	-2.03e-05	2.37e-06	-2.24e-07	-0.0833
$G_{6,0} \mathrm{ pT/cm^6}$	-1.18e-06	3.48e-08	-1.2e-06	7.5e-08	2.35e-08	0.285
X RMS Error pT	612		612			
Y RMS Error pT	46.8		60.3			
Z RMS Error pT	18	36	175			

Table 5.2: Goodness of fit parameters comparing fit of m = 0 using only vertical (Y) and both radial and vertical (Z and Y) sensors.

of the order $-B_r \sin(\beta)$ or $B_z(1 - \cos(\beta))$, which are negligibly small.

By comparison, components $m \neq 0$ produce components directly measurable in all three sensors, with the exception of the purely transverse modes such as $G_{1,\pm 2}$ and $G_{2,\pm 3}$, which produce no vertical component. As such, it is possible to access all such components using only the vertical and radial sensors.

As can be seen by the results presented in Table 5.2, comparing the fit to a typical B up map using only the vertical sensor to using both vertical and radial sensors, the numerical results for $G_{3,0}$ differ by up to 1σ . As the combined fit using both sensors is strongly dependent on the input value of $\beta(R)$, the fit using the vertical sensor is significantly more reliable. In this case, the RMS error on the vertical field is 47 pT, only a little more than the expected fluxgate noise. This is small compared to 186 pT RMS error for the predicted error on the radial sensor, which is consistant with arising from a 0.2 mrad random error in the angle β . As such, it was chosen to use only the vertical sensor in the fit of m = 0gradients. This results in 7 gradients $G_{0,0}$ to $G_{6,0}$ being fitted to 170 datapoints, for 163 degrees of freedom.

5.8.5 Assessment of fit accuracy and robustness

In order to interpret the results of each map, it is cruicial to understand the errors of each fitted coefficient and the covariances between them. These errors will arise from statistical sources such as noise or offset drift on the reading of the fluxgate, from positioning errors, or from methodological errors.

Assessment of the statistical fit error

In section 5.8, a two-stage fit was described: first a Fourier-like series was fitted to individual rings taken at constant R and z, with ϕ ranging from 0° to 360° every 10° and a repeat of the 0° position to give 38 points per ring in total. The coefficients obtained in this first fit are then used as the input for a second stage of fitting, where the gradients $G_{l,m}$ and $G_{l,-m}$ for $l = \max(m-1,0)$ to $l = l_{\max}$ are fitted to the coefficients $\cos(m\phi)$ and $\sin(m\phi)$ measured in each ring.

The most obvious source of statistical error in this procedure is noise on the individual fluxgate readings. As described in subsection 5.8.2, this can be explored through the measured Allan standard deviation of the data. It was previously found that the fit error within each ring is roughly consistent with what may be expected from the Allan deviation when fitting all orders up to $l_{\text{max}} = 6$. The same approach may be used to assess if the fit error can be explained by fluxgate noise. A single map lasted typically 4 to 5 hours, so the fit error would be expected to be of the order of the Allan standard deviation around this time. For one typical B_0 up map, the RMS error on the Y sensor was measured to be 47 pT, which is comparable to the expected error due to fluxgate stability over this integration time.

There are multiple methods to propagate this error forwards to the final gradients. The most obvious approach is to assume that the error on each \vec{A} used as input in a single fit is equal and independently distributed, and of the correct size to give $\chi^2/\nu = 1$, and then to use standard techniques to propagate this to find a final fit error [184]. However, these assumptions do not fully hold. In the event of a map with low field gradients, such as a map of B_0 , where the largest error comes from fluxgate drifts, there will be correlations between the readings which is not accounted for. Alternatively in the case of a map with large gradients, such as of a trimcoil designed to produce such gradients, the key source of error will be from positioning errors. In this case, the error will be most likely largest on the points furthest from the origin. These points have the greatest leverage on the determination of higher order gradients such as $G_{3,0}$, which causes the accuracy of these to be overestimated. This method yields an error on the final $G_{3,0}$ of around $\sigma_{G_{3,0}} = 0.0008$ pT/cm³.

As an alternative, more empirical, method, the 'bootstrap' [185, 186] can be used. In

short, this is a resampling technique, where the original dataset is reshuffled, choosing which data is included at random, with replacement. This results in some data points being double or triple weighted (or more), and others being not included at all. This new bootstrap dataset can then be fitted again to measure the parameter of interest again. By repeating this process many times, some distribution of the parameter of interest can be found. It can be shown that this technique successfully estimates the standard error of parameters of interest in many cases, without needing to make assumptions about the statistical nature of the underlying distribution or data. In this analysis, the bootstrap is applied at the intermediate step, selecting or deselecting entire rings to be shuffled into the final fit. Using this method, the estimated error on $G_{3,0}$ is around twice as large at $\sigma_{G_{3,0}} = 0.0015 \text{ pT/cm}^3$. This method also provides a natural way to estimate the error on the measured $\langle B_T^2 \rangle$, taking into account covariances between the all $G_{l,m}$ by bootstrapping. This yields an error on $\langle B_T^2 \rangle$ of $\sigma_{\langle B_T^2 \rangle} = 0.01 \text{ nT}^2$. The downside of this method is that it is computationally expensive, with the analysis needing to be repeated fifty times to find the error bars on each map.

In any case, the pure statistical error computed here is shown later in this chapter to be significantly less than other uncertainties such as the field reproducibility after degaussing and estimates of systematic errors in the mapping analysis. The uncertainties obtained using the bootstrap method are deemed to be definitive, as they avoid making the assumptions needed for the first method, and, being larger, they are also the conservative choice.

Fits with and without misalignments

It is also desirable to quantify how sensitive the results are to the accuracy of our determination of the misalignment angle functions $\alpha(R, z)$ and $\beta(R, z)$. To obtain an upper bound on the possible systematic, the results for a typical map are compared including the best determination of the misalignment angles, and when all misalignment angles are fixed to zero. The results for the m = 0 fits are presented in Table 5.3 and Table 5.4.As can be clearly seen, when only the Y sensor is included in the m = 0 fit the results are practically immune to the value of β , with the RMS fit error on the Y sensor improving from 47.2pT to 46.8pT when the misalignments are included. This might be expected, as the effect of these misalignments on the Y sensor in a vertical field acts like the cosine of a small angle. Meanwhile, the effect on the X and Z sensors is much more pronounced with the RMS error on the Z sensor reading improving from 3nT to 0.19nT, and the same

	No misalignments With misalignments					
	Value	Uncert	Value	Uncert	Diff	Diff σ
$G_{0,0} \ \mathrm{pT}$	1.03e+06	4.06	1.03e+06	4.02	10.3	1.8
$G_{1,0} \ \mathrm{pT/cm}$	21.8	0.413	22	0.409	-0.145	-0.249
$G_{2,0} \mathrm{ pT/cm^2}$	-8	0.0237	-8.02	0.0235	0.0222	0.665
$G_{3,0} \mathrm{ pT/cm^3}$	-0.0655	0.00103	-0.0654	0.00103	-0.000105	-0.0722
$G_{4,0} \ \mathrm{pT/cm^4}$	-0.00394	5.06e-05	-0.00393	5.01e-05	-9.98e-06	-0.14
$G_{5,0}~{ m pT/cm^5}$	-2.05e-05	1.28e-06	-2.05e-05	1.27e-06	9.45e-09	0.00525
$G_{6,0} \mathrm{ pT/cm^6}$	-1.19e-06	3.51e-08	-1.18e-06	3.48e-08	-1.85e-08	-0.374
X RMS Error pT	1.3e+03		612			
Y RMS Error pT	47.2		46.8			
Z RMS Error pT	3.01e	+03	18	36		

Table 5.3: Goodness of fit parameters comparing fit of m = 0 using only vertical (Y) sensor, with and without misalignment corrections

number for the X sensor improving from 1.3nT to 0.6nT. (only Y is included in the fit, but it is possible to predict values for all sensors). These numbers are consistant with the field and misalignments being correctly modelled, but with a nonreproducibility of β of 0.2mrad and a nonreproducibility of α of 0.5mrad. As the final fit values are practically unchanged by this extreme test, errors from this source will be neglected. Results for modes $m \neq 0$ are also not significantly altered.

	No misali	No misalignments		With misalignments		
	Value	Uncert	Value	Uncert	Diff	Diff σ
$G_{0,0} \ \mathrm{pT}$	1.03e+06	162	1.03e+06	10.7	-147	-0.902
$G_{1,0} \ \mathrm{pT/cm}$	-65.7	13.2	18.1	0.874	-83.9	-6.32
$G_{2,0} \mathrm{ pT/cm^2}$	-8.99	0.797	-8.1	0.0526	-0.89	-1.11
$G_{3,0} \mathrm{ pT/cm^3}$	-0.258	0.0328	-0.0669	0.00217	-0.191	-5.81
$G_{4,0} \ \mathrm{pT/cm^4}$	-0.00544	0.00146	-0.00397	9.66e-05	-0.00147	-1
$G_{5,0}~{ m pT/cm^5}$	-0.00017	3.58e-05	-2.03e-05	2.37e-06	-0.00015	-4.17
$G_{6,0} \mathrm{ pT/cm^6}$	-1.18e-06	1.14e-06	-1.2e-06	7.5e-08	2.14e-08	0.0188
X RMS Error pT	1.3e+03		612			
Y RMS Error pT	893		60.3			
Z RMS Error pT	2.66e	+03	175			

Table 5.4: Goodness of fit parameters comparing fit of m = 0 using vertical and radial sensors (Y and Z) sensor, with and without misalignment corrections

Drift Correction

Inaccuracies in the quality of the calibration measurement may also affect the result. Systematic errors in the measurement of the Y offset using the calibration procedure described in subsection 5.7.3 will manifest as a shift in $G_{0,0}$ (the accuracy of which is not particularly important), while the accuracy of the calibration of the X or Z offsets will have a similar effect to a small constant offset in the angles α or β , which have been previously established to have no significant effect on the fitted parameters.

As mapping measurements were taken in a non-randomised order, sequentially progressing from the lowest plane to the highest, it may be expected that a drift in the offset could manifest as a false gradient like $G_{1,0}$. In the mapping measurements of 2017, before and after each large map a calibration map was taken. This allows correction for a linear drift in the fluxgate offset. In Table 5.5 the fit results for m=0 are compared, compensating for a linear drift, or taking the offsets to be constant for the run. It might be expected that using multiple sensors to fit the gradients $G_{l,0}$ may give some additional immunity to drifts, as both the Y and Z sensors would be required to drift in a way so as to mimic a gradient. Comparison results for this test are shown in Table 5.6. It might also be expected that statistical errors in determining the drift during one run may bias the results of that run, these results may be interpreted as a constraint on that effect. As can be seen, the effect of correcting or not correcting for this bias in a run with typical

	No Drift C	Correction	Linear Drift Correction			
	Value	Uncert	Value	Uncert	Diff	Diff σ
$G_{0,0} \ \mathrm{pT}$	1.03e+06	4.01	1.03e+06	4.02	-1.25	-0.22
$G_{1,0} \ \mathrm{pT/cm}$	21	0.408	22	0.409	-0.997	-1.73
$G_{2,0} \mathrm{ pT/cm^2}$	-8.01	0.0234	-8.02	0.0235	0.0115	0.345
$G_{3,0} \mathrm{\ pT/cm^3}$	-0.0654	0.00102	-0.0654	0.00103	5.4e-05	0.0373
$G_{4,0} \mathrm{\ pT/cm^4}$	-0.00392	5e-05	-0.00393	5.01e-05	6.9e-06	0.0974
$G_{5,0} \mathrm{~pT/cm^5}$	-2.05e-05	1.26e-06	-2.05e-05	1.27e-06	1.47e-08	0.0082
$G_{6,0} \mathrm{~pT/cm^6}$	-1.17e-06	3.47e-08	-1.18e-06	3.48e-08	1.26e-09	0.0256
X RMS Error pT	606		612			
Y RMS Error pT	46.7		46.8			
Z RMS Error pT	19	2	186			

Table 5.5: Goodness of fit parameters comparing fit of m = 0 using only vertical Y sensor, with and without drift corrections

(small) drifts is negligible.

As each entire ring took only up to three minutes, any reasonable drift will not bias the Fourier coefficients for $m \neq 0$, so the gradients $m \neq 0$ are immune to this effect.

The results of performing the map analysis on a single representative map with or without misalignment correction, drift correction, and the data of the transverse sensor are displayed in Figure 5.15 to Figure 5.21.

	No Drift Correction		Linear Drift Correction			
	Value	Uncert	Value	Uncert	Diff	Diff σ
$G_{0,0} \text{ pT}$	1.03e+06	10.8	1.03e+06	10.7	-7.97	-0.524
$G_{1,0} \text{ pT/cm}$	17	0.882	18.1	0.874	-1.13	-0.911
$G_{2,0} \mathrm{ pT/cm^2}$	-8.2	0.0531	-8.1	0.0526	-0.0958	-1.28
$G_{3,0} \mathrm{ pT/cm^3}$	-0.0665	0.00219	-0.0669	0.00217	0.000365	0.118
$G_{4,0} \mathrm{ pT/cm^4}$	-0.00399	9.75e-05	-0.00397	9.66e-05	-1.59e-05	-0.116
$G_{5,0}~{ m pT/cm^5}$	-2.02e-05	2.39e-06	-2.03e-05	2.37e-06	6.12e-08	0.0182
$G_{6,0} \mathrm{ pT/cm^6}$	-1.22e-06	7.57e-08	-1.2e-06	7.5e-08	-2.12e-08	-0.199
X RMS Error pT	606		612			
Y RMS Error pT	67.3		60.3			
Z RMS Error pT	17	74	175			

Table 5.6: Goodness of fit parameters comparing fit of m = 0 using vertical Y and radial Z sensor, with and without drift corrections



Figure 5.15: Fitted value of terms $G_{0,m}$ for different choices of misalignment values, drift correction and sensor inclusion



Figure 5.16: Fitted value of terms $G_{1,m}$ for different choices of misalignment values, drift correction and sensor inclusion



Figure 5.17: Fitted value of terms $G_{2,m}$ for different choices of misalignment values, drift correction and sensor inclusion



Figure 5.18: Fitted value of terms $G_{3,m}$ for different choices of misalignment values, drift correction and sensor inclusion



Figure 5.19: Fitted value of terms $G_{4,m}$ for different choices of misalignment values, drift correction and sensor inclusion



Figure 5.20: Fitted value of terms $G_{5,m}$ for different choices of misalignment values, drift correction and sensor inclusion



Figure 5.21: Fitted value of terms $G_{6,m}$ for different choices of misalignment values, drift correction and sensor inclusion

Restricting Fit Volume

As a broad-spectrum test against a wide variety of possibly unknown systematic effects, the effect of restricting the range of points used in the fit was tested. As an example, this test would be sensitive to biases in the extracted gradients arising from a local dipole contamination on the mapper robot, or on the surface of the tank. The effect of restricting the analysis to points $R \leq 240$ mm,-80mm $\leq z$, $-80 \leq z \leq +80$ mm, $R \geq 240$ mm and simultaneously $R \leq 240$ mm and $-80 \leq z \leq +80$ mm were investigated. Naturally, restricting the domain of points will reduce the sensitivity of the fit, both because there are fewer points, and because points furthest from the origin have the greatest sensitivity to higher order gradients.

