

A New Measurement of the Electron Electric Dipole Moment Using Ytterbium Fluoride

Ian Joseph Smallman

Thesis submitted in partial fulfilment of the requirements
for the degree of Doctor of Philosophy.

Department of Physics

Imperial College London

August 2013

Declaration

I declare that this thesis is my own work. Where I have used the work of others the sources are appropriately referenced and acknowledged.

The copyright on this thesis rests with the author and is made available under a Creative Commons Attribution Non-Commercial No Derivatives licence. Researchers are free to copy, distribute or transmit the thesis on the condition that they attribute it, that they do not use it for commercial purposes and that they do not alter, transform or build upon it. For any reuse or redistribution, researchers must make clear to others the licence terms of this work.

Abstract

This thesis describes a new measurement of the electron electric dipole moment (eEDM, d_e) made using a pulsed beam of ytterbium fluoride (YbF) molecules. YbF molecules are used as they greatly enhance the eEDM interaction with an applied electric field. In addition they suppress interactions with magnetic fields in the plane perpendicular to the applied electric field. This is hugely beneficial for suppressing the systematic effect that limited previous atomic eEDM searches.

We measure the eEDM by performing a type of separated oscillating field interferometry, wherein the direction of applied electric and magnetic fields are reversed in between pulses of the molecular beam. From a dataset of 6194 individual eEDM measurements we find $d_e = (-2.4 \pm 5.7_{\text{stat}} \pm 1.5_{\text{syst}}) \times 10^{-28} \text{ e cm}$. This result is consistent with zero, so we set a new upper limit of $|d_e| < 10.6 \times 10^{-28} \text{ e cm}$ at the 90% confidence level. A complete analysis of the dataset is given, with a thorough account of all the supplementary tests that were performed to check for systematic error.

After publishing this world leading result we proceeded to upgrade the experiment to improve eEDM sensitivity and reduce certain systematic effects. This involved improving the rf polarisation along the parallel plate transmission line, shortening the rf pulse length and improving the magnetic shielding. A detailed discussion of the development and testing of the upgrades is given, including new measurements of the systematic uncertainties which will limit our next eEDM measurement.

Acknowledgements

I would like to thank my supervisors Jony Hudson and Ben Sauer for their guidance and assistance. I especially appreciated your patience when I was struggling to explain something. I would also like to thank my fellow EDM PhD students Dhiren Kara and Jack Devlin, with whom I have thoroughly enjoyed working. Thanks also to Ed Hinds and Mike Tarbutt for their insight and guidance during our weekly meetings.

Thanks to the rest of CCM and Ion Trapping for making my PhD an enjoyable experience, particularly the other residents of Bay 3: Tom W, Aki M, Valentina Z and Anne C. Our discussions were always excellent, the music was mostly good.

Thanks to Jon Dyne and Steve Maine for their technical advice and expertise, and Sanja Maricic for all the administrative support.

Thanks to my friends and family for their support and encouragement, and thanks to Charlotte for her continued understanding and support, particularly while writing this thesis.

Contents

Declaration	2
Abstract	3
Acknowledgements	4
List of Figures	8
List of Tables	10
1 Introduction	11
1.1 Symmetry in Physics	11
1.1.1 Baryogenesis	12
1.1.2 Searches for Symmetry Violating Physics	14
1.2 The Electric Dipole Moment	15
1.2.1 EDMs and Symmetry Violation	16
1.2.2 EDM Predictions	17
1.3 Measuring the Electric Dipole Moment	18
1.3.1 Schiff’s Theorem and Electron EDM Enhancement	19
1.3.2 Atoms and Molecules for EDM Measurement	21
1.3.3 Common Systematic Effects	22
1.3.4 Review of EDM Experiments	23
1.3.5 Prospects for Measuring an Electron EDM	25
1.3.6 The Ytterbium Fluoride Electron EDM Experiment	27
1.4 The YbF Molecule	28
2 The YbF Interferometer	30
2.1 Interferometer Overview	30
2.1.1 Extracting an EDM	34
2.2 Interferometer Theory	35
2.2.1 Evolution In Static Fields	35
2.2.2 RF Pulse Theory	37
2.2.3 Interference Lineshape	39
2.3 Apparatus	41
2.3.1 Vacuum Chamber	41
2.3.2 YbF Beam Source	43
2.3.3 Pump and Probe Laser System	45
2.3.4 Laser Induced Fluorescence Detection	47

2.3.5	Applied RF Fields	49
2.3.6	Applied Electric Field	54
2.3.7	Ambient Magnetic Fields	56
2.3.8	Applied Magnetic Fields	57
2.3.9	Computer	59
2.4	EDM Data Acquisition Cycle	63
3	Analysing Detector Signals	65
3.1	Switch Channels and Waveforms	65
3.2	Non-linear Channels	70
3.3	Important Channels	72
3.4	Lock Channels	74
3.5	Manual Reversals	76
3.6	EDM Calculation and the Lineshape Correction	77
3.7	Blind Analysis	80
3.8	Trimmed Mean and Statistical Uncertainty Calculations	81
3.9	Non-zero Channel Finder	83
4	The 2011 Dataset	84
4.1	Defining the PMT Gates	85
4.2	Empirical Corrections	85
4.2.1	Ambient Magnetic Field Correction	86
4.2.2	RF Detuning Phase Correction	87
4.3	Systematic Uncertainties	91
4.3.1	Uncancelled Asymmetric Electric Field Reversals	91
4.3.2	Electric Field Voltage Offset From Ground	94
4.3.3	Magnetic Fields	99
4.3.4	Geometric Phase	105
4.4	Non-Zero Channels	107
4.5	Other Correlations	112
4.5.1	Pump/Probe Laser Polarisation Dependence	112
4.5.2	Manual State Dependence	113
4.5.3	Molecular Arrival Time Dependence	113
4.6	EDM Sensitivity Variation with PMT Gates	114
4.7	EDM Value	116
5	Upgrading the YbF Interferometer	119
5.1	Upgraded Apparatus Setup	120
5.1.1	RF Transmission Line	120
5.1.2	High Voltage Plate Flattening	128
5.1.3	Shorter RF Pulses and a More Powerful RF Amplifier	129
5.1.4	New Magnetic Shields	130
5.1.5	Increased Electric Field	133
5.1.6	Independent Pump and Probe Laser Frequency Modulations	133
5.1.7	Third Turbo Pump on the Source Chamber	134
5.1.8	Updated Lock Channel Analysis Gates	135
5.2	Systematic Limitations of the Upgraded Interferometer	135

5.2.1	Detuning Phase Correlation	136
5.2.2	Asymmetric Electric Field Reversal	136
5.2.3	Electric Field Voltage Offset	137
5.2.4	Perpendicular Magnetic Field Test	139
5.2.5	Laser Tests	141
5.3	EDM sensitivity	145
6	Conclusions and Outlook	147
A	Expansion of PMT Signal	158
B	The Shot Noise Limit	161
C	The Upgraded RF Transmission Line	163
D	Parameter Timings and Saved Values	171

List of Figures

1.1	T and P reversals on both $\vec{d}$ and $\vec{s}$.	16
1.2	Variation of E_{eff} with E_{app} for YbF in its ro-vibrational ground state.	21
1.3	YbF energy levels.	29
2.1	An illustration of spin precession.	30
2.2	Overview of EDM experiment.	31
2.3	Sampled points for an EDM measurement.	34
2.4	Relative Stark shifts of ground state hyperfine levels.	35
2.5	Hyperfine level shifts due to electric and magnetic fields.	37
2.6	Schematic of molecular beam machine.	41
2.7	Schematic of the upper chamber.	42
2.8	Schematic of source chamber.	44
2.9	Source parameter scans.	45
2.10	Pump/probe laser polarisation rotation optics.	46
2.11	Typical probe PMT time of flight profile.	49
2.12	Schematic of rf system.	50
2.13	RF parameter scans.	51
2.14	Ramsey phase shift across molecular cloud.	53
2.15	Schematic of high voltage supply and plate coupling.	55
2.16	Typical interference curve as scanned with applied magnetic field.	58
2.17	Schematic of the magnetic field current supply.	59
2.18	Schematic of the computer–hardware interface.	60
2.19	The sequence of events that occur during molecular transit.	61
3.1	8 bit switch sequences that make up the simplest EDM measurement.	68
3.2	Illustration of $\{B\}$ for two average interferometer phases.	69
3.3	An Illustration of the multiplication/division of analysis channels.	71
3.4	Demonstration of experimental parameter locking.	75
3.5	PMT signal model.	78
3.6	Distribution of measured EDM values.	81
4.1	Correlation between d_e and $\widehat{M}_E \{E\}_{\text{Bart}}$.	86
4.2	Calibration of $\delta E \propto \{E \cdot RFiF\}/\{\delta B\}$.	89
4.3	Correlation between d_e and δE when $\bar{V} = 0.0 \text{ V}$.	90
4.4	Correlation between d_e and δE when $\bar{V} = 102.5 \text{ V}$.	92
4.5	Variation in d_e across the molecular cloud observed when $\bar{V} = 102.5 \text{ V}$.	96
4.6	Correlation between d_e and $\delta_{\text{rf}}^{(1)}$.	97
4.7	Variation in d_e across the molecular cloud observed when $\bar{V} = 0.0 \text{ V}$.	97

4.8	Correlation between d_e and $\bar{V}$.	99
4.9	Splitting between ($F = 1, m_F = \pm 1$) due to perpendicular magnetic fields.	100
4.10	A conservative model of leakage current flow.	101
4.11	Electric field rotations within machine.	106
4.12	Correlation between d_e and pump/probe laser polarisation.	112
4.13	Correlation between d_e and machine manual state.	113
4.14	EDM sensitivity vs analysis gate position/width.	115
5.1	Schematic of original rf coupling scheme.	121
5.2	Example of Zeeman split rf resonance lineshape for field mapping.	122
5.3	Comparison of rf polarisation maps.	123
5.4	Upgraded rf coupling schemes.	124
5.5	Loop antenna measurements of original and upgraded rf transmission lines.	125
5.6	Schematic of final iteration of rf transmission line.	126
5.7	Typical time of flight profile demonstrating rf pickup.	127
5.8	Comparison of electric field maps.	128
5.9	Lower half of science chamber following rf transmission line upgrade.	130
5.10	$\widehat{M}_E \{E\}_{\text{mag}}$ distributions before and after apparatus upgrade.	132
5.11	Correlation between d_e and δE in upgraded setup.	137
5.12	Correlation between d_e and $\delta_{\text{rf}}^{(1)}$ in upgraded setup.	138
5.13	Correlation between d_e and $\bar{V}$ in upgraded setup.	139
5.14	Variation in d_e with applied perpendicular magnetic field.	140
5.15	Optical spectrum of $Q(0)$ transition.	141
5.16	Illustration of optical pumping $F = 1$ with circularly polarised light.	143
5.17	Decomposition of laser polarisation ellipse into linear and circular components.	144
C.1	3D model of the electric and rf field plate assembly.	163
C.2	Cross section of the field plate assembly in xz -plane.	164
C.3	The rf transmission line plates.	165
C.4	Cross section of the plate support structure in xy -plane.	165
C.5	RF plate support extensions.	166
C.6	The rf plate support feet.	167
C.8	Photos of the field plate assembly.	168
C.9	RF coupling at the top of the plate assembly.	169
C.10	RF coupling at the bottom of the plate assembly.	170

List of Tables

1.1	Enhancement factors of various atoms and molecules.	20
3.1	Details of the experiment detectors.	67
3.2	Applied experimental parameter modulations.	69
3.3	The physical meaning of the most important analysis channels.	72
3.4	The channels used to optimise experimental hardware settings.	75
4.1	Non-zero channels.	108
4.2	Differences in d_e between manual states.	114
4.3	Summary of the uncertainties which limit our measurement of d_e	117
4.4	Bootstrap calculated bounds on $ d_e $	117
5.1	Reduction in $\widehat{M}_E \{E\}_{\text{Bart}}$ due to a third magnetic shield.	132
5.2	Summary of systematic uncertainties following apparatus upgrade.	145
D.1	Pattern Generator timing values.	171
D.2	Data acquired per block.	172

Chapter 1

Introduction

In this thesis I describe the recent progress of the experiment at the Centre for Cold Matter, Imperial College London, which measures the electron electric dipole moment (eEDM) using a pulsed beam of ytterbium fluoride (YbF) molecules.

The motivations for such a measurement are discussed in this chapter. This begins with a discussion of time reversal asymmetry and the need for symmetry violating physics. An explanation of how an EDM relates to time and parity symmetry violation is then presented, followed by a discussion of how an EDM is measured. A brief review of other EDM experiments around the world is then given.

1.1 Symmetry in Physics

While the subject of time has many aspects, one of its most interesting, and one might argue intuitively obvious, facets is the distinction between direction of flow. In everyday life there appears an obvious ‘arrow of time’. Consider, for example, a video recording of a snooker game, where the triangle of red balls is broken apart due to an initial collision with the cue ball. It is easy to tell whether the recording is played forward or in reverse.

Mathematically the time reversal operation, T , corresponds to the substitution $t \rightarrow -t$, which has no effect on the classical laws of physics. You cannot tell the direction of time flow from a video recording of a two balls colliding and yet in a snooker break each ball simply undergoes classical elastic collisions. This leads to a paradox: since complex, macroscopic systems are the aggregate of many microscopic systems obeying basic laws of physics, how can macroscopic systems appear so time asymmetric?