The results of the map analysis for the modes m = 0 restricting the domain of points by radius, height, or both radius and height are shown in Table 5.7, Table 5.8, Table 5.9 respectively. These same results are shown graphically for all gradients in Figure 5.22 to Figure 5.28.



Figure 5.22: Fitted value of terms $G_{0,m}$ for restricted domains of points

	Full Map		$R < 240 \mathrm{mm}$			
	Value	Uncert	Value	Uncert	Diff	Diff σ
$G_{0,0} \ \mathrm{pT}$	1.03e+06	4.02	1.03e+06	3.25	-1.87	-0.362
$G_{1,0} \text{ pT/cm}$	22	0.409	22.3	0.301	-0.35	-0.69
$G_{2,0} \mathrm{ pT/cm^2}$	-8.02	0.0235	-8.04	0.0228	0.0198	0.605
$G_{3,0} \mathrm{ pT/cm^3}$	-0.0654	0.00103	-0.067	0.00102	0.00157	1.09
$G_{4,0} \mathrm{ pT/cm^4}$	-0.00393	5.01e-05	-0.00405	6.1e-05	0.000121	1.54
$G_{5,0} \mathrm{ pT/cm^5}$	-2.05e-05	1.27e-06	-5.71e-06	2.9e-06	-1.48e-05	-4.68
$G_{6,0} \mathrm{ pT/cm^6}$	-1.18e-06	3.48e-08	-1.77e-06	1.17e-07	5.95e-07	4.87
X RMS Error pT	612		535			
Y RMS Error pT	46.8		32.6			
Z RMS Error pT	18	86	200			

Table 5.7: Goodness of fit parameters comparing fit of m = 0 using only the vertical (Y) sensor, using only measurement points with R < 240mm.

	Full	Full Map		-80 mm < z < 80 mm		
	Value	Uncert	Value	Uncert	Diff	Diff σ
$G_{0,0} \ \mathrm{pT}$	1.03e+06	4.02	1.03e+06	3.91	-7.82	-1.39
$G_{1,0}~{ m pT/cm}$	22	0.409	25.3	0.663	-3.3	-4.24
$G_{2,0} \mathrm{~pT/cm^2}$	-8.02	0.0235	-7.95	0.0445	-0.0716	-1.42
$G_{3,0}~{ m pT/cm^3}$	-0.0654	0.00103	-0.0705	0.00236	0.0051	1.98
$G_{4,0} \mathrm{\ pT/cm^4}$	-0.00393	5.01e-05	-0.00367	0.0001	-0.000262	-2.34
$G_{5,0}~{ m pT/cm^5}$	-2.05e-05	1.27e-06	-2.8e-05	2.05e-06	7.54e-06	3.13
$G_{6,0} \mathrm{ pT/cm^6}$	-1.18e-06	3.48e-08	-8.37e-07	6.76e-08	-3.39e-07	-4.46
X RMS Error pT	61	2	56	30		
Y RMS Error pT	46.8		25.3			
Z RMS Error pT	18	6	18	38		

Table 5.8: Goodness of fit parameters comparing fit of m = 0 using only the vertical (Y) sensor, using only height -80 < 0 < 80mm.

	Full Map		R < 240mm, $ z < 80$ mm			
	Value	Uncert	Value	Uncert	Diff	Diff σ
$G_{0,0} \ \mathrm{pT}$	1.03e+06	4.02	1.03e+06	4.24	-8.03	-1.37
$G_{1,0} \ \mathrm{pT/cm}$	22	0.409	25.3	0.773	-3.35	-3.83
$G_{2,0} \mathrm{ pT/cm^2}$	-8.02	0.0235	-8.01	0.0725	-0.0153	-0.2
$G_{3,0} \mathrm{ pT/cm^3}$	-0.0654	0.00103	-0.0656	0.00534	0.000113	0.0207
$G_{4,0} \ \mathrm{pT/cm^4}$	-0.00393	5.01e-05	-0.00403	0.000312	0.000101	0.32
$G_{5,0}~{ m pT/cm^5}$	-2.05e-05	1.27e-06	-1.43e-05	1.02e-05	-6.16e-06	-0.6
$G_{6,0}~{ m pT/cm^6}$	-1.18e-06	3.48e-08	-1.49e-06	3.54e-07	3.13e-07	0.879
X RMS Error pT	612		539			
Y RMS Error pT	46.8		26.5			
Z RMS Error pT	18	6	167			

Table 5.9: Goodness of fit parameters comparing fit of m = 0 using only the vertical (Y) sensor, using only measurement points with R < 240mm and height -80 < 0 < 80mm.



Figure 5.23: Fitted value of terms $G_{1,m}$ for restricted domains of points



Figure 5.24: Fitted value of terms $G_{2,m}$ for restricted domains of points



Figure 5.25: Fitted value of terms $G_{3,m}$ for restricted domains of points



Figure 5.26: Fitted value of terms $G_{4,m}$ for restricted domains of points



Figure 5.27: Fitted value of terms $G_{5,m}$ for restricted domains of points



Figure 5.28: Fitted value of terms $G_{6,m}$ for restricted domains of points
5.8.6 Effect of Position Offsets

Although the greatest care was taken to properly align the mapper at the start of datataking each year, it is possible that there is some systematic offset in the precise position measured in the radial and vertical axes. The vertical axis was positioned relative to a large vacuum flange at the bottom-centre of the tank, and then aligned vertically relative to gravity as described in subsection 5.7.2. The radial position was set by manually positioning the mapper cart completely to the end stop of the radial rail, which was fixed to be +400mm, according to the design.

The absolute effect was quantified by adding offsets of up to ± 5 mm to the true recorded values of the radius and z coordinate, repeating the analysis and looking at the change in the result. The results of this analysis are shown in Figure 5.29, Figure 5.30 and Figure 5.31.



Figure 5.29: Fitted value of $G_{3,0}$ from map 12112X15 and RMS fit error for small offsets in radius.

As can be seen, a shift of over 1 mm in either direction is required to cause a 1σ change in the values of the relevant parameters, and a shift of over 3 mm is needed before this uncertainty begins to infringe upon the accuracy requirements of the measurement. However, considering all steps in the installation of the mapper and in the installation of the precession chamber of the apparatus before nEDM datataking, we are confident that the positioning uncertainty is better than 0.5 mm.

Although an offset in the ϕ axis would cause some issues with the raw values of the



Figure 5.30: Fitted value of $\langle B_T^2 \rangle$ from map 12112X15 and RMS fit error for small offsets in radial zero position.

gradients, causing some shift between the $G_{l,m}$ and $G_{l,-m}$ modes, the precession chamber was cylindrically symmetric (with the exception of the mercury shutter and windows), so any shift will cancel out in the relevant results ($\langle B_T \rangle^2$ and $G_{l,0}$).



Figure 5.31: Fitted value of $G_{3,0}$ and $\langle B_T^2 \rangle$ from map 12112X15 for small offsets in the z zero position.

5.9 Reproducibility of Magnetic Field

The overarching aim of the mapping measurement is to estimate the magnetic field configuration during datataking from a series of offline measurements. Mapping measurement campaigns took place in 2013, 2014 and 2017, while nEDM datataking took place in August-December 2015 and May-December 2016. During the intervening time periods, the μ -metal shields and vacuum tank were opened and closed many times. Additionally, the μ -metal shields were degaussed several times during datataking when changing the magnetic configuration and inverting the direction of the B field - this process is not exactly reproducible. It is therefore necessary to quantify how reproducible the magnetic field is between measurements. A specific aim of the 2017 campaign was to measure this effect. As described in section 5.4, the majority of datataking consisted of a 24-hour cycle of a B field inversion degaussing followed by 3-4 maps, including at least one B_0 map. With such a strategy, 17 maps of a pure B up field and 11 maps of a pure B down field were obtained after repeated degaussing. Smaller numbers of similar maps were taken in 2013 and 2014. The fitted gradients to B up and B down maps taken in 2017 are shown in Figure 5.32 to Figure 5.47. By comparing the measured field in each map, the reproducibility of degaussing can be estimated. During the 2017 mapping measurement campaign, the shield and tank were opened once and several maps taken before and after. By comparing these sets of maps, the reproducibility of the field arising between opening the tank and manipulating the apparatus can be estimated. As a final check, the magnetic field configuration can be compared between years. Comparison of the measurements of 2017 to 2013 and 2014 are shown in Figure 5.48 to Figure 5.63. During the mapping campaign of 2014, a magnetic contamination was detected on the vacuum tank within a blind threaded hole used to attach the lid to a trolley for removal and installation. This was removed partway through the 2014 mapping campaign. Maps taken before the removal of the contamination are marked '2014 Bad Helicoils', and those from after are marked simply '2014'. Additionally, between the 2014 mapping campaign and the start of EDM datataking, the configuration of the μ -metal shield was altered to facilitate developments in the mercury magnetometry system: a tube of μ -metal which previously reached from the central shield outwards to protect the readout PMT was removed. The degaussing proceedure was also slightly altered. From 2015, the guiding field coils which lead along the neutron guides to the edge of the vacuum tank to maintain the neutron polarisation during transport were additionally turned on. Finally, during the 2014 mapping campaign the large 5T superconducting magnet was shut down for the mapping period, while in 2017 and during nEDM datataking, the magnet was switched on. An attempt was made to replicate the field configuration of 2014 as part of the 2017 mapping campaign in order to analyse the success of this, and ensure that any difference in field can be explained.



Figure 5.32: Fitted gradients l = 0 for B up, considering only maps taken in 2017.



Figure 5.33: Fitted gradients l = 1 for B up, considering only maps taken in 2017.



Figure 5.34: Fitted gradients l = 2 for B up, considering only maps taken in 2017.



Figure 5.35: Fitted gradients l = 3 for B up, considering only maps taken in 2017.



Figure 5.36: Fitted gradients l = 4 for B up, considering only maps taken in 2017.



Figure 5.37: Fitted gradients l = 5 for B up, considering only maps taken in 2017.



Figure 5.38: Fitted gradients l = 6 for B up, considering only maps taken in 2017.



Figure 5.39: $\langle B_T^2 \rangle$ computed for each B up, considering only maps taken in 2017.



Figure 5.40: Fitted gradients l = 0 for B down, considering only maps taken in 2017.



Figure 5.41: Fitted gradients l = 1 for B down, considering only maps taken in 2017.



Figure 5.42: Fitted gradients l = 2 for B down, considering only maps taken in 2017.



Figure 5.43: Fitted gradients l = 3 for B down, considering only maps taken in 2017.



Figure 5.44: Fitted gradients l = 4 for B down, considering only maps taken in 2017.



Figure 5.45: Fitted gradients l = 5 for B down, considering only maps taken in 2017.



Figure 5.46: Fitted gradients l = 6 for B down, considering only maps taken in 2017.



Figure 5.47: $< B_T^2 >$ computed for each B down map, considering only maps taken in 2017.



Figure 5.48: Fitted gradients l = 0 for B up, comparing maps taken in 2013, 2014 and 2017, including the 2017 maps taken in the 2014 setup.



Figure 5.49: Fitted gradients l = 1 for B up, comparing maps taken in 2013, 2014 and 2017, including the 2017 maps taken in the 2014 setup.



Figure 5.50: Fitted gradients l = 2 for B up, comparing maps taken in 2013, 2014 and 2017, including the 2017 maps taken in the 2014 setup.



Figure 5.51: Fitted gradients l = 3 for B up, comparing maps taken in 2013, 2014 and 2017, including the 2017 maps taken in the 2014 setup.



Figure 5.52: Fitted gradients l = 4 for B up, comparing maps taken in 2013, 2014 and 2017, including the 2017 maps taken in the 2014 setup.



Figure 5.53: Fitted gradients l = 5 for B up, comparing maps taken in 2013, 2014 and 2017, including the 2017 maps taken in the 2014 setup.



Figure 5.54: Fitted gradients l = 6 for B up, comparing maps taken in 2013, 2014 and 2017, including the 2017 maps taken in the 2014 setup.



Figure 5.55: $\langle B_T^2 \rangle$ computed for each B up map, comparing maps taken in 2013, 2014 and 2017, including the 2017 maps taken in the 2014 setup.



Figure 5.56: Fitted gradients l = 0 for B down, comparing maps taken in 2013, 2014 and 2017, including the 2017 maps taken in the 2014 setup.



Figure 5.57: Fitted gradients l = 1 for B down, comparing maps taken in 2013, 2014 and 2017, including the 2017 maps taken in the 2014 setup.



Figure 5.58: Fitted gradients l = 2 for B down, comparing maps taken in 2013, 2014 and 2017, including the 2017 maps taken in the 2014 setup.



Figure 5.59: Fitted gradients l = 3 for B down, comparing maps taken in 2013, 2014 and 2017, including the 2017 maps taken in the 2014 setup.



Figure 5.60: Fitted gradients l = 4 for B down, comparing maps taken in 2013, 2014 and 2017, including the 2017 maps taken in the 2014 setup.



Figure 5.61: Fitted gradients l = 5 for B down, comparing maps taken in 2013, 2014 and 2017, including the 2017 maps taken in the 2014 setup.



Figure 5.62: Fitted gradients l = 6 for B down, comparing maps taken in 2013, 2014 and 2017, including the 2017 maps taken in the 2014 setup.



Figure 5.63: $\langle B_T^2 \rangle$ computed for each B down map, comparing maps taken in 2013, 2014 and 2017, including the 2017 maps taken in the 2014 setup.

5.10 Fit of Trimcoil Fields

In normal operation, the magnetic field produced by the main B_0 coil is modified by a set of N trim coils and M guiding field coils. The trim coils are used to apply gradients in order to homogenise the field inside the precession chamber to maintain excellent neutron polarisation times (thus having a direct impact on statistical sensitivity), and to control magnetic field related systematic effects. Guiding field coils are used to maintain high neutron polarisation during transport from the 5T superconducting magnet polariser to the chamber, and from the chamber to the USSA spin analysis system (see chapter 3). Most of the guiding field coils are outside the μ -metal shielding and therefore will not cause large gradients within the tank, however the vertical guiding field coil (which leads from the switch to the bottom of the vacuum tank) penetrates the μ -metal shield and is known to cause large gradients inside the tank. The current used in each guiding field coil was optimised for each direction of B_0 at the beginning of neutron datataking each year, while the trimcoil currents are optimised individually after each degaussing using a technique described in [127]. Additionally, the coils TTC and BTC are used together to change the vertical gradient within one 'base configuration'. Typical trimcoil currents used were of the order of hundreds of μA to a few mA.

To determine the gradients of an arbitrary configuration, the field produced by each trimcoil as a function of the current applied must be measured. In order to do this, a series of dedicated measurements were performed. First, the apparatus was degaussed, with the B_0 coil switched off, and the remnant field was measured. The gradients $G_{l,m}$ of this field are determined using the standard analysis described in section 5.8. Second, a known current, typically +20mA, was applied to one chosen trimcoil, and the field mapped and gradients fitted again. A separate linear fit was performed of each extracted gradient $G_{l,m}$ versus applied current to that coil. By this method, a matrix of trimcoil gradients per unit current for each coil can be obtained, which when combined with knowledge of each coil current. An uncertainty on this gradient can be obtained by considering the error on the fit parameters and the χ^2 of this fit.

In 2013 and 2014, full sets of trimcoil maps (with associated remnant maps) was taken, with the exception of the vertical guide coil, the mounting of which was mechanically incompatible with the mapper. In 2017, a mounting was devised for the coil, and the three subcoils of the vertical guiding field coil were measured. A limited subset of the main trimcoils were mapped to validate the results of previous years. The results of fitting three selected gradients against the current in the coil 'TTC' are shown in Figure 5.64.