The behaviour of a macroscopic system, however, is not only a result of its microscopic processes, but also its initial boundary conditions. It is these initial conditions that give

rise to the apparent time asymmetry. Consider again the snooker example above; if the balls were to begin in the disordered broken apart state, while it is possible that they all might coalesce and reform a perfect triangle of balls, this situation is highly improbable. As Boltzmann explained in 1895 [1], a complex system moves towards an equilibrium corresponding to the most probable macroscopic state, attained by the most number of microstates.

Nonetheless, it is believed that the behaviour of many physical systems is fundamentally asymmetric in the direction of time flow. In fact, time symmetry violation is just one of the three important discrete transformation operations in which symmetry may not hold. The two other transformations are charge conjugation (C) — the replacement of all particles with their complementary antiparticle, and parity reversal (P) — a reversal of all coordinate axes.

Of course these discrete transformations need not be considered independently. A system which is asymmetric when both charge and parity are reversed is known as CP violating, for example. Actually, under the simultaneous transformations of C, P and T, it is believed that symmetry does hold. This is the so called CPT invariance theorem. If, for example, CP symmetry was violated, via the CPT theorem, this would imply a symmetry violation of T.

In the 1950s Lüders and Pauli showed that CPT invariance follows from the very basic assumptions of Lorentz invariance, quantum mechanics, and that particle interactions are represented by fields [2, 3]. All popular modern quantum field theories abide by these assumptions. It is worth highlighting, though, that as CPT invariance is based on assumptions, it is still subject to experimental tests like all other discrete symmetry transformations. To date no experimental evidence exists for CPT violation; a thorough review is presented in reference [4].

1.1.1 Baryogenesis

The very nature of our universe requires symmetry violating physics. 13 billion years after the big bang we observe an imbalance in the amount of matter and antimatter present in the universe. The evidence for this so called ‘baryon asymmetry problem’ is quite strong. The absence of γ -radiation from annihilation reactions both locally and in intergalactic space is the predominant argument. The number density of antiprotons present in cosmic rays is also consistent with their secondary production in accelerator-like processes such as $p + p \rightarrow 3p + \bar{p}$, so it is unlikely they are primordial. An in depth discussion is given by Steigman in [5].

By measuring the spatial fluctuations in the cosmic microwave background, Bennett *et al.* [6] have measured the current baryon asymmetry, η_0 , to be

$$\eta_0 = \frac{n_b - n_{\bar{b}}}{n_\gamma} = (6.19 \pm 0.14) \times 10^{-10}, \quad (1.1)$$

where n_b and $n_{\bar{b}}$ are the respective number densities of baryons and antibaryons, and n_γ is the number density of photons. This value is in excellent agreement with the prediction from a theory known as primordial nucleosynthesis, which can estimate the baryon asymmetry from the known relative abundances of the light elements H, D, He, Li, Be and B.

A possible explanation for the observed asymmetry is that it was present when the universe began. But this possibility still leaves many questions unanswered, such as what caused the preference for matter over antimatter? Creationist arguments aside, the observed asymmetry is unlikely to be the result of an initial asymmetry due to the predictions of the widely popular cosmic inflation model. This model predicts that shortly after the big bang there followed a period of exponential expansion of the early universe [7]. The rapid inflation provides a solution to many problems in cosmology, however, it would completely dilute any initial baryon asymmetry. We therefore require symmetry violating physics from which a baryon asymmetry can be established over time, without the need to impose an asymmetry in the initial conditions of the universe.

In 1967 Sakharov outlined three requirements that all ‘baryogenesis’ models must obey to generate the baryon asymmetry that we observe today [8]: 1) baryon number (the number of baryons minus the number of anti-baryons) cannot be preserved; 2) C and CP (or T) symmetries must be violated 3) reaction processes must occur out of thermal equilibrium.

The first condition is obvious, since in the beginning $\eta = 0$ and now $\eta \neq 0$. We require the second criterion as if C or CP (or T) are exact symmetries then it can be shown that the total rate for any process which produces an excess of baryons is equal to the rate of the complementary process that produces an excess of antibaryons. Even if the baryon number is not preserved in a reaction, no net baryon number will result unless C and CP (or T) are violated.

The third condition is a little complicated. In thermal equilibrium, the abundance of any given baryon (antibaryon), n_X ($n_{\bar{X}}$), is dictated only by its mass, m_X ($m_{\bar{X}}$) and temperature [9]. Since $m_X = m_{\bar{X}}$, this implies that, if temperature is held constant, $n_X = n_{\bar{X}}$ irrespective of the rates of the processes involving X and $\bar{X}$, so no baryon asymmetry is possible. Consider now the case that temperature is not constant, indeed lets assume, as above, that because of expansion the universe temperature is cooling.

Consider also the baryon number and CP violating abstract decays $P \rightarrow Q + R$ and $\bar{P} \rightarrow \bar{Q} + \bar{R}$. If the cooling is rapid compared to the decay rates, then reactions cannot respond and the abundances n_P and $n_{\bar{P}}$ do not remain in thermal equilibrium. An excess of P and $\bar{P}$ is created above the equilibrium abundance. Over time P and $\bar{P}$ eventually decay, but because the decay rates are not equal, and because the decays do not preserve baryon number, a preference for matter over antimatter may result.

1.1.2 Searches for Symmetry Violating Physics

The first evidence for symmetry violating physics was found in 1957 by Wu *et al.* [10]. They observed that electrons were preferentially emitted in the opposite direction to the nuclear spin of the parent cobalt-60 isotope, thus demonstrating the violation of parity symmetry. Many experiments have since demonstrated P violation in various physical systems (see for example [11, 12]), but at the time the result was entirely surprising. P violation led to the assertion of C violation in order to preserve the invariance of C and P combined. However, in 1964 Christenson *et al.* demonstrated that the long-lived neutral kaon, K_L^0 , occasionally decayed into two charged pions [13] — a decay that is strictly forbidden by CP symmetry. Like P violation, CP violation is now a well established phenomenon and has been observed in a few physical systems. A thorough review is given in reference [14].

Naturally physicists assume that CP violation is an indication of T asymmetry but, as Purcell and Ramsey pointed out [15, 16], this can only be verified experimentally. The CPLEAR experiment at CERN [17], and the BaBar detector at SLAC provide the only known evidence for T asymmetry, without reference to the CPT theorem. The former experiment looked for a difference in rates between the supposedly time symmetric $K^0 \rightarrow \bar{K}^0$ transformations and vice versa, whereas the latter experiment studied the rates at which neutral B mesons change quantum state [18]. Both experiments are incredibly complex. Clearly the importance of verifying T-reversal asymmetry in other, preferably simpler, systems with different systematic effects is undeniable. The YbF eEDM experiment and other experiments aim to do just this by looking for a permanent electric dipole moment (EDM) in fundamental particles [19].

1.2 The Electric Dipole Moment

For a simple system of two opposite charges, $-q$ and $+q$, separated by the vector $\vec{r}$, originating at the negative charge, the electric dipole moment (EDM, $\vec{d}$) is defined by the product $q\vec{r}$. However, for a particle with an inhomogeneous charge density, $\rho(\vec{r})$, the EDM takes the more complex form [20]:

$$\vec{d} = \int_V \rho(\vec{r}) \vec{r} \partial V, \quad (1.2)$$

where $\vec{r}$ is now measured from the centre of mass, and ∂V is a differential volume element. For fundamental (point) particles such as the electron, this equation describes the permanent polarisation of the vacuum field surrounding the particle.

Perhaps the most general definition of $\vec{d}$ is that it is simply a property of a particle defined through its interaction with an applied electric field, $\vec{E}$. This interaction can be described with the Hamiltonian:

$$H_d = -\vec{d} \cdot \vec{E}. \quad (1.3)$$

For a particle with spin $\vec{s}$, the orientation of $\vec{d}$ relative to $\vec{s}$ is constrained by the projection theorem, a specific case of the Wigner-Eckart theorem, which states that [21]:

$$\langle s, m_s | \vec{d} | s, m_s \rangle = \langle s, m_s | \vec{s} | s, m_s \rangle \frac{\langle s, m_s | \vec{d} \cdot \vec{s} | s, m_s \rangle}{\hbar^2 s(s+1)} = c_s \langle s, m_s | \vec{s} | s, m_s \rangle, \quad (1.4)$$

where the quantum state $|s, m_s\rangle$ is specified by the eigenvalues of s^2 and s_z . Equation 1.4 shows that $\vec{d}$ must lie either parallel or anti-parallel to $\vec{s}$ as the matrix element $\langle s, m_s | \vec{d} \cdot \vec{s} | s, m_s \rangle$ is independent of the spin projection, m_s . $\vec{d} \cdot \vec{s}$ is a scalar operator, $[\vec{d} \cdot \vec{s}, s_z] = 0$, so has independent eigenvalues to s_z .

Throughout this thesis I will often refer to an EDM with the scalar variable d . Unless otherwise stated this assumes the convention:

$$\vec{d} = d \frac{\vec{s}}{|\vec{s}|}. \quad (1.5)$$

1.2.1 EDMs and Symmetry Violation

As a particle's EDM must lie either parallel or anti-parallel to its spin, there are four possible degenerate states that it may occupy. These are illustrated in figure 1.1. However, the Pauli exclusion principle¹ states that spin-1/2 fermions may only have two degenerate states. Under the assumption of rotational invariance, this means that $\vec{d}$ may only have one relative orientation to $\vec{s}$. The two possibilities are labelled type A and B in the figure. As illustrated, under a parity transformation the direction of $\vec{d}$ reverses whereas $\vec{s}$ does not. Conversely, under a time-reversal transformation $\vec{s}$ reverses whereas $\vec{d}$ does not. Both these transformations take the particle into a forbidden state indicating that $\vec{d}$ must vanish else both P and T symmetries would be violated. Measurement of a non-zero EDM in spin-1/2 particles, such as the electron, proton or neutron, therefore provides direct evidence for both P and T symmetry violation.




FIGURE 1.1: A dipole moment, $\vec{d}$, may lie either parallel or anti-parallel to the particle spin $\vec{s}$. Under a parity reversal the direction of $\vec{d}$ reverses and $\vec{s}$ remains unaffected. For a time reversal the direction of $\vec{s}$ is reversed and $\vec{d}$ remains unchanged.

Put mathematically, for both parity and time symmetries to hold we require that $\langle \psi | H | \phi \rangle = \langle P\psi | H | P\phi \rangle$ (P is unitary) and $\langle \psi | H | \phi \rangle = \langle T\phi | H | T\psi \rangle$ (T is anti-unitary) for general quantum states $|\psi\rangle$ and $|\phi\rangle$. As usual H is a (Hermitian) Hamiltonian. Consider the following:

¹There is very strong evidence for the Pauli exclusion principle, see for example [22] and references therein. One such demonstration of the two internal states an electron may occupy is given by the structure of the periodic table of elements.

$$\begin{aligned}
\langle \psi | H | \phi \rangle &= \langle \psi | P^\dagger P H P^\dagger P | \phi \rangle \\
&= \langle P\psi | P H P^\dagger | P\phi \rangle. \quad (1.6)
\end{aligned}
\qquad
\begin{aligned}
\langle \psi | H | \phi \rangle &= \langle \psi | T^\dagger T H T^\dagger T | \phi \rangle \\
&= \langle T\psi | T H T^\dagger | T\phi \rangle^* \\
&= \langle T\phi | T H T^\dagger | T\psi \rangle. \quad (1.7)
\end{aligned}$$

So we require $H_d = T H_d T^\dagger$ and $H_d = P H_d P^\dagger$. It is often quoted (see for example [23]) that an EDM violates time and parity symmetries as the Hamiltonian $H_d = -\vec{d} \cdot \vec{E} = -d \frac{\vec{s}}{|\vec{s}|} \cdot \vec{E}$ changes sign under both transformations. This argument, however, implies that there are no T and P conjugate particles with dipole moment $-d \frac{\vec{s}}{|\vec{s}|}$. If there were we could write a more complicated Hamiltonian such as $H_d = c(\vec{s} \cdot \vec{E})(\vec{s} \cdot \vec{d})$ and symmetry would be restored. Here c is some real constant and the scalar product $\vec{s} \cdot \vec{d}$ accounts for the orientation of $\vec{d}$ relative to $\vec{s}$. But note that we must use the Pauli exclusion principle (or some other relevant restriction on the number of states) and state that spin-1/2 fermions have no extra degeneracies, so $H_d = 0$ else both time and parity symmetries are violated.

To confuse matters a little, many molecules, such as YbF, have a permanent electric dipole moment in the reference frame of the molecule. For diatomic heteronuclear molecules, the dipole moment lies along the internuclear axis. However, in the absence of any fields there is no dipole moment in any laboratory axis as molecular rotation averages the molecular dipole moment to zero. An electric dipole moment can be induced in the laboratory frame by applying an electric field. Naturally this dipole reverses with the applied electric field. Permanent dipole moments in fundamental particles do not reverse with an applied electric field because they lie along their internal angular momentum. When an external electric field is applied to fundamental particles, the electric dipole moment and angular momentum vectors precess about the axis of the electric field. Reversing the direction of the applied electric field reverses the handedness of the precession, but the projection of the dipole moment onto the axis of the electric field remains constant. Crucially, molecular dipole moments (or indeed induced atomic dipole moments) do not violate time and parity symmetries as there is no linear Stark shift in the limit of weak applied electric fields.