Figure 5.64: Example fits of selected gradients versus currents, for maps of TTC. Two sets of maps were measured, which are fitted with a common trimcoil field (directly proportional to applied current) and independent remnant fields, allowed to vary freely between degaussings.

5.11 Linearity

As an alternative route to establishing the field during each configuration measured during nEDM datataking, specific maps were taken of most 'Base Configurations' used during datataking. These base configurations describe one optimised uniform field after a degaussing, which are only modified by the addition of a vertical gradient comprised mainly of $G_{1,0}$ using a combination of the TTC and BTC coils.

As a cross-check, the results of these maps were compared to the results obtained by combining the mapped pure B_0 field with the trimcoil maps. To avoid any effect arising from the degaussing process, B_0 and configuration maps were only compared within one degaussing. For each gradient $G_{l,m}$, the difference between the B_0 map and the base configuration map was compared to the field predicted to be produced by a linear combination of the trimcoil gradients. The results of this analysis for the gradient $G_{3,0}$ are displayed in Figure 5.65, with the residuals plotted as a histogram in Figure 5.66. As $\langle B_T^2 \rangle$ is not a linear observable, it is not possible to analyse in this way, instead the measured $\langle B_T^2 \rangle$ from the base configuration map is compared to that predicted from the corresponding B_0 map plus trimcoils. The results of this analysis are plotted in Figure 5.67, with the residuals plotted as a histogram in Figure 5.67, with the residuals plotted as a histogram in Figure 5.68. The RMS error in both cases (assuming the trimcoil map error is negligible) is expected to be $\sqrt{2}$ times the statistical error on a single map, being the difference of two maps.



Figure 5.65: $G_{3,0}$ predicted from B_0 and trimcoil maps, versus measured $G_{3,0}$ from base configuration maps



Figure 5.66: Histogram of $G_{3,0}$ residuals when comparing configuration maps to values predicted from B_0 maps and trimcoil maps



Figure 5.67: $\langle B_T^2 \rangle$ predicted from B_0 and trimcoil maps, versus measured $\langle B_T^2 \rangle$ from base configuration maps



Figure 5.68: Histogram of $\langle B_T^2 \rangle$ residuals when comparing configuration maps to values predicted from B_0 maps and trimcoil maps

5.12 Evaluation of Gradients During nEDM Datataking

There are two possible approaches to predict the gradients within the chamber during datataking. The first approach is to start with a measurement of the unmodified B_0 field, then add the effect of the trimcoils in proportion to the currents used during datataking. This relies on the accuracy of each individual trimcoil map, and on there being no real nonlinearities in the trimcoil response (this clearly holds in free space, but one might imagine the complex behaviour of μ -metal leading to some unexpected effect). Alternatively, one could use the individual base configuration maps, which are modified only by adding small currents in the TTC and BTC coils to add a small vertical gradient $G_{1,0}$. Both methods were shown to produce compatible results in section 5.11.

The uncertainty on the gradients obtained arises from several sources assessed throughout this chapter:

- Statistical error on the mapping measurements (section 5.8.5)
- Systematic error on the mapping measurements (subsection 5.8.5)
- Changes in magnetic configuration between 2014 and 2017
- Reproducibility of the magnetic field within the chamber between degaussings (section 5.9)
- Accuracy of trimcoil maps and validity of reconstruction methodology (section 5.11)

Of these, the leading systematic uncertainty arises from the error in the mapper process itself, in particular that found by restricting the domain of points analysed, which serves as a very general test (section 5.8.5). The exact cause has not been established, but this may be, for example, due to positioning uncertainties or nonlinearities, such as a scaling error in one axis. This gives a systematic error in the determination of $G_{3,0}$ of ± 0.0051 pT/cm³. The leading statistical error arises from the reproducibility of the magnetic field after degaussing. The mean $\langle B_T^2 \rangle$ measured for a B_0 up map (without trimcoils) was 1.83 nT², with a standard deviation of 0.27 nT². For B_0 down, this figure was 3.00nT² with a standard deviation of 0.25 nT². For $G_{3,0}$, the mean figure measured for B_0 up was -0.0698 pT/cm³ with a standard deviation of 0.0033 pT/cm³, while the corresponding figure for B_0 down was 0.0622pT/cm³ with a standard deviation of 0.0053 pT/cm³. The statistical error on each map is not a limiting factor in the analysis, being substantially less than the reproducibility. In order to minimise the error on the obtained gradients due to the degaussing reproducibility, it was chosen to use the path starting from the raw B_0 maps. As base configuration maps were taken only, the reproducibility error would come twice (causing an increase of $\sqrt{2}$ in the statistical error), as both the reproducibility error incurred during the mapping campaign and the reproducibility error incurred during nEDM datataking must be considered. As many B_0 maps were taken, these can be averaged to reduce the effect of the reproducibility during mapping by a factor $1/\sqrt{N}$, resulting in a modest but welcome improvement in the final statistical sensitivity.

As the results from this analysis are used as a direct input to the main nEDM analysis, the uncertainty in the result will cause a loss in sensitivity for the entire experiment (see section 4.7). The corrections applied are

$$d_{\rm n}^{corr} = d_{\rm n}^{\rm meas} - \frac{\hbar \left|\gamma_{\rm n}\gamma_{\rm Hg}\right|}{8c^2} r^2 \left(\frac{R^2}{4} + \frac{H^2}{10}\right) G_{3,0}$$
(5.29)

$$R^{\rm corr} = R^{\rm meas} / (1 + \delta_T + \delta_{\rm Earth} + \delta_{\rm other}).$$
(5.30)

Therefore the error on the determination of $G_{3,0}$ will be:

$$\sigma_{d_{n}} = \frac{\hbar |\gamma_{n} \gamma_{Hg}|}{8c^{2}} r^{2} \left(\frac{R^{2}}{4} + \frac{H^{2}}{10}\right) \sigma_{G_{3,0}}$$
(5.31)

$$\approx 5.44 \times 10^{-25} \text{ecm} \frac{\sigma_{G_{3,0}}}{1 \text{pT/cm}^3}.$$
 (5.32)

Therefore, the achieved systematic error on $G_{3,0}$ of 0.0051 pT/cm³ will correspond to a final systematic error contribution of approximately 2×10^{-27} e cm, substantially below the statistical sensitivity of the experiment. As this is a systematic error, it must add in quadrature with the final EDM error.

Errors in the correction of $\langle B_T^2 \rangle$ will enter indirectly via R. The error on δ_T will be

$$\sigma_{\delta_T} = \frac{\sigma_{\langle B_T^2 \rangle}}{2B_0^2}.$$
(5.33)

The gradient of the crossing lines plot is

$$\frac{\partial d_{\rm n}}{\partial R} = B_0 \frac{\hbar \gamma_{\rm Hg}^2}{8c^2 \langle z \rangle} \tag{5.34}$$

where $\langle z \rangle$ is the height difference between the mercury and the neutrons. It will be fixed at $\langle z \rangle = 0.35$ cm for this calculation. Combining these, the error contributed to $d_{\rm n}$ from the $\langle B_T^2 \rangle$ correction will be

$$\sigma_{d_{n}} = \frac{\partial d_{n}}{\partial R} \sigma_{R} \tag{5.35}$$

$$= B_0 \frac{\hbar \gamma_{\rm Hg}^2}{8c^2 \langle z \rangle} \left(\delta_T * \left| \frac{\gamma_{\rm n}}{\gamma_{\rm Hg}} \right| \right)$$
(5.36)

$$\approx 6.28 \times 10^{-27} \mathrm{ecm} \frac{\sigma \langle B_T^2 \rangle}{\mathrm{1nT}^2}.$$
(5.37)

For the attained reproducibility error on $\langle B_T^2 \rangle$ of $\approx 0.3 \text{ nT}^2$, this will correspond to an error on d_n of $\approx 2 \times 10^{-27}$ e cm. However, errors on the reproducibility of the field are statistical: the shield was degaussed approximately once per month, and at each degaussing the field was reset at a new configuration. The statistical sensitivity of each period of datataking between degaussings was of the order several $\times 10^{-26}$ ecm. Therefore, this effect will be reduced by a factor of approximately $1/\sqrt{N_{\text{degaussings}}}$ over the entire nEDM datataking period, reducing it to insignificance.

The gradients obtained from the mapping analysis using this technique are used directly as an input to the final crossing lines analysis, to compensate for these errors on a run-by-run basis.

5.13 Conclusion

In this chapter, offline measurements using a specially designed field mapping robot were used to measure the magnetic field within the vacuum chamber of the experiment to high precision and accuracy. The resulting maps were analysed and decomposed in a basis of gradients in order to extract quantities relevant to the compensation of systematic effects in the experiment. In order to understand the accuracy of this compensation, the accuracy of the mapping process itself, as well as the reproducibility of the field and the accuracy of reconstruction of the gradients of real trimcoil configurations were investigated. The results of this analysis are used to compensate systematic effects in the main experiment, reducing the residual systematic error on the nEDM from such magnetic field effects to a level of $\sigma d_{n,mapping} \approx 2 \times 10^{-27}$ ecm, well below the statistical sensitivity of $\sigma d_{n,stat} \approx$ 1×10^{-26} ecm and within the error budget for this effect of $\sigma d_{n,mapping} \approx 5 \times 10^{-27}$ ecm. This is a substantial improvement over previous efforts to measure this effect in 2013 and 2014, since which time the systematic uncertainty on $G_{3,0}$ (which could have been the limiting factor in the overall experimental result) has been improved by a factor 4.

Chapter 6

Search for Axion-Like Matter Through Nuclear Spin Precession in Electric and Magnetic Fields

The axion is a hypothetical particle that arises when the standard model is modified by adding the Peccei-Quinn symmetry. This symmetry was first proposed in 1977 by Peccei and Quinn [23–27], in order to explain the lack of CP violation in the strong force - the socalled strong CP problem. In addition, particles with these characteristics are attractive dark matter candidates [187]. Until recently, experiments searching for axions have mostly focused on the coupling of axions to [17], and relied on the conversion of axions to photons (and vice-versa) occurring - meaning the experimental observable is proportional to a very small coupling constant squared or even to the fourth power. However, in recent years it has been suggested that the axion may be observed by its coupling to gluons [188–190]. Additionally, given the non-detection of any WIMP dark matter candidates at the time of writing [17], the suggestion that the dark matter particle(s) may be light (significantly below the mass of nucleons) has gained in popularity. Therefore, a new generation of axion searches using laboratory experiments has begun, searching for the influence of dark matter axions on precision systems. These searches have the benefit that the observable is linear in the coupling constant[188-192], and the axions do not need to be produced: they are already there in the background. In this chapter a novel analysis placing new constraints on the couplings of cosmologically-interesting axions to gluons by searching for an oscillating neutron electric dipole moment is presented. These results have been previously published in [161].

6.1 Motivation and Theory

6.1.1 The Strong CP Problem and Axions

In Quantum Chromodynamics (QCD), the theory of the strong force, it is possible to write a term [18, 19]

$$\mathcal{L} = \frac{g_s^2}{32\pi^2} \bar{\theta} G^a_{\mu\nu} \tilde{G}^{a\mu\nu}, \tag{6.1}$$

where g_s is the strong coupling constant, $G^a_{\mu\nu}$ is the gluon field and $\tilde{G}^{a\mu\nu}$ is its dual and $\bar{\theta}$ is a *CP*-violating phase. This term does not affect physics through processes which are described by Feynman diagrams as it is a total derivative, but produces so-called instanton effects. It can be shown that QCD can only depend on the combination [19]

$$\bar{\theta} = \theta - \arg \det m_q, \tag{6.2}$$

where m_q is the quark mass matrix and θ is another *CP*-violating phase. It is therefore possible to 'move' *CP* violation between the quark mass matrix and θ . If θ is nonzero, or the quark masses are complex, then CP and P are violated by this term as $G^a_{\mu\nu}\tilde{G}^{a\mu\nu}$ is CP-odd. This also implies that any physical effects of $\bar{\theta}$ depend on $\bar{\theta}$ modulo 2π . This also implies that the CP violation and dependence on θ vanish if at least one quark is massless (though this possibility is strongly excluded [17]).

In practice, no CP violation has been observed in the strong sector. The strongest limits on this come from experimental limits on the (static) EDM of the neutron [21], and of mercury-199 [22]. A nonzero $\bar{\theta}$ would induce a CP violating pion-nucleon coupling, with constant $|g_{\pi NN}| \approx 0.038 |\bar{\theta}|$ [193]. This is small in comparison to the CP conserving coupling $g_{\pi NN} \approx 13.4$. Calculations using chiral perturbation theory showed that this mechanism can lead to a neutron EDM of $d_n = 5.2 * 10^{-16} \bar{\theta}$ e cm [193]. A more refined calculation using chiral perturbation theory gives $|g_{\pi NN}| \approx 0.016 |\bar{\theta}|$ [194]. This confirms a previous result using a QCD sum rules technique [195] that $d_n = 1.2 * 10^{-16} \bar{\theta}$ e cm (to an estimated accuracy of 45-50%).

By considering the current limit on the nEDM from the previous generation Sussex-RAL-ILL experiment [21] of $|d_n| \leq 3.0 \times 10^{-26}$ ecm, it is possible to find a bound on $\bar{\theta}$ of approximately $|\bar{\theta}| \leq 3 \times 10^{-10}$. As $\bar{\theta}$ is an angle, it would naturally be expected to be of order 1. This represents a significant fine-tuning type problem, the so-called Strong CP Problem.

A full exposition of the Peccei-Quinn theory is beyond the scope of this thesis. However, in short an additional $U(1)_{PQ}$ symmetry is added to the Standard Model. This symmetry is spontaneously broken at a scale f_a , which results in a new Nambu-Goldstone boson - the axion. This dynamic field (labelled a) then effectively replaces the static angle θ . A term

$$\mathcal{L} = \xi \frac{a}{f_a} \frac{g_s^2}{32\pi^2} \bar{\theta} G^a_{\mu\nu} \tilde{G}^{a\mu\nu}$$
(6.3)

is added to the Lagrangian, with similar form to the QCD θ term. The potential is minimised at $\langle a \rangle = -\frac{f_a}{\xi} \bar{\theta}$, which leads to a cancellation of the physical effects of the $\bar{\theta}$ term.

The mass and couplings of the axion are then model specific. The original Peccei Quinn model proposed that the symmetry was broken at the electroweak scale, $f_a \approx 250$ GeV, however this model was quickly ruled out by experiment [17]. However, two new major types of so-called 'invisible axion' models have been proposed: KSVZ[27, 196] and DFSZ[197, 198], named for their creators. It can be shown that in each of these models, the axion mass is related to the symmetry breaking scale by [18]:

$$m_a \approx 6.3 \frac{10^6 \text{GeV}}{f_a} \text{eV}$$
 (6.4)

The couplings to matter are also dependent on f_a , however the exact dependence of each coupling is model dependent [18]. Constraints will be found on the coupling constants as a function of axion mass, giving model-independent results. The symmetry breaking scale can be extremely large, and thus axion masses very light.