1.2.2 EDM Predictions

The standard model of particle physics is a very well established theory that provides a description of the fundamental constituents of the universe and their interactions. There is almost no experimental evidence that contradicts the predictions of the standard model, however, few would argue that the theory is complete. It cannot, for example, incorporate gravity, nor account for the baryon asymmetry that we observe in the universe today [24].

In fact the standard model only incorporates a very limited amount of CP (T) violation, so predicts very small EDM values. It predicts the EDM of the proton, d_p , and neutron, d_n , to lie in the range $10^{-32} - 10^{-31}$ e cm [25, 26], and the electron EDM, d_e , to be $< 10^{-38}$ e cm [27]. These values are well below the current experimental upper limit, but if a non-zero EDM value were to be measured to lie within a few orders of magnitude of the current upper limit, it would provide new evidence for physics beyond the standard model.

There are many new theories in particle physics that attempt to address the limitations of the standard model. Examples include the various forms of supersymmetry. These new models include extra sources of symmetry violation and therefore predict dipole moments in the range $10^{-25} - 10^{-30}$ e cm [28]. EDM values in this range are now becoming measurable with modern experimental techniques. So by improving the measured EDM limits we can better constrain, or even prove the invalidity of these new models.

1.3 Measuring the Electric Dipole Moment

Fundamentally all EDM experiments attempt to measure the energy shift, $U = \langle -\vec{d} \cdot \vec{E} \rangle$, that is induced by an applied electric field, $\vec{E}$. This is very difficult because the shift is extremely small and is easily obscured by the magnetic dipole interaction $\langle -\vec{\mu} \cdot \vec{B} \rangle$.

Consider, for example, an electron EDM at the current experimental upper limit $d_e = 10^{-27}$ e cm. Even with a large applied electric field of 1 GV/cm, the EDM shift is only 0.25 mHz. This is the same as the Zeeman shift caused by a magnetic field of only 17 fT. For comparison, the magnetic field surrounding the Earth is about 50 mT in amplitude and its fluctuations are certainly larger than a few fT. For this reason EDM experiments must be extremely well shielded from external magnetic fields and particular care must be taken to control those which result from changes in the applied electric field.

Most modern EDM searches experiment on neutral matter such as neutrons, atoms or molecules, as the large applied electric fields would accelerate any charged particles. Heavy atoms and molecules provide a particularly nice system for measurement of the electron EDM as they are able to enhance the EDM interaction making it more easily measurable. This is an amazing result that is completely counterintuitive, as you would expect the charge within the atom or molecule to rearrange itself to screen the external electric field. However, in 1963 Schiff showed that under certain circumstances non-zero proton, neutron or electron EDMs can have an observable impact on the interaction of an atom/molecule with an applied electric field [29]. This is discussed in the following sections.

1.3.1 Schiff's Theorem and Electron EDM Enhancement

In his famous paper [29] Schiff began by demonstrating that for an electrostatically bound system of point charges, in the absence of relativistic effects, the energy eigenstates of the system do not depend upon the EDMs of its constituent particles. However, real atoms and molecules are not collections of point particles, the nucleus in particular has a finite size, and relativistic effects are important. Schiff went on to show that when these effects are taken into account, non-zero proton, neutron and electron EDMs can contribute to an overall atomic/molecular permanent EDM.

In diamagnetic atoms electron pairing cancels any contribution from d_e . In these atoms the finite size of the nucleus permits an atomic EDM, d_A , proportional to both d_p and d_n , albeit attenuated by $\sim 10^4$ [30]. Conversely, paramagnetic atoms contain at least one unpaired electron, so d_A may result from a non-zero d_e . In fact, as Sandars discovered in 1965 [31], in particularly heavy atoms the relativistic motion of unpaired electrons can lead to a proportionality

$$R = \frac{d_A}{d_e} \gg 1. \quad (1.8)$$

The ‘enhancement factor’, R , depends upon the polarisation of the electronic wavefunction. In atoms the polarisation comes from a mixing of opposite parity states. As the interaction with the electric field is tiny in comparison to the energy level splitting a perturbative calculation provides an excellent approximation. To lowest order the mixing of opposite parity states and hence electronic polarisation is linear in applied electric field, $\vec{E}_{\text{app}}$. So we can write the EDM interaction as

$$H_d = -R \vec{d}_e \cdot \vec{E}_{\text{app}}, \quad (1.9)$$

where it can be shown [32–34] that

$$R \approx 10Z^3\alpha^2. \quad (1.10)$$

As is usual, Z is the atomic number and α is the fine structure constant.

In polar molecules the wavefunction is naturally mixed, which results in a strong polarisation along the internuclear axis. The polarisation is particularly large if the molecule contains a highly electronegative element such as fluorine. So by combining fluorine with a heavy element such as ytterbium, a heavy polar molecule, such as YbF, can produce a relativistic enhancement of the electron EDM that is over 1000 times larger than any atom [35]. Of course molecules rotate in free space so their polarisation averages to zero. A modest applied electric field is therefore required to align the molecule. However, the polarisation is not linear in applied field, so the convention is to ascribe the

enhancement to the electric field. Accordingly we write the EDM interaction as

$$H_d = -\vec{d}_e \cdot \vec{E}_{\text{eff}}, \quad (1.11)$$

and define the effective electric field as

$$\vec{E}_{\text{eff}} = E_{\text{eff}}^{\text{max}} \eta(E_{\text{app}}) \hat{z}, \quad (1.12)$$

where $\hat{z}$ is a unit vector parallel to the applied field, $\vec{E}_{\text{app}}$. $E_{\text{eff}}^{\text{max}}$ is a constant that depends upon the molecular structure.

$$\eta(E_{\text{app}}) = \langle \hat{n} \cdot \hat{z} \rangle \quad (1.13)$$

is a (dimensionless) polarisation factor as defined with the unit vector $\hat{n}$ pointing along the molecular dipole moment, i.e. from the negative ion to the positive². For YbF in its rotational and vibrational ground state, $\eta(E_{\text{app}})$ is easily calculated by considering the (DC) Stark shift of a rigid rotor³.

Table 1.1 lists the enhancements for the atoms and molecules discussed in the following sections. Figure 1.2 demonstrates the variation of E_{eff} with E_{app} for YbF in its rotational (and vibrational) ground state. As discussed in section 2, we typically operate our experiment with $E_{\text{app}} = 10 \text{ kV/cm}$. Here $\eta(E_{\text{app}}) = 0.558$ and $E_{\text{eff}} = -14.5 \text{ GV/cm}$, an enhancement of -1.45×10^6 !