6.1.2 Ultralight Axions as a Dark Matter Candidate

Dark matter was first proposed by the astronomer Fritz Zwicky in 1933 to explain anomalies in the rotation of galaxies [199]. Since then, astronomers and cosmologists have confirmed these observations[28], and dark matter remains a standard feature of most cosmological models. However, the question still stands: what is dark matter made of? Many have tried over the years to answer this question, however still dark matter has never been directly detected. One of the dominant hypotheses has been that dark matter consists of heavy WIMPs (*Weakly Interacting Massive Particles*), which interact through only gravity and the weak force, having a mass in the 100 GeV range. Supersymmetric theories predict many hypothetical particles, the lightest of which (the so-called 'Lightest Supersymmetric Particle') could be a stable, heavy dark matter candidate. However, with the non-discovery of Supersymmetry or any WIMP Dark Matter candidate particle at the LHC [17], such theories are coming under question.

One alternate hypothesis, is that dark matter instead consists of extremely light axions, or similar axionlike particles. This is plausible over a wide range of mass scales[187, 200, 201], however particles with masses in the range of $10^{-24} - 10^{-18} eV$ are the focus of this chapter. In such mass ranges, the number-density of particles is sufficiently high that the axion field will appear as a coherent wave in a classical field, oscillating at a frequency proportional to the mass of the particle. Such dark matter is additionally able to solve several issues with structure formation[187, 202]. Additionally, a feasible nonthermal production mechanism in the early universe exists in the form of Vacuum Misalignment[203-205]. Requiring that the axion dark matter does not have a de-Broglie wavelength larger than the smallest Dwarf galaxies results in a lower bound on the possible mass of $m_a \geq 10^{-22} eV[206]$.

6.1.3 Effects of Axionic Dark Matter in nEDM Experiments

As previously observed in subsection 6.1.1, adding the axion to the Standard Model serves to negate the physical effects of a nonzero $\bar{\theta}$ by absorbing $\bar{\theta}$ into the field a, when the potential is at its minimum. This works in the low density limit, but once the density of axionic dark matter is considered, a oscillates about this minimum depending on the local dark matter density. Since this interferes with the cancellation, it then effectively acts like an oscillating value of θ , which leads to measurable effects in sensitive laboratory experiments.

Two specific interaction terms were considered, representing the axion-gluon and the axion-nucleon couplings:

$$\mathcal{L}_{\rm int} = \frac{C_G}{f_a} \frac{g_s^2}{32\pi^2} a G^b_{\mu\nu} \tilde{G}^{b\mu\nu} - \frac{C_N}{2f_a} \partial_\mu a \bar{N} \gamma^\mu \gamma^5 N \tag{6.5}$$

where N is the nucleon field and C_G and C_N are model dependent dimensionless parameters characterising the coupling strength of axions to gluons and nucleons.

The gluonic coupling looks like the QCD θ term, with a proportionality to the axion field. By the same mechanism as the original θ term [193–195], this term leads to the value of the measured neutron EDM being proportional to the axion field. If axions make up a substantial portion of the local dark matter density, behaving as a coherent oscillation, this means that one would expect to see an oscillation in the measured value of the neutron EDM. Therefore, by constraining a possible oscillation in the nEDM, one can constrain the amplitude of the coupling between axionic dark matter and gluons[188–191].

The coupling to nucleons is a derivative type coupling. If the form of the axion field is like $a = a_0 \cos(m_a t - \vec{\sigma_N} \cdot \vec{p_a})$ (where $\vec{\sigma_N}$ is the spin vector of the nucleons and $\vec{p_a}$ is the momentum of the axion dark matter), then this will result in a energy shift in a system of polarised nucleons of

$$H_{\rm int} = \frac{C_N}{2f_a} \sin(m_a t) \vec{\sigma}_N \cdot \vec{p}_a.$$
(6.6)

This will result in a shift in the measured precession frequency depending on the phase of the axion field and upon the angle between the quantisation axis of the system and the momentum. For experiments not aligned along the Earth's rotation axis, this will give an additional 24-hour sidereal modulation. This can be seen by expressing in a nonrotating celestial coordinate system (see [207]):

$$\vec{\sigma}_N \cdot \vec{p}_a = \hat{m}_F f(\sigma_N) m_a |\vec{v}_a| \\ \times \left[\cos(\chi) \sin(\delta) + \sin(\chi) \cos(\delta) \cos(\Omega_{\rm sid} t - \eta) \right]$$
(6.7)

where χ is the angle between Earth's axis of rotation and the spin quantization axis $(\chi = 49.4^{\circ} \text{ at the ILL}, \text{ while } \chi = 42.5^{\circ} \text{ at the location of the PSI, for systems with the magnetic field oriented vertically}, <math>\delta \approx -48^{\circ}$ and $\eta \approx 138^{\circ}$ are the declination and right ascension of the galactic axion DM flux relative to the Solar System, $\Omega_{\text{sid}} \approx 7.29 \times 10^{-5} \text{ s}^{-1}$ is the daily sidereal angular frequency, $\hat{m}_F = m_F/F$ is the normalized projection of the total angular momentum onto the quantization axis, and $f(\sigma_N) = +1$ for the free neutron, while $f(\sigma_N) = -1/3$ for the ¹⁹⁹Hg atom in the Schmidt (single-particle) model. It may be noted that by varying the orientation of the magnetic field, the value of χ may be varied, which will tune the relative amplitudes of the daily sidereal modulation and the constant term. The sign of the modulation will invert with the sign of the applied magnetic field. This effect is also referred to as the axion wind effect, and often is likened to a 'pseudomagnetic field'.

6.1.4 Other Experimental Probes of Axions and Axionic Dark Matter

Axions are predicted to couple to matter in many ways. The most commonly probed is the axion to two photon coupling. A summary of current constraints on this coupling published by the PDG [17] is reproduced in Figure 6.1. Some of the first experiments probing this coupling are so-called 'light shining through walls' experiments, and similar efforts continue to this day. In short, this type of experiment relies on the Primakoff production of axions from photons in a high magnetic field region, followed by transmission through a wall impermeable to photons, followed by a conversion back to a photon on the other side. Such experiments are sensitive to axions, but also a broad spectrum of other light weakly coupled particles, such as dark photons. A review is presented in [208]. The largest limitation of this type of experiment is that it relies on an extremely rare process occurring twice, meaning that the observable is then proportional to the very small coupling constant to the power 4.

One additional result of an axion-photon coupling would be axion production inside stars, through the Primakoff effect. This can be probed by both astrophysical observations, and attempts to directly detect the axions produced inside the sun. Such experiments are called Axion Helioscopes, and set limits on the coupling across a wide mass range. The most recent is CAST [209] (CERN Axion Solar Telescope), which utilises a refurbished LHC dipole magnet (9T) which can be aimed at the sun. In the high field region, axions produced in the sun could be converted back to photons which are then detected in the cavity. The sun would be an abundant source of axions, however the observable is then still only proportional to the small axion-photon coupling to the fourth power, as both production and conversion must occur.

Another interesting experiment probing the same coupling, ADMX [210–213], instead attempts to search for the same cosmic dark-matter axions produced in the early universe as this work, though at a different mass scale. The experiment utilises a 136L microwave cavity, tunable over a range of resonances, within a superconducting magnet, operated at fields of up to 7.6 T. The oscillating axion field interacting with the strong magnetic field would drive the resonant cavity, producing a signal which could be detected using a sensitive SQUID microwave detector. The apparatus had previously been used to exclude a portion of the mass range of the KSVZ axion, between 1.9 and $3.3\mu eV$, corresponding to frequencies of a few hundred MHz. More recently, the apparatus has been upgraded, and used to probe axion masses between 2.66 and $2.81 \mu eV$ with a sensitivity great enough to probe the parameter space of the more weakly coupled DFSZ axion. [213] As this experiment probes already-existing dark matter axions, of a density set by the assumption that they make up the majority of the local dark matter density, only one axion-photon conversion is necessary. This means that the observable (the power of the signal generated) is then proportional to only the small coupling constant squared, rather than to the power 4, a clear improvement in the long term prospects. The largest drawback of this kind of tuned-cavity search is that the sensitive frequency range must be scanned slowly to acquire sensitivity to the entire mass range, and that the tunable range of such a cavity (while retaining a high Q) is relatively limited. There are also practical limitations on the largest size of cavity that can be produced, as they must be maintained at cryogenic temperatures to reduce noise.

Until this work, the coupling of axions to gluons remained unexplored by laboratory



Figure 6.1: Existing laboratory and observational constraints on axion-photon coupling, figure taken from [17], based on work cited within. Does not include latest results from ADMX in [213].

experiments. However, the CASPEr collaboration proposes a dedicated experiment [214, 215] to probe this parameter space. Much like this work, the CASPEr experiment essentially searches for an oscillating EDM, in this case of ²⁰⁷Pb. The sample is placed in a large magnetic field, of order 10T, and becomes magnetised along that axis. This field is then reduced to a set level and a transverse electric field is applied. If the oscillation frequency of the oscillating EDM matches the Larmor frequency of the spins in the given magnetic field, then the spins will begin to acquire a transverse polarisation, and precess about the axis of the magnetic field. This magnetisation can be read out using a SQUID magnetometer. The sensitive frequency can tuned by varying the magnetic field over a very large range. The original CASPEr proposal proposes to cover a mass range from 10^{-6} eV to 10^{-14} eV[214]. The experiment is foreseen to run in two stages, with the first using only known technologies to reach a sensitivity below that needed to resolve KSVZ axions, while a second stage [215] would achieve the sensitivity needed to exclude KSVZ and DFSZ axions below 10^{-9} eV down to 10^{-14} eV. An alternative measurement scheme utilising sidebands was proposed to cover the range from 10^{-14} to 10^{-17} eV, meeting with the upper mass range probed by the work presented in this chapter.

It should be noted that astrophysical and cosmological observations also provide powerful constraints and motivations on the couplings of such particles, such as those described in the review [187].

6.2 Experiment

6.2.1 The Sussex-RAL-ILL Room Temperature EDM Experiment

The Sussex-RAL-ILL experiment room temperature nEDM experiment [20][86][87][88][21] sets (at the time of writing) the strongest limit on the static electric dipole moment of the neutron, though it is set to be superseded by its spiritual successor, the PSI nEDM experiment. The Sussex-RAL-ILL experiment took data from 1998 to 2002, and achieved a final result of $d_n = -0.21 \pm 1.82 \times 10^{-26} ecm$, corresponding to an upper limit of $d_n < 3.0 \times 10^{-26} ecm$ (at 90% CL). The experiment is extremely similar in principle to the PSI experiment. It ran from the turbine UCN source at the PF2 beamline at the Institut Laue-Langevin in Grenoble, France.

6.3 Analysis of the Sussex-RAL-ILL Dataset

It has been established in section 6.1 that if dark matter is comprised entirely of axions or similar axionlike particles, this would manifest as a time oscillation in the real measured value of the EDMs of the neutron of

$$d_{\rm n}(t) \approx +2.4 \times 10^{-16} \ \frac{C_G a_0}{f_a} \cos(m_a t) \ e \cdot {\rm cm} \,,$$
 (6.8)

and of mercury atoms of

$$d_{\rm Hg}(t) \approx +1.3 \times 10^{-19} \ \frac{C_G a_0}{f_a} \cos(m_a t) \ e \cdot {\rm cm} \,.$$
 (6.9)

A time series of the experimental data from the previous generation Sussex-RAL-ILL experiment and the current-generation PSI nEDM experiment has been analysed to test for a statistically significant oscillation, and then to constrain the amplitude of a possible oscillating EDM, and finally to interpret this as a constraint on the possible coupling between axions and gluons, assuming that these axions constitute the entirety of the local dark matter density. First, the Least Squares Spectral Analysis (LSSA) technique (analogous to a discrete Fourier transform) will be outlined, then the methods to test if a signal is statistically significant will be explained, and finally the use of the CL_S technique used for limit-setting will be described.
6.3.1 Preparation of a Time Series from the Sussex-RAL-ILL Experimental Data

In order to extract information about any oscillations in the measured value of the nEDM, a preliminary analysis is needed to turn the raw data into a time series of the measured neutron EDM. This stage of the analysis is shared with the static nEDM analysis as documented in detail in [21], and very similar to the general analysis strategy used to fit the data from the PSI experiment documented in chapter 4, but will be briefly summarised here. The raw experimental data from the apparatus is at its core a measurement of the neutron counts and Hg magnetometer readout from each four minute experimental cycle. To convert this to a single measurement of the neutron EDM, many cycles must be combined to measure a Ramsey curve, from which the neutron precession frequency can be extracted. By comparing measurements at different electric field states, a value for the neutron electric dipole moment can be extracted. The data is binned into segments of continuous running with no change applied to the magnetic field environment. This results in a single value for the EDM measurement, constituting the average over several hours, up to two days. Runs were not synchronised in start time or in length, and therefore vary in sensitivity by an order of magnitude. The extracted values for the EDM are corrected for systematics using the crossing lines technique described more fully in [21] and chapter 4. The result of this is a time series of measured neutron EDM as a function of time. In total, 545 measurements were collected between 1998 and 2002, with the majority of the sensitivity accumulated during the final two years.

6.3.2 Least Squares Spectral Analysis and the Lomb-Scargle Periodogram

The most well-known technique to search for an oscillation in an experimental dataset is the discrete Fourier transform, which converts a series of (possibly complex) numbers $x_0, x_1, \ldots, x_{N-1}$ to a sum of frequencies $\omega_n = \frac{2n\pi}{N} f_s$ (where f_s is the sampling frequency) with complex amplitudes $X_0, X_1, \ldots, X_{N-1}$ (where the argument determines the phase of each frequency component)

$$X_k = \sum_{n=0}^{N-1} x_n \cdot \exp\left(-i2\pi kn/N\right) \,. \tag{6.10}$$

In order to remove the effect of phase, one can instead compute the *Schuster Periodogram* or *Classical Periodogram* [216, 217]

$$P_S(f) = \frac{1}{N} \left| \sum_{n=0}^{N-1} x_n \cdot \exp\left(-i2\pi kn/N\right) \right|^2$$
(6.11)

$$= \frac{1}{N} \left[\left(\sum_{n=0}^{N-1} x_n \cdot \cos\left(2\pi f t_n\right) \right)^2 + \left(\sum_{n=0}^{N-1} x_n \cdot \sin\left(2\pi f t_n\right) \right)^2 \right].$$
(6.12)

This *periodogram* is an estimator of the *power spectrum* of the input data. The classical periodogram has the useful property that when applied to Gaussian noise, the results are chi-squared distributed, allowing for simple statistical tests of the significance of a signal. While the classical periodogram can be calculated for unevenly spaced datasets, unfortunately the statistical properties are not simple in this case.

It was modified by Scargle [218] to produce a periodogram technique with more useful statistical properties, given a full exposition in the cited work but summarised below. This was proven to be equivalent to the linear fitting of sines and cosines to the dataset at the target frequency, as investigated by Lomb [219], leading it to be named the *Lomb-Scargle* periodogram.

$$P_{LS}(f) = \frac{1}{2} \left[\frac{\left(\sum_{n=0}^{N-1} x_n \cdot \cos\left(2\pi f(t_n - \tau)\right)\right)^2}{\sum_{n=0}^{N-1} \cos^2(2\pi f(t_n - \tau))} + \frac{\left(\sum_{n=0}^{N-1} x_n \cdot \sin\left(2\pi f(t_n - \tau)\right)\right)^2}{\sum_{n=0}^{N-1} \cos^2(2\pi f(t_n - \tau))} \right],$$
(6.13)

with τ specified at each frequency to be

$$\tau = \frac{1}{4\pi f} \arctan\left(\frac{\sum_{n=0}^{N-1} x_n \cdot \sin(4\pi f t_n)}{\sum_{n=0}^{N-1} x_n \cdot \cos(4\pi f t_n)}\right).$$
 (6.14)

This formulation differs from the classical periodogram only in the denominator of each of the two sum terms, therefore it produces similar results to the classical periodogram in many cases of interest. In the case of equally spaced measurements, it reduces to the classical periodogram.