Species	Enhancement	Reference
Hg	$d(^{199}\text{Hg}) = 5.5 \times 10^{-5} d_p + 5.3 \times 10^{-4} d_n$	[37]
Tl	$d_{Tl} = -585 d_e$	[38, 39]
Cs	$d_{Cs} = 124 d_e$	[38]
Gd IV	$d_{Gd IV} \sim -3.3 d_e$	[40]
YbF	$E_{\text{eff}}^{\text{max}} = -26 \text{ GV/cm}$	[41–46]
PbO	$E_{\text{eff}}^{\text{max}} = 26 \text{ GV/cm}$	[47]
ThO	$E_{\text{eff}}^{\text{max}} = 104 \text{ GV/cm}$	[48]
WC	$E_{\text{eff}}^{\text{max}} = -54 \text{ GV/cm}$	[49]
HfF ⁺	$E_{\text{eff}}^{\text{max}} = 24 \text{ GV/cm}$	[50]

TABLE 1.1: Enhancement factors for various atoms and molecules.

²Some papers use the convention that $\hat{n}$ points from the heavy nucleus to the light nucleus. In YbF this vector would point in the opposite direction to molecular dipole moment.

³This is defined by the Hamiltonian $H = B J^2 - \mu E_{\text{app}} \cos(\theta)$, where B is a rotational constant dependent upon the molecular moment of inertia, $\vec{J}$ is the molecular angular momentum and μ is the molecular dipole moment [36]. $\eta(E_{\text{app}}) = \langle \cos(\theta) \rangle$ so is proportional to gradient $\langle \partial H / \partial E_{\text{app}} \rangle$.




FIGURE 1.2: Variation of E_{eff} with E_{app} for YbF in its ro-vibrational ground state. At our usual operating field, $E_{\text{app}} = 10 \text{ kV/cm}$, $\eta(E_{\text{app}}) = 0.558$ giving $E_{\text{eff}} = -14.5 \text{ GV/cm}$. This is illustrated by the dashed line.

1.3.2 Atoms and Molecules for EDM Measurement

Atoms particularly the alkalis, are a natural choice for experimentation — they are easy to manipulate, are spectroscopically well understood and precise calculations can predict their behaviour. Furthermore, as was demonstrated in the last section, they can amplify the electron EDM interaction.

Many polar molecules also have these benefits and can produce enhancement factors that are many orders of magnitude larger than atoms. Polar molecules have a further benefit in that they interact with electric fields much more strongly than they do magnetic fields. This is in direct contrast to atoms where magnetic field interactions dominate over electric. As will be shown in section 4.3.3, applying an electric field to polar molecules drastically suppresses any interaction with magnetic field components perpendicular to the axis of the electric field. Only magnetic fields in the axis of the applied electric field are able to cause systematic error.

As shown in table 1.2 there are a number of polar molecules which are suitable for an electron EDM measurement. However, the size of $E_{\text{eff}}^{\text{max}}$ is not the only requirement for consideration. We must also consider: how easily $\eta(E_{\text{app}})$ saturates; the difficulty of producing enough molecules in the required quantum state; the lifetime of the state of interest; safety concerns such as whether the molecule is toxic or radioactive; as well as other practical requirements such as the availability of lasers at the required frequencies.

1.3.3 Common Systematic Effects

Before reviewing the major EDM experiments in turn, it is worth briefly describing a few common sources of systematic error that are present in many, if not all, EDM experiments, so that the relative advantages of each may be compared. A comprehensive discussion of potential systematic errors is given by Lamoreaux and Golub in [51], and a more thorough discussion of the various systematic errors regarding the YbF electron EDM experiment is given in section 4.3.

In general there are two main types of EDM experiment, namely confinement type and beam type experiments. In confinement type experiments particles, be they neutrons or atoms etc., are contained within a well defined space where their movement is limited. In beam type experiments, particles pass through a vacuum chamber and are manipulated and probed at various positions along the machine. Typically they travel at a high velocity and have a narrow velocity spread.

Of primary concern are magnetic fields which depend on the applied electric field. Energy shifts due to such fields can masquerade as an EDM induced shift, as the magnetic dipole interaction, $-\vec{\mu} \cdot \vec{B}$, perturbs a particle in exactly the same way as the electric dipole interaction, $-\vec{d} \cdot \vec{E}$.

Leakage currents that flow between high voltage electrodes can generate a magnetic field. In the worst case scenario the magnetic field not only aligns in part with the electric field, but its direction also reverses following polarity changes in $\vec{E}$. Leakage currents can be quite troublesome in confinement type experiments as the cell walls are often used to separate the high voltage electrodes. In beam type experiments the high voltage feeds are quite far away from the interaction region, furthermore the electrode support structure is not so conducive to circulating current.

Another troublesome source of systematic error is the ‘motional magnetic field’. As the particles move through the applied electric field, $\vec{E}$, with velocity $\vec{v}$, a magnetic field $\vec{B}_v = \vec{E} \times \vec{v}/c^2$ is generated in the rest frame of the particles. This field is in addition to a static magnetic field, $\vec{B}$, which is applied, in most cases, to align the particle. If $\vec{B}$ is slightly misaligned from the applied electric field, $\vec{E}$, the static magnetic field would gain a component of $\vec{B}_v$, thus making the Zeeman energy shift dependent on the direction of $\vec{E}$. This type of systematic affects beam experiments most because of the low spread in particle velocities. It can, however, affect confinement experiments in a few subtle ways despite the random particle motions present therein (see reference [51]).

Another common source of systematic error is geometric phase. This can imitate an EDM interaction if the particles move through inhomogeneous electric and magnetic fields that

change when the applied electric field is reversed. Since this systematic depends upon the motions of individual particles it is particularly hard to track in confinement type experiments.

1.3.4 Review of EDM Experiments

In this section I give an overview of the three experiments which have set the best limits on the proton, neutron and electron EDM. These experiments are the most recent generation in a line EDM experiments stretching back to the late 1950s. For an in depth review of the history of EDM experiments see references [32, 52].

Ultracold Neutron Storage Experiment

Some of the first attempts to measure a permanent electric dipole moment were performed on neutron beams [53]. Since then the techniques for the slowing and confinement of neutrons have improved considerably. The most recent measurement by Baker *et al.* used a cell of ultracold neutrons (UCNs) to limit $|d_n| < 2.9 \times 10^{-26} \text{ e cm}$ with a 90% confidence level [54]. Their method follows Ramsey’s separated oscillating fields technique [55] which allowed them to measure the change in the Larmor precession frequency as the relative orientation of static electric and magnetic fields were changed from parallel to anti-parallel. This technique is similar to the YbF electron EDM experiment technique described in sections 2.1–2.2.