It can be shown that the probability distribution function (PDF) of the power $Z = P_{LS}(f)$ at an arbitrary frequency f, given an input of Gaussian noise with a variance σ_0^2 , is

$$p(Z=z)dz = \exp\left(\frac{z}{\sigma_0^2}\right)dz.$$
(6.15)

The mean power is then

$$\bar{Z} = \int_0^\infty z p(Z=z) dz = \int_0^\infty z \exp\left(\frac{z}{\sigma_0^2}\right) dz$$
(6.16)

$$=\sigma_0^4. \tag{6.17}$$

The cumulative distribution function (CDF) can be shown to be

$$p(Z < z) = 1 - \exp\left(\frac{z}{\sigma_0^2}\right). \tag{6.18}$$

It therefore follows that the statistical significance of a large excess is

$$p(Z > z) = \exp\left(\frac{z}{\sigma_0^2}\right). \tag{6.19}$$

The Lomb-Scargle periodogram also has the interesting property that it is equivalent to the least-squares fitting of a sinusoid of the given frequency and of arbitrary phase to the data, in the case where errors on each observation are equal. This immediately suggests a method to extend the Lomb-Scargle periodogram to cover the case of non-equal errors on each observation, while staying as close to the convenient statistical properties of the Lomb-Scargle periodogram: the least squares fitting of a sinusoid to the dataset. However, in the case of nonuniform errors, it is not immediately clear how to replace the $1/\sigma_0^2$ denominator in Equation 6.15, Equation 6.18 and Equation 6.19. This will be analysed by using a Monte Carlo technique.

6.3.3 Estimation of the CDF at one Frequency Under the Null Hypothesis

In the absence of a suitable analytic expression, it was decided to use a Monte Carlo approach to establish the CDF. N Monte Carlo fake datasets were generated, with the same time structure as the real data, consisting of Gaussian noise centred at zero at each measured data point, with the standard deviation of the Gaussian equal to the statistical error on each run. At each frequency f_i to be investigated, a least squares fit of

$$d_n(t) = A_i \cos(2\pi f_i t) + B_i \sin(2\pi f_i t)$$
(6.20)

was done to each generated time series. As shown above, the Lomb-Scargle Periodogram is equivalent to least squares fitting in the case of equal and uncorrelated errors, so using the result of [218], the least-squares power $Z_i = A_i^2 + B_i^2$ at a frequency f_i follows a form related to Equation 6.18:

$$p(Z < z) = 1 - \exp(\lambda_i z), \qquad (6.21)$$

where λ_i is a coefficient to be determined for each frequency.

6.3.4 Search for Oscillations at any Frequency

As the frequency of any potential axion dark matter signal is not known a priori, it is necessary to compute the periodogram at wide range of frequencies. An appropriate upper and lower frequency and an appropriate spacing between frequencies must be selected. When dealing with a regularly sampled signal, the Nyquist sampling theorem states that the complete frequency content of a signal may be reconstructed completely only if band limited to half the sampling frequency f_s , giving an obvious choice of maximum frequency. If the observation window is limited to a length T, any frequency peaks in the power spectrum will be smeared to a sinc function of width $\Delta f = 1/T$. This means therefore that there are only $N = T/f_s$ independent frequencies to sample [217].

In the case of unevenly sampled data, the choice of frequencies is not so trivial. It can be possible to meaningfully access higher than the Nyquist frequency (for a measurement with the same number of samples taken over the same total measurement time). Additionally, there is the complication that our measurements are of finite precision, and take a finite time to perform (typically 1 day). Therefore, it is impossible to define analytically a convenient grid of *independent* frequencies to test.

It was decided to define the grid to scan to be linearly spaced with a frequency spacing of $\Delta f = 1/T$. The lowest frequency was chosen arbitrarily, to meet certain astrophysical limits on the hypothetical axion mass of 1×10^{-24} eV. The highest frequency was also arbitrarily chosen to correspond to a period of 0.1 days, to match where the sensitivity of the experiment to real signals would be significantly diminished due to the time averaging effect over individual measurement runs, each taking a minimum of several hours.

The expected distribution of power at a specific frequency has previously been found, however as it is desired to search across a large number of frequencies, one must account for the fact that when making N statistical tests, one would be likely to be significant at a p-value of 1/N. This is the so-called *problem of multiple comparisons*, or the *look-elsewhere* effect. This is corrected for by applying a correction to convert between the *local* p-value of a particular observation at a specified frequency, and the *global* p-value of a particular observation in the context of testing a large number of frequencies. For N independent tests, the conversion is

$$p_{\text{global}} = 1 - (1 - p_{\text{local}})^N.$$
 (6.22)

This formula is used to define so-called false alarm thresholds: the required local p-value in order to claim significance at a specified global p-value.

In deriving this formula, the assumption has been made that all of these tests are statistically independent. However, as the real dataset deviates substantially from the simplest case of uniformly spaced data with uniform error bars, this may no longer be the case, as the set of frequencies investigated is arbitrary. In particular, one might consider the case where the frequency range is significantly oversampled: in this case Equation 6.22 would significantly overestimate the false alarm levels, because frequencies spaced less than the $\Delta f = 1/T$ linewidth would be strongly correlated. As a result of this, it was chosen to instead define an 'effective number of frequencies' N_{eff} . This was defined by analysing the same Monte Carlo null hypothesis datasets used to evaluate the CDF of power for each frequency. From each periodogram, the least probable peak was selected. The *p*-value of each peak was collected together. A maximum likelihood estimator of N_{eff} was derived

$$N_{\rm eff} = \frac{-N_{\rm Periodograms}}{\Sigma_i \ln(1 - p_{i,\min})},\tag{6.23}$$

where $p_{i,\min}$ is local *p*-value of the least probable excess of each periodogram *i*. Using this, false alarm thresholds that are consistent even when oversampling the frequency space can be found. The effective number of frequencies was fitted to be $N_{\text{eff}} = 1026$, compared to N = 1334 frequencies tested in total.

The results of this analysis are presented in Figure 6.2. The global p-value of the most significant peak is 0.53, consistent with non-detection.



Figure 6.2: The results of a detection analysis as a function fitted amplitude versus trial frequency. The black line represents the actual measured periodogram of the dataset, while green indicates the mean power measured in each frequency across all generated Monte Carlo datasets. The orange line indicates the computed power necessary for significance at levels of 1-5 σ , for the fitted $N_{\text{eff}} = 1026$ (the true number of frequencies tested was N = 1334). If the black line crosses the *n*th orange line, it is significant at $n\sigma$.

6.3.5 Exclusion of oscillations

The probability distribution of the power under the null hypothesis was previously established in subsection 6.3.3. In order to estimate the probability distribution under a particular signal hypothesis, a similar Monte Carlo technique was used. Fake datasets are generated, however a sinusoidal signal of specified frequency ω_{test} and amplitude A_{test} is injected in addition to the expected Gaussian noise. The fake data is of the form

$$d_{n,i} = \frac{A_{\text{test}}}{t_{\text{end},i} - t_{\text{start},i}} \left(\sin(\omega_{\text{test}} t_{\text{end},i} + \delta) - \sin(\omega_{\text{test}} t_{\text{start},i} + \delta) \right) + N(0,\sigma_i^2)$$
(6.24)

where δ is a phase, drawn from $U(0, 2\pi)$. This is equivalent to a signal of form $A_{\text{test}} \cos \omega_{\text{test}} t_i + \delta$ averaged over the total observation time $t_{\text{end},i} - t_{\text{start},i}$ plus Gaussian noise. Following this, the Lomb-Scargle periodogram was computed at only the test frequency ω_{test} to extract the power that would be measured at that frequency. At each pair ($\omega_{\text{test}}, A_{\text{test}}$), 200 fake periodograms were generated. Then, these results were used directly to estimate the distribution, without smoothing or extrapolation.

As can be seen in Figure 6.3, in some frequency ranges, the apparent exclusion when considering the *p*-value of the signal hypothesis is anomalously low - significantly lower than the limit on the static EDM. This arises because the power measured in that frequency fluctuated anomalously low - causing an artificially strong exclusion. It was decided to use the CL_s technique [17, 220] to find an exclusion on the amplitude of a potential oscillation. This is a frequentist technique popular in detection-type particle physics analyses which aim to search for a small signal in noisy data. The statistic is given by (using the notation of [220])

$$CL_s = \frac{CL_{s+b}}{CL_b} = \frac{1 - p(z \ge Z|\text{signal})}{1 - p(z \ge Z|\text{background})}.$$
(6.25)

The effect of this is to modify the measured *p*-values to avoid claiming exclusions in regions where the analysis would not be able to claim a detection.

It was chosen to compute the exclusion at 200 frequencies spaced logarithmically across the entire range of frequencies analysed in the detection analysis. As the measurement timings were not consistent, no sharp resonant regions of increased or decreased sensitivity were expected. This results in a significant decrease in the amount of computation power needed to compute the exclusion region. The CL_S statistic is computed at 100 logarithmically spaced amplitudes from 1×10^{-26} ecm to 100×10^{-26} ecm. To find a 95% exclusion region, the 0.05 isocontour of CL_S was found. This is presented in Figure 6.3.



Figure 6.3: Comparison of exclusion results obtained using the CL_S technique(below), versus a naïve P(data|signal) (above). The blue line represents the boundary of the 95% exclusion region. The area in black is strongly excluded, while the yellow regions are allowed. The simple method results in claims of anomalously strong (stronger than the static EDM limit) exclusions in regions where the power spectrum fluctuated low; claims of such sensitivity must be unphysical.

6.4 Analysis of the PSI dataset

6.4.1 Fit Model

A closely related and complementary analysis was performed on the dataset arising from the current-generation nEDM experiment[91, 221], which collected data from 2015-2017. Only the dataset from 2015 and 2016 was included in the analysis. The dataset spanned 506 days. In order to probe a higher frequency range than was probed by the analysis of the ILL dataset, a different strategy was employed. It was decided, that rather than analysing the measured EDM value for entire runs (meaning that signals with period comparable to or shorter than the total run length would be averaged out), the ratio of neutron to mercury precession frequencies extracted from single five-minute experimental cycles would be analysed. By this method, drifts in the mean B_0 field would be automatically corrected. Additionally, drifts in the linear gradient $G_{1,0}$ were corrected using data from the caesium magnetometer array [154].

One complication is that the gradient $G_{1,0}$ was deliberately changed every few days within the apparatus. This step change would cause a step change in the frequency ratio R, causing spurious excesses in the periodogram if not accounted for. As described more thoroughly in chapter 4, the caesium magnetometer array is precise and sensitive enough to usefully compensate for drifts in the gradient $G_{1,0}$, but not accurate enough to use as the only measure of the vertical gradient $G_{1,0}$. Thus, it was chosen to additionally fit a free offset in the value of R to each individual run, in addition to the hypothetical oscillation which would absorb these step-changes. Fitting these free offsets results in a large drop in sensitivity to oscillations slower than the typical run length, as these can be largely absorbed into the offsets. However, this region is already well constrained by the analysis of the Sussex-RAL-ILL dataset.

The model fitted to the data was

$$R_{\text{cycle}} = A\cos(\omega_{\text{test}}t) + B\sin(\omega_{\text{test}}t) + \delta R_{\text{run}}, \qquad (6.26)$$

where $A^2 + B^2$ is the power as in the previous analysis, and δR_{run} is a set of free offsets fitted to each run.

The dataset was split into three portions according to the relative direction of magnetic and electric fields: E parallel to B ($E \uparrow \uparrow B$), E antiparallel to $B(E \uparrow \downarrow B)$, and Eoff(E = 0).

6.4.2 Detection Analysis

Interpretation of these datasets in terms of detections and exclusions also mirrors the approach used in the analysis of the ILL dataset. However, as the frequency ratio R rather than the measured EDM was used, which is also sensitive to many non EDM-like frequency shifts, several tests were proposed to ensure that any significant signal really does have the characteristics of an oscillating EDM before claiming a credible candidate detection.

- A significant signal should be detected in both of the $E \neq 0$ datasets.
- No signal should be observed in the E = 0 dataset.
- The signal should have a phase difference of π between the $E \uparrow\uparrow B$ and $E \uparrow\downarrow B$ datasets.
- The signal should be a very narrow peak rather than a broad band excess in power (which would more likely indicate some unaccounted for environmental disturbance, such as some day-night cycle connected systematic)
- The signal should not coincide with periodic events in the operation of the apparatus

The range of frequencies to scan was selected to cover the axion mass range from 10^{-22} eV to 10^{-17} eV. The frequencies tested were spaced by the spectral resolution of the dataset (as for the ILL dataset) - 1/506 days = 23 nHz, giving 156198 trial frequencies in all, approximately 100× more than in the analysis of the Sussex-RAL-ILL dataset. The results of this analysis for each E state (E = 0, $E \uparrow B$ and $E \uparrow B$) are presented in Figure 6.4.

As normal experimental running in this fashion involves several periodic features, there are several areas where increases in the average power are expected. In particular, the typical cycle repetition rate of 303mHz (1/300 seconds) and the electric field reversal at 28μ Hz (corresponding to a 10 hour period - the HV was reversed each 5 hours). In the dataset where $E \uparrow\uparrow B$, there are 5 frequencies above the 3σ false alarm threshold. Two of these are within 100μ Hz of the cycle repetition frequency, and can therefore be discarded. Meanwhile, the other three excesses are in the low frequency region which is already very strongly excluded by analysis of the ILL dataset. In the dataset with $E \uparrow\downarrow B$, there are 3 significant excesses in the same regions, which can be disregarded. In the control dataset E = 0 there are no significant excesses. These results are consistent with non-detection.



Figure 6.4: Periodograms computed for the PSI dataset, with features as described in text. The black line is the periodogram computed for each electric field state, while the dark green shows the mean of all generated Monte Carlo periodograms under the null hypothesis. The three progressively lighter green bands represent the 1,2 and 3σ deviations extracted from the Monte Carlo data. Finally, the five orange lines indicate the 1,2,3,4 and 5 σ false alarm thresholds, as defined within the text. In each insert, the area around the inverse cycle spacing (1/(5minutes)) is shown, where a large increase in power is not unexpected. Apparent detections within this reason should be treated sceptically.

6.4.3 Exclusion Analysis

Proceeding with the analysis, the CL_S procedure was again used to obtain an exclusion in terms of oscillating EDM amplitude as a function of frequency, which can later be interpreted as limits on the axion-gluon coupling as a function of axion mass. The CL_S statistic is evaluated for 200 frequencies logarithmically spaced across the entire range tested, at 100 amplitudes logarithmically spaced between 1 and 100×10^{-26} ecm, and the 0.05 isocontour is computed. In generating the Monte Carlo samples, it is assumed that the measurement is averaged uniformly over the 180s free precession period of each measurement cycle. This computed exclusion region is presented in Figure 6.5. Combining this result with the ILL analysis results in complementary constraints over a large frequency range, as shown in Figure 6.6.



Figure 6.5: Computed exclusion region for EDM amplitude as a function of frequency using (left) the naïve P(data|signal) approach and (right) the CL_S method. Amplitude in units of 10^{-26} ecm. Compare with Figure 6.3.



Figure 6.6: Combined 95% exclusion region using CL_S analysis technique on both datasets.

6.4.4 Wind effect analysis

A very closely related analysis can be performed to constrain the axion-nucleon coupling. The coupling would cause a spin-dependent shift in the two energy levels of the neutron as per Equation 6.6, which would manifest as a modulation of the measured precession frequency. The modulation depends on the dot product of the mean momentum of the axion field and the spin quantisation axis of the neutrons, which is conveniently expressed by Equation 6.7. This includes a component which is constant, and a component with a 24-hour modulation, the relative magnitudes of which depend on the direction of the magnetic field. For the configuration at the PSI, both of these components are similar in size. This would produce three peaks in the power spectrum, the largest at the frequency corresponding to the axion mass m_a , and additional components at $m_a + \Omega_{\rm sid}$ and $|m_a - \Omega_{\rm sid}| = 0$ $\Omega_{\rm sid}$ For simplicity, it was chosen to neglect the portion of the signal with a sidereal modulation - especially given that there are many possible 24-hour effects which could give fake signals, such as daily temperature cycles. The sign of the modulation inverts when the magnetic field direction is inverted, but has no relation with the electric field. As such, the dataset was split according to magnetic field direction. The two datasets were analysed in the same way as the oscillating-EDM datasets. Two overlapping 3σ excesses at 3.42969 μ Hz and 3.32568 μ Hz were found, however these two signals do not have the correct phase relation to be consistent with an axion signal.

6.5 Conclusion

In this chapter the first laboratory limits on a coupling of dark-matter axions to gluons were presented. These limits were obtained by constraining the size of a possible oscillating electric dipole moment of the neutron. Combining the results from the ILL and PSI experimental analysis, large regions of the parameter space can be experimentally probed. These limits are up to 1000 times stronger than previous astrophysical limits for cosmologically interesting 10^{-22} eV axions. Additionally, the axion-wind coupling between axions and neutrons was probed, strengthening laboratory bounds by up to a factor 40. The mass range probed is complementary to the ranges accessible by 'on-resonance' experiments. The published limits are presented in Figure 6.7.



Figure 6.7: Final exclusion regions obtained for the axion-gluon coupling (above) and axion-neutron wind type coupling below. The axion-gluon result represents the first laboratory bound on this quantity, improving by up to a factor 1000 for cosmologically interesting 10^{-22} eV axions, while the axion-nucleon result represents an improvement by up to a factor 40 over previous laboratory limits over several orders of magnitudes of mass, but is still less restrictive than previous astrophysical limits.

Chapter 7

Closing Words

The PSI nEDM experiment collected enough data in 2015-16 to become the most sensitive neutron EDM measurement to date. Analysis of the dataset and consideration of the systematic corrections is well on the way to completion, with publication of a new world limit planned in 2019. A great many experiments around the world are in competition to build the next-generation of nEDM experiments, including the n2EDM experiment by the same collaboration. Construction of n2EDM has already begun, with the building of a large magnetically shielded room in the PSI UCN experimental area. The apparatus is scheduled to start taking useful nEDM data in 2021. A final sensitivity of the order 10^{-27} e cm is envisaged; a measurement on that scale which would put severe strain on several popular theories of new physics.

The search for axion-like dark matter has set the first experimental limits on the axiongluon coupling, improving in sensitivity over previous astrophysical across several orders of magnitude in axion mass. However, even across the limited mass range probed, the sensitivity is not yet sufficient to challenge some of the best-motivated regions. New, dedicated, experiments are proposed to extend the range of masses probed, including the CASPEr experiment, which promises to eventually reach the sensitivity to probe the best-motivated models.

Bibliography

- M. Thomson. Modern Particle Physics 1st ed. (Cambridge University Press, Cambridge, 2013).
- C. S. Wu, E. Ambler, R. W. Hayward, D. D. Hoppes & R. P. Hudson. Experimental Test of Parity Conservation in Beta Decay. *Physical Review* 105, 1413–1415 (1957).
- E. Fermi. Versuch einer Theorie der Beta-Strahlen. I. Zeitschrift f
 ür Physik 88, 161–177 (1934).
- T. D. Lee & C. N. Yang. Question of Parity Conservation in Weak Interactions. *Physical Review* 104, 254–258 (1956).
- R. P. Feynman & M. Gell-Mann. Theory of the Fermi Interaction. *Physical Review* 109, 193–198 (1958).
- E. C. G. Sudarshan & R. E. Marshak. Chirality Invariance and the Universal Fermi Interaction. *Physical Review* 109, 1860–1862 (1958).
- T. W. Kibble. History of electroweak symmetry breaking. Journal of Physics: Conference Series 626, 012001 (2015).
- N. Cabibbo. Unitary Symmetry and Leptonic Decays. *Physical Review Letters* 10, 531–533 (1963).
- S. L. Glashow, J. Iliopoulos & L. Maiani. Weak Interactions with Lepton-Hadron Symmetry. *Physical Review D* 2, 1285–1292 (1970).
- J. H. Christenson, J. W. Cronin, V. L. Fitch & R. Turlay. Evidence for the 2π Decay of the K⁰₂ Meson. *Physical Review Letters* 13, 138–140 (1964).
- H. Burkhardt *et al.* First evidence for direct CP violation. *Physics Letters B* 206, 169–176 (1988).
- M. Kobayashi & T. Maskawa. CP-Violation in the Renormalizable Theory of Weak Interaction. Progress of Theoretical Physics 49, 652–657 (1973).

- B. Aubert *et al.* Observation of CP Violation in the B0 Meson System. *Physical Review Letters* 87, 091801 (2001).
- K. Abe *et al.* Observation of Large CP Violation in the Neutral B Meson System. *Physical Review Letters* 87, 091802 (2001).
- R. Aaij et al. Search for CP violation in the phase space of D⁰ → π+π-π+π- decays. Physics Letters B 769, 345–356 (2017).
- K. Abe *et al.* Search for CP Violation in Neutrino and Antineutrino Oscillations by the T2K Experiment with 2.2 × 10²¹ Protons on Target. *Physical Review Letters* 121, 171802 (2018).
- 17. M. Tanabashi et al. Review of Particle Physics 2018. arXiv: 0601168 [astro-ph].
- R. D. Peccei. The strong CP problem and axions. Lecture Notes in Physics 741, 3–17 (2008).
- 19. P. Sikivie. The strong CP problem. Comptes Rendus Physique 13, 176–179 (2012).
- C. A. Baker *et al.* Improved Experimental Limit on the Electric Dipole Moment of the Neutron. *Physical Review Letters* 97, 131801 (2006).
- J. M. Pendlebury *et al.* A Revised Experimental Upper Limit on the Electric Dipole Moment of the Neutron. *Physical Review D* 092003, 1–23 (2015).
- 22. B. Graner, Y. Chen, E. G. Lindahl & B. R. Heckel. Reduced Limit on the Permanent Electric Dipole Moment of Hg 199. *Physical Review Letters* **116**, 161601 (2016).
- R. D. Peccei & H. R. Quinn. CP conservation in the presence of pseudoparticles. *Physical Review Letters* 38, 1440–1443 (1977).
- R. D. Peccei & H. R. Quinn. Constraints imposed by CP conservation in the presence of pseudoparticles. *Physical Review D* 16, 1791–1797 (1977).
- 25. S. Weinberg. A new light boson? Physical Review Letters 40, 223–226 (1978).
- F. Wilczek. Problem of strong P and T invariance in the presence of instantons. *Physical Review Letters* 40, 279–282 (1978).
- J. E. Kim. Weak-Interaction Singlet and Strong CP Invariance. *Physical Review Letters* 43, 103–107 (1979).
- 28. P. A. R. Ade et al. Plank 2015 results. Astronomy & Astrophysics 594, A13 (2016).
- A. D. Sakharov. Violation of CP in variance, C asymmetry, and baryon asymmetry of the universe. *Soviet Physics Uspekhi* 34, 392–393 (1991).

- D. E. Morrissey & M. J. Ramsey-Musolf. Electroweak baryogenesis. New Journal of Physics 14, 125003 (2012).
- A. Bochkarev & M. Shaposhnikov. Electroweak Production of Baryon Asymmetry and Upper Bounds on the Higgs and Top Masses. *Modern Physics Letters A* 02, 417–427 (1987).
- K. Kajantie, M. Laine, K. Rummukainen & M. Shaposhnikov. The electroweak phase transition: a non-perturbative analysis. *Nuclear Physics B* 466, 189–258 (1996).
- J. A. Heister et al. Search for the standard model higgs boson at LEP. Physics Letters, Section B: Nuclear, Elementary Particle and High-Energy Physics 565, 61-75 (2003).
- 34. G. Aad *et al.* Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC. *Physics Letters B* **716**, 1–29 (2012).
- 35. S. Chatrchyan *et al.* Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC. *Physics Letters B* **716**, 30–61 (2012).
- D. McKeen, M. Pospelov & A. Ritz. Modified Higgs branching ratios versus CP and lepton flavor violation. *Physical Review D* 86, 1–8 (2012).
- T. Cohen, D. E. Morrissey & A. Pierce. Electroweak baryogenesis and Higgs signatures. *Physical Review D* 86, 013009 (2012).
- C.-Y. Chen, S. Dawson & Y. Zhang. Complementarity of LHC and EDMs for exploring Higgs CP violation. *Journal of High Energy Physics* 2015, 56 (2015).
- A. M. Baldini *et al.* Search for the lepton flavour violating decay μ+→ e+γ with the full dataset of the MEG experiment: MEG Collaboration. *The European Physical Journal C* 76, 434 (2016).
- G. Pezzullo. The Mu2e experiment at Fermilab: a search for lepton flavor violation. Nuclear and Particle Physics Proceedings 285-286, 3-7 (2017).
- P. Huet & E. Sather. Electroweak baryogenesis and standard model CP violation. *Physical Review D* 51, 379–394 (1995).
- K. Blum, C. Delaunay, M. Losada, Y. Nir & S. Tulin. CP violation beyond the MSSM: baryogenesis and electric dipole moments. *Journal of High Energy Physics* 2010, 101 (2010).

- S. Huber, T. Konstandin, T. Prokopec & M. Schmidt. Baryogenesis in the MSSM, nMSSM and NMSSM. *Nuclear Physics A* 785, 206–209 (2007).
- 44. S. J. Huber, T. Konstandin, T. Prokopec & M. G. Schmidt. Electroweak phase transition and baryogenesis in the nMSSM. *Nuclear Physics B* **757**, 172–196 (2006).
- J. Kozaczuk, S. Profumo, M. J. Ramsey-Musolf & C. L. Wainwright. Supersymmetric electroweak baryogenesis via resonant sfermion sources. *Physical Review D* 86, 096001 (2012).
- W. Buchmüller, P. Di Bari & M. Plümacher. Leptogenesis for pedestrians. Annals of Physics 315, 305–351 (2005).
- R. Barbieri, P. Creminelli, A. Strumia & N. Tetradis. Baryogenesis through leptogenesis. *Nuclear Physics B* 575, 61–77 (2000).
- E. Molinaro. Leptogenesis: Recent Developments. Nuclear and Particle Physics Proceedings 265-266, 180–182 (2015).
- F. R. Joaquim, I. Masina & A. Riotto. Observable Electron EDM and Leptogenesis. International Journal of Modern Physics A 22, 6253–6278 (2007).
- 50. C. Smith & S. Touati. EDM with and beyond flavor invariants. arXiv: 1707.06805 (2017).
- 51. G. Greene *et al.* A new measurement of the magnetic moment of the neutron. *Physics Letters B* **71**, 297–300 (1977).
- G. L. Greene *et al.* Measurement of the neutron magnetic moment. *Physical Review* D 20, 2139–2153 (1979).
- G. L. Greene *et al.* An Improved Derived Value for the Neutron Magnetic Moment in Nuclear Magnetons. *Metrologia* 18, 93 (1982).
- F. J. M. Farley *et al.* New Method of Measuring Electric Dipole Moments in Storage Rings. *Physical Review Letters* 93, 052001 (2004).
- 55. V. Anastassopoulos *et al.* A storage ring experiment to detect a proton electric dipole moment. *Review of Scientific Instruments* **87**, 115116 (2016).
- R. Chislett. The muon EDM in the g-2 experiment at Fermilab. EPJ Web of Conferences 118 (eds G. D'Ambrosio et al.) 01005 (2016).
- G. W. Bennett *et al.* Improved limit on the muon electric dipole moment. *Physical Review D* 80, 052008 (2009).

- P. G. H. Sandars. Measurability of the Proton Electric Dipole Moment. *Physical Review Letters* 19, 1396–1398 (1967).
- L. I. Schiff. Measurability of Nuclear Electric Dipole Moments. *Physical Review* 132, 2194–2200 (1963).
- E. M. Purcell & N. F. Ramsey. On the Possibility of Electric Dipole Moments for Elementary Particles and Nuclei. *Physical Review* 78, 807–807 (1950).
- 61. C.-P. Liu & J. Engel. Schiff screening of relativistic nucleon electric-dipole moments by electrons. *Physical Review C* **76**, 028501 (2007).
- V. V. Flambaum & A. Kozlov. Screening and finite-size corrections to the octupole and Schiff moments. *Physical Review C* 85, 068502 (2012).
- R. H. Parker *et al.* First Measurement of the Atomic Electric Dipole Moment of ²²⁵Ra. *Physical Review Letters* **114**, 233002 (2015).
- N. Auerbach, V. V. Flambaum & V. Spevak. Collective T and P -Odd Electromagnetic Moments in Nuclei with Octupole Deformations. *Physical Review Letters* 76, 4316–4319 (1996).
- 65. V. A. Dzuba, V. V. Flambaum, J. S. M. Ginges & M. G. Kozlov. Electric dipole moments of Hg, Xe, Rn, Ra, Pu, and TlF induced by the nuclear Schiff moment and limits on time-reversal violating interactions. *Physical Review A* 66, 012111 (2002).
- J. Dobaczewski & J. Engel. Nuclear Time-Reversal Violation and the Schiff Moment of ²²⁵Ra. *Physical Review Letters* 94, 232502 (2005).
- V. Spevak, N. Auerbach & V. V. Flambaum. Enhanced T,P-odd electromagnetic moments in reflection asymmetric nuclei. *Physical Review C* 56, 1357–1369 (1997).
- J. Engel, M. J. Ramsey-Musolf & U. van Kolck. Electric dipole moments of nucleons, nuclei, and atoms: The Standard Model and beyond. *Progress in Particle and Nuclear Physics* **71**, 21–74 (2013).
- V. Andreev et al. Improved limit on the electric dipole moment of the electron. Nature 562, 355–360 (2018).
- J. J. Hudson *et al.* Improved measurement of the shape of the electron. *Nature* 473, 493–496 (2011).
- P. Sandars. The electric dipole moment of an atom. *Physics Letters* 14, 194–196 (1965).