It was essential that magnetic fields inside the chamber were well characterised, so spin polarised ^{199}Hg atoms were also contained along with the UCNs. The Hg spin precession was constantly measured with circularly polarised UV, that was absorbed in proportion to the x component of its spin vector. As the EDM of ^{199}Hg had been measured to be much smaller than the experimental sensitivity of this experiment, its EDM interaction with $\vec{E}$ could be neglected. It, therefore, acted as an excellent in-situ magnetometer (a ‘comagnetometer’) [56].

The neutron velocity was very low, so motional systematic errors were small. Geometric phase effects contributed the largest uncertainty to the systematic error budget. Despite being a confinement type experiment, magnetic fields generated by leakage current constituted $\ll 1\%$ of the total systematic error budget.

Mercury Vapour Experiment

For a number of years Fortson *et al.* have been running an EDM experiment on a cell of mercury atoms at the University of Washington. Their latest measurement: $|d(^{199}\text{Hg})| < 3.1 \times 10^{-29} \text{ e cm}$ (95% CL) is the lowest EDM limit of any kind recorded yet [57]. Through Schiff moment calculations [37], their result implies a neutron EDM, $|d_n| < 5.8 \times 10^{-26} \text{ e cm}$, and a proton EDM, $|d_p| < 7.9 \times 10^{-25} \text{ e cm}$.

Mercury has a very long spin polarisation lifetime time which permits long probing periods and hence gives a correspondingly high statistical sensitivity. Fortson *et al.* took advantage of this, using a UV laser to both prepare the atomic state and monitor the spin precession rate about aligned electric and magnetic fields over prolonged periods of 100–200 s. The (atomic) EDM was measured by comparing the Larmor precession frequency for aligned and anti-aligned electric and magnetic fields.

Their latest apparatus contained four vapour cells positioned in a common magnetic field. The two middle cells had oppositely directed electric fields, allowing simultaneous measurement of the EDM interaction. The outer two cells were enclosed by the electrodes so experienced no electric field. These cells permitted precise monitoring of magnetic field fluctuations and cancellation of field gradients. The three main systematic uncertainties in their latest result were due to leakage currents, dataset cuts when leakage current sparks exceeded 100 pA, and EDM variations with experimental parameters such as laser power, frequency, spin lifetimes etc. By tipping the static magnetic field by $\pm 10^\circ$, motional magnetic field effects were found to be almost completely negligible.

Thallium Beam Experiment

Until very recently, the best limit on the electron EDM was set by Commins *et al.* at the University of California using beams of thallium atoms. This experiment was developed over 15 years, culminating in the final result of $|d_e| < 1.6 \times 10^{-27} \text{ e cm}$, as was published in 2002 [58]. The experiment had reached its limit in statistical sensitivity and the systematic uncertainties could not be negated or accounted for to any higher precision.

As the velocity of the thallium atoms was over 400 m/s, the motional magnetic field was a major source of systematic error. To combat this effect Commins *et al.* used two pairs of counter propagating beams. One pair of counter propagating beams would have been sufficient to control this effect, but the second pair enabled the simultaneous measurement with both orientations of $\vec{E}$ relative to $\vec{B}$. This provided common-mode noise rejection and control of certain systematic effects. Each beam contained thallium and sodium atoms and Ramsey’s method was simultaneously performed on both elements. Since

sodium is a relatively light atom it has a negligible EDM enhancement factor so served as an excellent comagnetometer.

In their final measurement the systematic uncertainty was entirely dominated by motional magnetic fields despite the Herculean effort to account for this error source with counter propagating beams. This perhaps best demonstrates the benefit of using polar molecules over atoms. The effects of motional magnetic fields are almost completely negligible due to the strong suppression from the (tensor) Stark shift.

1.3.5 Prospects for Measuring an Electron EDM

The YbF electron EDM experiment is not the only remaining experiment that aims to measure the electron EDM. In this section I will give a brief overview of the other experiments which have the potential to improve the limit on the electron EDM by at least an order of magnitude.

Other Polar Molecule Searches

The DeMille, Doyle and Gabrielse research groups are collaborating to build an experiment at Harvard University to measure d_e with ThO molecules (the Advanced Cold Molecule EDM experiment — ACME) [59]. This experiment utilises recent developments in buffer gas molecular sources to provide a cold slow beam of ThO with a brightness that far exceeds conventional supersonic beam sources [60]. The ThO molecule can be aligned to give a huge effective electric field of $E_{\text{eff}} \approx 100 \text{ GV/cm}$ for $< 100 \text{ V/cm}$ of applied field. In addition the state for EDM measurement has relatively long radiative lifetime of $\approx 2 \text{ ms}$ which permits a long phase evolution period, giving a correspondingly high experimental sensitivity. The salient feature of ThO is its Ω -doublet⁴ level structure that allows the molecules to be aligned parallel or anti-parallel to the applied electric field. This is advantageous as the EDM and Zeeman shifts are in the same direction for one relative orientation, and are opposed for the other, so some important systematic effects can be cancelled out by comparing the shifts in each orientation. One final benefit of ThO is that its magnetic moment is nearly zero so magnetic interactions are severely suppressed. The ThO measurement scheme involves comparison of the fluorescence generated by two perpendicularly polarised lasers. In the first generation experiment the research team expect to be achieve a statistical sensitivity of $4 \times 10^{-29} \text{ e cm } \sqrt{\text{day}}$ [59].

Leanhardt *et al.* are currently building an experiment at Michigan University to measure the electron EDM using a beam of WC molecules [49]. WC only requires a very small

⁴ Ω is the total (spin+orbital) electronic angular momentum projected onto the inter-molecular axis.

applied electric field of a few mV/cm to become completely polarised, so leakage currents should be very small. Furthermore, like ThO, WC also possesses an Ω -doublet state which will help eliminate certain systematics. The molecular source is still in an early stage of development, but Leanhardt is confident that they can produce enough molecules to be competitive with other polar molecule EDM experiments.

Cornell *et al.* at Colorado University are attempting to trap the molecular ion HfF^+ in an rf quadrupole [61, 62]. Ramsey's method will be performed in the presence of rotating electric and magnetic fields, which are necessary to trap the molecules. HfF^+ also benefits from an Ω -doublet state, and because the ion is trapped, motional fields are small and long spin coherence times of approximately 1 s are attainable. However, the rapidly changing fields may result in systematic errors which are hard to account for. In addition, the applied electric field cannot be reversed. Nevertheless, Cornell believes that the experiment can achieve a sensitivity of $6 \times 10^{-29} \text{ e cm } \sqrt{\text{day}}$.