- I. B. Khriplovich & S. K. Lamoreaux. CP Violation Without Strangeness: The Electric Dipole Moments of Particles (Springer, 2004).
- T. Chupp & M. Ramsey-Musolf. Electric dipole moments: A global analysis. *Physical Review C* 91, 035502 (2015).
- L. Bian, T. Liu & J. Shu. Cancellations Between Two-Loop Contributions to the Electron Electric Dipole Moment with a CP-Violating Higgs Sector. *Physical Review Letters* 115, 021801 (2015).
- 75. G. W. Bennett *et al.* Final report of the E821 muon anomalous magnetic moment measurement at BNL (2006).
- A. Crivellin, M. Hoferichter & P. Schmidt-Wellenburg. Combined explanations of (g-2)_{mu,e} and implications for a large muon EDM. arXiv: 1807.11484 (2018).
- 77. D. Polder. On the dipole moment of ammonia. Physica 9, 908–914 (1942).
- E. Shabalin. The Electric Dipole Moments Of Baryons In The Kobayashi-Maskawa CP Noninvariant Theory. Sov.J.Nucl.Phys. 32, 443–447 (1979).
- M. Pospelov & A. Ritz. Electric dipole moments as probes of new physics. Annals of Physics 318, 119–169 (2005).
- C.-Y. Seng. Reexamination of the standard model nucleon electric dipole moment. *Physical Review C* 91, 025502 (2015).
- J. H. Smith, E. M. Purcell & N. F. Ramsey. Experimental Limit to the Electric Dipole Moment of the Neutron. *Physical Review* 108, 120–122 (1957).
- N. F. Ramsey. A Molecular Beam Resonance Method with Separated Oscillating Fields. *Physical Review* 78, 695–699 (1950).
- W. B. Dress, P. D. Miller, J. M. Pendlebury, P. Perrin & N. F. Ramsey. Search for an electric dipole moment of the neutron. *Physical Review D* 15, 9–21 (1977).
- 84. D. Griffiths. Introduction to Electrodynamics (Prentice-Hall, 1999).
- I. S. Altarev *et al.* A search for the electric dipole moment of the neutron using ultracold neutrons. *Nuclear Physics, Section A* 341, 269–283 (1980).
- S. K. Lamoreaux & R. Golub. Comment on "Improved Experimental Limit on the Electric Dipole Moment of the Neutron". *Physical Review Letters* 98, 149101 (2007).
- C. A. Baker *et al.* Reply: Comment on "Improved Experimental Limit on the Electric Dipole Moment of the Neutron". *Physical Review Letters* **98**, 149102 (2007).

- C. A. Baker *et al.* Apparatus for measurement of the electric dipole moment of the neutron using a cohabiting atomic-mercury magnetometer. *Nuclear Instruments and Methods in Physics Research Section A* 736, 184–203 (2014).
- C. Baker *et al.* The search for the neutron electric dipole moment at the Paul Scherrer Institute. *Physics Proceedia* 17, 159–167 (2011).
- 90. A. Anghel et al. The PSI ultra-cold neutron source. Nuclear Instruments and Methods in Physics Research, Section A: Accelerators, Spectrometers, Detectors and Associated Equipment 611, 272–275 (2009).
- C. Abel *et al.* nEDM experiment at PSI: data-taking strategy and sensitivity of the dataset. arXiv: 1811.04012v1 (2018).
- 92. C. Abel *et al.* The n2EDM experiment at the Paul Scherrer Institute. arXiv: 1811.02340 (2018).
- 93. N. J. Ayres et al. Monte Carlo simulations for the optimization and data analysis of experiments with ultracold neutrons in Proceedings of the International Conference on Neutron Optics (NOP2017) (Journal of the Physical Society of Japan, 2018). arXiv: 1806.10778.
- 94. R. Picker. How the Minuscule Can Contribute to the Big Picture: The Neutron Electric Dipole Moment Project at TRIUMF in Proceedings of the 14th International Conference on Meson-Nucleon Physics and the Structure of the Nucleon (MENU2016) (2017). arXiv: 1612.00875.
- 95. T. M. Ito *et al.* Performance of the upgraded ultracold neutron source at Los Alamos National Laboratory and its implication for a possible neutron electric dipole moment experiment. *Physical Review C* 97, 12501 (2018).
- F. M. Piegsa *et al.* New source for ultracold neutrons at the Institut Laue-Langevin. *Physical Review C* 90 (2014).
- A. P. Serebrov et al. UCN source with superfluid helium at WWR-M reactor. J. Phys.: Conf. Ser. 798, 012147 (2017).
- A. P. Serebrov *et al.* New search for the neutron electric dipole moment with ultracold neutrons at ILL. *Physical Review C* 92 (2015).
- 99. V. V. Fedorov *et al.* Measurement of the neutron electric dipole moment via spin rotation in a non-centrosymmetric crystal. *Physics Letters B* **694**, 22–25 (2010).

- V. V. Fedorov, V. V. Voronin & Y. P. Braginetz. Search for the neutron EDM by crystal-diffraction method. Test experiment and future progress. *Physica B: Condensed Matter* 406, 2370–2372 (2011).
- C. A. Baker *et al.* CryoEDM: a cryogenic experiment to measure the neutron Electric Dipole Moment. *Journal of Physics: Conference Series* 251, 012055 (2010).
- 102. N. Fomin et al. Fundamental neutron physics beamline at the spallation neutron source at ORNL. Nuclear Inst. and Methods in Physics Research, A 773, 45–51 (2014).
- 103. T. M. Ito & the nEdm Collaboration. Plans for a Neutron EDM Experiment at SNS. Journal of Physics: Conference Series 69, 012037 (2007).
- R. Golub & S. K. Lamoreaux. Neutron electric-dipole moment, ultracold neutrons and polarized 3He. *Physics Reports* 237, 1–62 (1994).
- 105. F. M. Piegsa. New concept for a neutron electric dipole moment search using a pulsed beam. *Physical Review C* 88, 045502 (2013).
- 106. R. Hall-Wilton. Status of the European Spallation Source ESS AB, the Instrument Selection Process, and a Fundamental Physics Beamline at the ESS. *Physics Procedia* 51, 8–12 (2014).
- R. Golub, D. Richardson & S. Lamoreaux. Ultra-Cold Neutrons 1st ed. (IOP Publishing, 1991).
- B. Fornal & B. Grinstein. Dark Matter Interpretation of the Neutron Decay Anomaly. *Physical Review Letters* 120, 191801 (2018).
- 109. Z. Berezhiani. Neutron lifetime puzzle and neutron mirror neutron oscillation. arXiv: 1807.07906.
- G. Ban et al. Direct Experimental Limit on Neutron-Mirror-Neutron Oscillations. *Physical Review Letters* 99, 161603 (2007).
- 111. P. Mohan Murthy. In Preparation PhD thesis (ETH Zürich, 2018).
- 112. J. S. Nico. Neutron beta decay. Journal of Physics G: Nuclear and Particle Physics 36, 104001 (2009).
- 113. A. Serebrov et al. Superconducting UCN Polarizer for a New EDM Spectrometer. Journal of research of the National Institute of Standards and Technology 110, 185– 8 (2005).

- J. M. Pendlebury & D. J. Richardson. Effects of gravity on the storage of ultracold neutrons tech. rep. (1994), 504–511.
- 115. P. G. Harris, J. M. Pendlebury & N. E. Devenish. Gravitationally enhanced depolarization of ultracold neutrons in magnetic-field gradients. *Physical Review D* 89, 016011 (2014).
- S. Afach *et al.* Gravitational depolarization of ultracold neutrons: Comparison with data. *Physical Review D* 92, 052008 (2015).
- 117. S. Afach *et al.* Observation of Gravitationally Induced Vertical Striation of Polarized Ultracold Neutrons by Spin-Echo Spectroscopy. *Physical Review Letters* 162502, 1–6 (2015).
- D. Roulier et al. Status of the GRANIT Facility. Advances in High Energy Physics 2015, 730437 (2015).
- 119. H. Abele *et al.* QBounce, the quantum bouncing ball experiment. *Physics Procedia* 17, 4–9 (2011).
- I. I. Rabi, J. R. Zacharias, S. Millman & P. Kusch. A New Method of Measuring Nuclear Magnetic Moment. *Physical Review* 53, 318–318 (1938).
- 121. P. Harris. The Neutron EDM Experiment. arXiv: 0709.3100v3 (2007).
- 122. M. Rawlik. Tuning Neutron RF-Pulse Frequency in the nEDM Experiment (Jagiellonian University, 2012).
- 123. A. Steyerl et al. A new source of cold and ultracold neutrons. Physics Letters A 116, 347–352 (1986).
- 124. G. Bison et al. An ultracold neutron storage bottle for UCN density measurements. Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment 830, 449–453 (2016).
- 125. G. Bison *et al.* Comparison of ultracold neutron sources for fundamental physics measurements. *Physical Review C* 95, 45503 (2017).
- 126. M. Kuzniak. The Neutron Electric Dipole Moment Experiment: Research and Development for the New Spectrometer PhD thesis (Jagiellonian University, 2008).
- 127. E. Wursten. In Preparation PhD thesis (KU Leuven, 2018).
- 128. D. Anicic et al. A fast kicker magnet for the PSI 600 MeV proton beam to the PSI ultra-cold neutron source. Nuclear Instruments and Methods in Physics Research A 541, 598–609 (2005).

- 129. M. Bopp, H. Fitze, P. Sigg & L. Stingelin. Upgrade concepts of the PSI accelerator RF systems for a projected 3 mA operation in AIP Conference Proceedings 600 (AIP, 2001), 300–302.
- M. Wohlmuther & G. Heidenreich. The spallation target of the ultra-cold neutron source UCN at PSI. Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment 564, 51–56 (2006).
- 131. H. Becker et al. Neutron production and thermal moderation at the PSI UCN source. Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment 777, 20–27 (2015).
- R. Golub & K. Boning. Condensed New Type of Low Temperature Source of Ultra-Cold Neutrons and Production of Continous Beams of UCN tech. rep. (1983), 95– 98.
- 133. M. Daum et al. Transmission of ultra-cold neutrons through guides coated with materials of high optical potential. Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment 741, 71–77 (2014).
- 134. K. Bodek et al. Storage of ultracold neutrons in high resistivity, non-magnetic materials with high Fermi potential. Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment 597, 222–226 (2008).
- 135. F. Atchison et al. Measurement of the Fermi potential of diamond-like carbon and other materials. Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms 260, 647–656 (2007).
- S. Afach *et al.* A device for simultaneous spin analysis of ultracold neutrons. *The European Physical Journal A* 51, 143 (2015).
- G. Ban et al. Ultracold neutron detection with ⁶Li-doped glass scintillators. The European Physical Journal A 52, 326 (2016).
- 138. A. Serebrov, A. Aldushchenkov, M. Lasakov, I. Kuznetsov & I. Stepanenko. New method for precise determination of neutron beam polarization. Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment 357, 503–510 (1995).

- 139. S. Afach *et al.* Dynamic stabilization of the magnetic field surrounding the neutron electric dipole moment spectrometer at the Paul Scherrer Institute. *Journal of Applied Physics* **116**, 084510 (2014).
- 140. P. G. Harris & J. M. Pendlebury. Dipole-field contributions to geometric-phaseinduced false electric-dipole-moment signals for particles in traps. *Physical Review* A 73, 014101 (2006).
- 141. C. Abel *et al.* Magnetic field uniformity in neutron electric dipole moment experiments. arXiv: 1811.06085 (2018).
- 142. O. Bourrion, G. Pignol, D. Rebreyend & C. Vescovi. Development of a multifunction module for the neutron electric dipole moment experiment at PSI. Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment 701, 278–284 (2013).
- 143. K. Green et al. Performance of an atomic mercury magnetometer in the neutron EDM experiment. Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment 404, 381–393 (1998).
- 144. S. V. Komposch. Realization of a high-performance laser-based mercury magnetometer for neutron EDM experiments PhD thesis (ETH Zürich, 2017).
- 145. J. A. Thorne. *Electric Field Optimisation for Cryogenic nEDM Experiments* PhD thesis (2018).
- 146. M. C. Fertl. A laser based mercury co-magnetometer for the neutron electric dipole moment search PhD thesis (ETH Zürich, 2013).
- 147. G. Ban et al. Demonstration of sensitivity increase in mercury free-spin-precession magnetometers due to laser-based readout for neutron electric dipole moment searches. Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment 896, 129–138 (2018).
- 148. J. M. Pendlebury *et al.* Geometric-phase-induced false electric dipole moment signals for particles in traps. *Physical Review A* **70**, 032102 (2004).
- 149. S. K. Lamoreaux & R. Golub. Detailed discussion of a linear electric field frequency shift induced in confined gases by a magnetic field gradient: Implications for neutron electric-dipole-moment experiments. *Physical Review A - Atomic, Molecular, and Optical Physics* **71**, 1–10 (2005).

- 150. G. Pignol & S. Roccia. Electric-dipole-moment searches: Reexamination of frequency shifts for particles in traps. *Physical Review A* **85**, 042105 (2012).
- 151. S. Afach *et al.* Measurement of a false electric dipole moment signal from 199 Hg atoms exposed to an inhomogeneous magnetic field. *European Physical Journal D* 69. arXiv: arXiv:1503.08651v1 (2015).
- 152. G. Pignol, M. Guigue, A. Petukhov & R. Golub. Frequency shifts and relaxation rates for spin-1/2 particles moving in electromagnetic fields. *Physical Review A* 92, 053407 (2015).
- 153. S. Groeger, G. Bison, J.-L. Schenker, R. Wynands & A. Weis. A high-sensitivity laser-pumped Mx magnetometer. *European Physical Journal D* 38, 239–247 (2006).
- 154. S. Gröger. Laser-pumped cesium magnetometers for the PSI-nEDM experiment PhD thesis (University of Fribourg).
- 155. P. Knowles et al. Laser-driven Cs magnetometer arrays for magnetic field measurement and control. Nuclear Instruments and Methods in Physics Research, Section A: Accelerators, Spectrometers, Detectors and Associated Equipment 611, 306–309 (2009).
- 156. S. Groeger, G. Bison & A. Weis. Design and performance of laser-pumped Csmagnetometers for the planned UCN EDM experiment at PSI. Journal of Research of the National Institute of Standards and Technology 110, 179 (2005).
- 157. J. Zenner. The search for the neutron electric dipole moment PhD thesis (Johannes Gutenberg-Universität, Mainz, 2013).
- 158. C. Abel *et al.* Statistical sensitivity of the nEDM apparatus at PSI to neutron mirror-neutron oscillations. arXiv: 1811.01906 (2018).
- 159. S. Afach *et al.* Constraining interactions mediated by axion-like particles with ultracold neutrons. *Physics Letters B* **745**, 58–63 (2014).
- 160. I. Altarev *et al.* Testing isotropy of the universe using the Ramsey resonance technique on ultracold neutron spins. *Physica B: Condensed Matter* 406, 2365–2369 (2011).
- 161. C. Abel *et al.* Search for Axionlike Dark Matter through Nuclear Spin Precession in Electric and Magnetic Fields. *Physical Review X* 7, 041034 (2017).
- 162. The PSI n2EDM Collaboration. n2EDM Project Overview Report tech. rep. (2018).

- 163. M. Rawlik. *FID signal analysis and new DAQ system in the nEDM experiment* (Jagiellonian University, 2014).
- 164. The Git Distributed Version Control System. Git
- 165. MathWorks Inc. MATLAB R2015a
- 166. The Python Software Foundation. The Python Programming Langauge
- 167. T. Oliphant. A guide to NumPy (Trelgol Publishing, 2006).
- E. Jones, T. Oliphant, P. Peterson et al. SciPy: Open source scientific tools for Python
- 169. W. Mckinney. Data Structures for Statistical Computing in Python in Proceedings of the 9th Python in Science Conference (Scipy 2010) (2010), 51.
- 170. The HDF Group. The HDF5 Library and File Format
- 171. LPC Caen. Fast Acquisition System for nuclEar Research (FASTER)
- 172. Y. Chibane, S. K. Lamoreaux, J. M. Pendlebury & K. F. Smith. Minimum variance of frequency estimations for a sinusoidal signal with low noise. *Measurement Science* and Technology 6, 1671–1678 (1995).
- 173. V. Hélaine. Mesure du moment dipolaire électrique du neutron: analyse simultanée de spin et analyse préliminaire des données PhD thesis (Université de Caen Basse-Normandie, 2014).
- S. Afach *et al.* A measurement of the neutron to 199Hg magnetic moment ratio. *Physics Letters B* 739, 128–132 (2014).
- C. A. Baker *et al.* Reply: Comment on "Improved Experimental Limit on the Electric Dipole Moment of the Neutron". *Physical Review Letters* 98, 149102 (2007).
- 176. A. Y. Silenko. Taking into account the Earth's rotation in experiments on the search of the electric dipole moment of the neutron. *Physics of Particles and Nuclei Letters* 4, 468–470 (2007).
- 177. J. R. Piller. Accurate optical magnetometry based on precessing spin-aligned atoms PhD thesis (University of Fribourg, 2016).
- 178. G. Wyszynski. Development of mangetic field control systems in the nEDM Experiment PhD thesis (Jagiellonian University, 2016).
- 179. S. Afach *et al.* Highly stable atomic vector magnetometer based on free spin precession. *Optics Express* **23**, 22108 (2015).

- S. Afach. Development of a cesium vector magnetometer for the neutron EDM experinemt PhD thesis (ETH Zürich, 2014).
- 181. D. W. Allan. Statistics of Atomic Frequency Standards tech. rep. 2 (1966).
- 182. IEEE Standards Coordinating Committee. IEEE standard definitions of physical quantities for fundamental frequency and time metrology-random instabilities (IEEE standard 1139-2008) (2008).
- 183. J. A. Barnes. Tables of bias functions, B1 and B2, for variances based on finite samples of processes with power law spectral densities - Volume 375 of NBS technical note (U.S. Dept. of Commerce, National Bureau of Standards, 1969).
- E. Ziegel, W. Press, B. Flannery, S. Teukolsky & W. Vetterling. Numerical Recipes 3rd Edition: The Art of Scientific Computing arXiv: arXiv:1011.1669v3 (2007).
- B. Efron & R. Tibshirani. An Introduction to the Bootstrap (Chapman & Hall/CRC, 1993).
- B. Efron. Bootstrap Methods: Another Look at the Jackknife. The Annals of Statistics 7, 1–26 (1979).
- 187. D. J. E. Marsh. Axion Cosmology. Physics Reports 643, 87 (2015).
- 188. P. W. Graham & S. Rajendran. Axion dark matter detection with cold molecules. *Physical Review D* 84, 1–13 (2011).
- 189. P. W. Graham & S. Rajendran. New observables for direct detection of axion dark matter. *Physical Review D* 88, 1–13 (2013).
- 190. Y. V. Stadnik & V. V. Flambaum. Axion-induced effects in atoms, molecules, and nuclei: Parity nonconservation, anapole moments, electric dipole moments, and spingravity and spin-axion momentum couplings. *Physical Review D* 89, 043522 (2014).
- P. W. Graham *et al.* Spin precession experiments for light axionic dark matter. *Physical Review D* 97, 055006 (2018).
- 192. P. W. Graham, D. E. Kaplan, J. Mardon, S. Rajendran & W. A. Terrano. Dark matter direct detection with accelerometers. *Physical Review D* 93, 1–22 (2016).
- 193. R. Crewther, P. Di Vecchia, G. Veneziano & E. Witten. Chiral estimate of the electric dipole moment of the neutron in quantum chromodynamics. *Physics Letters* B 88, 123–127 (1979).
- 194. J. De Vries, E. Mereghetti & A. Walker-Loud. Baryon mass splittings and strong CP violation in SU(3) chiral perturbation theory. *Physical Review C* 92, 45201 (2015).

- 195. M. Pospelov & A. Ritz. Theta-Induced Electric Dipole Moment of the Neutron via QCD Sum Rules. *Physical Review Letters* 83, 2526–2529 (1999).
- 196. M. Shifman, A. Vainshtein & V. Zakharov. Can confinement ensure natural CP invariance of strong interactions? *Nuclear Physics B* 166, 493–506 (1980).
- 197. M. Dine, W. Fischler & M. Srednicki. A simple solution to the strong CP problem with a harmless axion. *Physics Letters B* 104, 199–202 (1981).
- A. Zhitnitsky. On Possible Suppression of the Axion Hadron Interactions. (In Russian). Soviet Journal of Nuclear Physics 31, 497–504 (1980).
- F. Zwicky. Die Rotverschiebung von extragalaktischen Nebeln. Helvetica Physica Acta 6, 110–127 (1933).
- 200. L. D. Duffy & K. Van Bibber. Axions as dark matter particles. New Journal of Physics 11, 105008 (2009).
- 201. J. E. Kim & G. Carosi. Axions and the strong CP problem. Reviews of Modern Physics 82, 557–601 (2010).
- 202. H.-Y. Schive, T. Chiueh & T. Broadhurst. Cosmic structure as the quantum interference of a coherent dark wave. *Nature Physics* 10, 496–499 (2014).
- 203. J. Preskill, M. B. Wise & F. Wilczek. Cosmology of the invisible axion. *Physics Letters B* 120, 127–132 (1983).
- 204. L. Abbott & P. Sikivie. A cosmological bound on the invisible axion. *Physics Letters B* 120, 133–136 (1983).
- 205. M. Dine & W. Fischler. The not-so-harmless axion. *Physics Letters B* 120, 137–141 (1983).
- 206. R. Hlozek, D. Grin, D. J. E. Marsh & P. G. Ferreira. A search for ultralight axions using precision cosmological data. *Physical Review D* 91, 103512 (2015).
- 207. V. A. Kostelecký & C. D. Lane. Constraints on Lorentz violation from clockcomparison experiments. *Physical Review D* 60, 116010 (1999).
- 208. J. Redondo & A. Ringwald. Light shining through walls. Contemporary Physics ISSN: 52, 211–236 (2011).
- V. Anastassopoulos *et al.* New CAST limit on the axion–photon interaction. *Nature Physics* 13, 584–590 (2017).
- 210. S. J. Asztalos *et al.* Improved rf cavity search for halo axions. *Physical Review D*69. arXiv: 0310042 [astro-ph] (2004).

- 211. S. Asztalos et al. Design and performance of the ADMX SQUID-based microwave receiver. Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment 656, 39–44 (2011).
- J. Hoskins *et al.* Modulation sensitive search for nonvirialized dark-matter axions. *Physical Review D* 94, 082001 (2016).
- N. Du *et al.* Search for Invisible Axion Dark Matter with the Axion Dark Matter Experiment. *Physical Review Letters* 120 (2018).
- 214. D. Budker, P. W. Graham, M. Ledbetter, S. Rajendran & A. O. Sushkov. Proposal for a cosmic axion spin precession experiment (CASPEr). *Physical Review X* 4, 1–10 (2014).
- 215. A. Garcon *et al.* The cosmic axion spin precession experiment (CASPEr): A dark-matter search with nuclear magnetic resonance. *Quantum Science and Technology* 3, 014008 (2018).
- 216. A. Schuster. On the investigation of hidden periodicities with application to a supposed 26 day period of meteorological phenomena. *Journal of Geophysical Research* 3, 13 (1898).
- J. T. VanderPlas. Understanding the Lomb–Scargle Periodogram. The Astrophysical Journal Supplement Series 236, 16 (2018).
- J. D. Scargle. Studies in astronomical time series analysis. II Statistical aspects of spectral analysis of unevenly spaced data. *The Astrophysical Journal* 263, 835 (1982).
- N. R. Lomb. Least-squares frequency analysis of unequally spaced data. Astrophysics and Space Science 39, 447–462 (1976).
- 220. A. L. Read. Presentation of search results: the CL s technique tech. rep. (2002), 2693–2704.
- P. Schmidt-Wellenburg. The quest for an electric dipole moment of the neutron. AIP Conference Proceedings 1753, 060002 (2016).

Appendix A

The Standardised Adequate Basis of Polynomials

Within the precession chamber of the experiment, the magnetic field is nominally $\pm 1 \ \mu T$ in the vertical direction. This field is exceptionally uniform during nEDM running, however the characterisation of small inhomogeneities and gradients is crucial to the computation of systematic effects in the measurement. To this end, a particular basis for the expansion of the magnetic field was constructed. While the choice of basis is in principle arbitrary, the one presented here is particularly convenient for our purposes. This formulation was proposed in [91]

The field is expressed as a series of 'gradients' $G_{l,m}$, where l is often referred to as the order. Gradients with m = 0 are cylindrically symmetrical, and include the 'linear vertical gradient' $G_{1,0}$ and the 'cubic vertical gradient' $G_{3,0}$, which are responsible for the false EDM effect. The magnetic field can be reconstructed from these gradients in the form

$$\vec{B}(\vec{r}) = \sum_{l,m} G_{l,m} \begin{pmatrix} \Pi_{x,l,m}(\vec{r}) \\ \Pi_{y,l,m}(\vec{r}) \\ \Pi_{z,l,m}(\vec{r}) \end{pmatrix}$$
(A.1)

where the functions (or *modes*) $\vec{\Pi}_{l,m}$ are harmonic polynomials in x, y, z of degree l and $G_{l,m}$ are the expansion coefficients.

Modes of arbitrarily high order can be trivially computed using a computer algebra program using the procedure and formulae from [91]. For convenience, the modes up to order 3 are presented in Cartesian and cylindrical coordinates in Table A.1 and Table A.2. Modes up to order 6 are used in the analysis of field map data in chapter 5, however these expressions are too large to be conveniently printed.

l	m	Π_x	Π_y	Π_z
0	-1	0	1	0
0	0	0	0	1
0	1	1	0	0
1	-2	y	x	0
1	-1	0	z	y
1	0	$-\frac{1}{2}x$	$-rac{1}{2}y$	z
1	1	z	0	x
1	2	x	-y	0
2	-3	2xy	$x^2 - y^2$	0
2	-2	2yz	2xz	2xy
2	-1	$-\frac{1}{2}xy$	$-\frac{1}{4}(x^2 + 3y^2 - 4z^2)$	2yz
2	0	-xz	-yz	$z^2 - \frac{1}{2}(x^2 + y^2)$
2	1	$-\frac{1}{4}(3x^2 + y^2 - 4z^2)$	$-\frac{1}{2}xy$	2xz
2	2	2xz	-2yz	$x^2 - y^2$
2	3	$x^2 - y^2$	-2xy	0
3	-4	$3x^2y - y^3$	$x^3 - 3xy^2$	0
3	-3	6xyz	$3(x^2z - y^2z)$	$3x^2y - y^3$
3	-2	$-\frac{1}{2}(3x^2y + y^3 - 6yz^2)$	$-\frac{1}{2}(x^3 + 3xy^2 - 6xz^2)$	6xyz
3	-1	$-\frac{3}{2}xyz$	$-\frac{1}{4}(3x^2z + 9y^2z - 4z^3)$	$3yz^2 - \frac{3}{4}(x^2y + y^3)$
3	0	$\frac{3}{8}(x^3 + xy^2 - 4xz^2)$	$\frac{3}{8}(x^2y + y^3 - 4yz^2)$	$z^3 - \frac{3}{2}z(x^2 + y^2)$
3	1	$-\frac{1}{4}(9x^2z + 3y^2z - 4z^3)$	$-\frac{3}{2}xyz$	$3xz^2 - \frac{3}{4}(x^3 + xy^2)$
3	2	$-x^3 + 3xz^2$	$-3yz^2 + y^3$	$3(x^2z - y^2z)$
3	3	$3(x^2z - y^2z)$	-6xyz	$x^3 - 3xy^2$
3	4	$x^3 - 3xy^2$	$-3x^2y + y^3$	0

Table A.1: The basis of harmonic polynomials sorted by degree.

l	m	$\Pi_{ ho}$	Π_{ϕ}	Π_z
0	-1	$\sin\phi$	$\cos\phi$	0
0	0	0	0	1
0	1	$\cos\phi$	$-\sin\phi$	0
1	-2	$ ho\sin 2\phi$	$ ho\cos 2\phi$	0
1	-1	$z\sin\phi$	$z\cos\phi$	$ ho\sin\phi$
1	0	$-\frac{1}{2} ho$	0	z
1	1	$z\cos\phi$	$-z\sin\phi$	$ ho\cos\phi$
1	2	$ ho\cos 2\phi$	$-\rho\sin 2\phi$	0
2	-3	$\rho^2 \sin 3\phi$	$ ho^2 \cos 3\phi$	0
2	-2	$2\rho z \sin 2\phi$	$2\rho z\cos 2\phi$	$ ho^2 \sin 2\phi$
2	-1	$\frac{1}{4}(4z^2 - 3\rho^2)\sin\phi$	$\frac{1}{4}(4z^2-\rho^2)\cos\phi$	$2\rho z\sin\phi$
2	0	- ho z	0	$-\tfrac{1}{2}\rho^2 + z^2$
2	1	$\tfrac{1}{4}(4z^2-3\rho^2)\cos\phi$	$\tfrac{1}{4}(\rho^2 - 4z^2)\sin\phi$	$2 ho z\cos\phi$
2	2	$2\rho z\cos 2\phi$	$-2\rho z\sin 2\phi$	$ ho^2\cos 2\phi$
2	3	$ ho^2 \cos 3\phi$	$- ho^2\sin 3\phi$	0
3	-4	$\rho^3 \sin 4\phi$	$ ho^3 \cos 4\phi$	0
3	-3	$3\rho^2 z \sin 3\phi$	$3 ho^2 z\cos 3\phi$	$ ho^3 \sin 3\phi$
3	-2	$\rho(3z^2-\rho^2)\sin 2\phi$	$\frac{1}{2}\rho(6z^2-\rho^2)\cos 2\phi$	$3\rho^2 z \sin 2\phi$
3	-1	$\frac{1}{4}z(4z^2-9\rho^2)\sin\phi$	$\frac{1}{4}z(4z^2-3\rho^2)\cos\phi$	$\rho(3z^2 - \frac{3}{4}\rho^2)\sin\phi$
3	0	$\frac{3}{8}\rho(\rho^2 - 4z^2)$	0	$\frac{1}{2}z(2z^2-3\rho^2)$
3	1	$\frac{1}{4}z(4z^2-9\rho^2)\cos\phi$	$\frac{1}{4}z(3\rho^2-4z^2)\sin\phi$	$\rho(3z^2 - \frac{3}{4}\rho^2)\cos\phi$
3	2	$\rho(3z^2 - \rho^2)\cos 2\phi$	$\frac{1}{2}\rho(\rho^2 - 6z^2)\sin 2\phi$	$3\rho^2 z \cos 2\phi$
3	3	$3 ho^2 z\cos 3\phi$	$-3\rho^2 z \sin 3\phi$	$ ho^3\cos 3\phi$
3	4	$ ho^3 \cos 4\phi$	$- ho^3\sin4\phi$	0

Table A.2: The basis of harmonic polynomials sorted by degree in cylindrical coordinates.