Solid State Searches

Hunter *et al.* [63] and a collaboration between Liu and Lamoreaux [64, 65] are using two complementary techniques to look for a linear magnetoelectric effect in gadolinium solids. As the electron EDM must be aligned with its magnetic moment, upon application of a polarising magnetic (electric) field, the sample must produce a electric (magnetic) field of its own. Hunter measures the induced electric field with JFET amplifiers whereas Liu and Lamoreaux measure the induced magnetic field with a SQUID. The Gd crystal structures were chosen because: they are easily spin polarised; they have a relatively high enhancement factor (see table 1.1); they are good insulators so do not suffer large leakage currents; and in the absence of T-violation their symmetry does not permit a linear or quadratic magnetoelectric effect. These experiments have the potential to improve the current electron EDM limit by a factor of 10^3 [64], however the complex crystal structures do make systematic errors hard to quantify.

Caesium Experiments

There are number of experiments that aim to measure the electron EDM using caesium atoms. To overcome the motional magnetic field systematic effect that limited the thallium beam experiment Weiss, Heinzen and Chu plan to experiment on optically trapped caesium [66–68]. However, optically trapping atoms does create a few additional problems, such as light shifts, that will need to be overcome. Gould plans to use a caesium fountain [69] which should also suppress motional systematic effects, as the rise and fall

of each atom results in zero average velocity. Both approaches have the potential to limit d_e to approximately 10^{-29} e cm [68, 69].

1.3.6 The Ytterbium Fluoride Electron EDM Experiment

The electron EDM experiment at the Centre for Cold Matter, Imperial College London studies the polar molecule ^{174}YbF , which combines the heavy element ytterbium ($Z = 70$) with fluorine. At the normal operating applied electric field of 10 kV/cm, the molecule gives an effective electric field of -14.5 GV/cm, as shown in figure 1.2. The experiment has been running for a number of years now, and has reached a certain level of refinement which places it the current forerunner of electron EDM experiments.

In 2002 the group published their first result, measuring $d_e = (-0.2 \pm 3.2) \times 10^{-26}$ e cm [70, 71]. While this was not the most precise value at the time, it was limited by counting statistics rather than systematic errors, demonstrating the advantages of using molecules rather than atoms for an EDM measurement. In the 10 years that followed, the group has improved the molecule source, quantum state control and experimental automation, which has increased the sensitivity that can be achieved per day. The developments are described in detail in the theses of Condylis (2006) [72], Ashworth (2008) [73] and Kara (2010) [74].

Unfortunately with improved statistical sensitivity new systematics became apparent. Characterising and overcoming these limitations has formed a major part of the time spent in recent years. In 2011 the group published the most precise measurement of the electron EDM to date: $d_e = (-2.4 \pm 5.7_{\text{stat}} \pm 1.5_{\text{syst}}) \times 10^{-28}$ e cm (68.3% confidence) [75]. This is a significant result as it is the first time that a new limit on the electron EDM has been set using molecules rather than atoms.

I joined the YbF eEDM experiment in January 2010. At this time all of the experimental hardware and software had been built, and the experiment was fully capable of collecting large datasets consisting of many thousands of individual EDM measurements. During my PhD I was directly responsible for the setup, maintenance and running of the experiment. I also analysed the acquired data and presented the results to my supervisors during our weekly progress meetings.

When I joined the EDM team, the experiment was suffering from a number of systematic effects that were preventing us from undertaking a full data run. The early months of my PhD were spent running diagnostic experiments to try and identify the sources of error. I was heavily involved in deciding upon what experiments were needed to be run to better understand or constrain the systematic errors. Sometimes the experiments were

simple functionality tests on specific pieces of apparatus, but more often the experiments involved taking large datasets with certain experimental parameters detuned from their ideal settings. Running these tests was a slow process, as it would typically take a few days to collect enough data just to see an effect.

By September 2010 we were finally ready to start taking data with the experiment configured for ‘normal running’. The first part of this thesis presents a detailed account of our 2011 electron EDM measurement as published in [75, 76]. The experiment is described in chapter 2, and chapter 3 details how we extract measurements from our raw detector signals. In chapter 4 a thorough analysis of the published dataset is presented. I discuss the various tests which I performed on the dataset to check for anomalies and inconsistencies. I also provide details of the supplementary experiments that were performed to calibrate and constrain the various systematic errors, and explicitly demonstrate how we associate a systematic uncertainty to each.

In parallel to running the EDM experiment, during the data acquisition phase I spent a lot of time working on upgrades to the experiment which would both improve EDM sensitivity and decrease the effects of systematic dependencies. After publishing in 2011, the focus of my PhD switched to the implementation and characterisation of these upgrades with the aim to make a new measurement at the 2×10^{-28} e cm level. Chapter 5 details the upgrades which were implemented, including new measurements of the systematic effects which will limit our next electron EDM measurement.

Before discussing the details of the YbF electron EDM experiment, it is worth providing a brief overview of the YbF molecule itself.

1.4 The YbF Molecule

In essence the ^{174}YbF molecule can be considered an ionically bound core of $^{174}\text{Yb}^{2+}$ and F^- , orbited by a single unpaired (valence) electron (which originated from the ^{174}Yb $6s^2$ orbital) [77, 78]. The ground electronic, vibrational and rotational molecular state, labelled $X^2\Sigma^+(v=0, N=0)$, is similar to an alkali atom. This state is split by the hyperfine interaction between the electron spin ($S=1/2$) and the fluorine nuclear spin ($I=1/2$) into a singlet ($F=0$) and triplet ($F=1$), separated by approximately 170 MHz. The isotope ^{174}Yb has no nuclear spin, so doesn’t contribute to the total angular momentum ($\vec{F}$).

The aim of our experiment is to measure the EDM induced energy splitting between two states of opposite electron spin. We measure this splitting between the $F=1$, $m_F=\pm 1$

levels (the stretched states where I and S are aligned) using a variant of Ramsey's separated oscillating fields interference method, as described in sections 2.1–2.2.

In order to prepare and readout the molecular state we drive $Q(0)$ transitions⁵ to the first excited electronic state labelled $A^2\Pi_{1/2}(v=0, N=0)$ with a 553 nm laser, as illustrated in figure 1.3. This state is also split into $F=0$ and $F=1$ levels, but the separation is so small (~ 3 MHz [36]) that we normally treat it as a single upper state. The A state has a lifetime of (28 ± 2) ns and a $v'=0 \rightarrow v''=0$ Franck-Condon factor of 0.933 ± 0.003 [79], so it almost instantly decays predominantly to the vibrational and electronic ground state, branching between the $N=0$ and $N=2$ rotational levels with 2/3 and 1/3 probabilities respectively^{6,7}. The $X^2\Sigma^+(v=0, N=2)$ state is considered to be a dark state as it has a lifetime which is much longer than the interferometer free evolution time period (see section 2.1). All population that falls into this state is lost to us and plays no role in the experiment.




FIGURE 1.3: The YbF energy levels which are relevant to the EDM experiment. Half of the population in $N=0, F=1$ is pumped into $N=0, F=0$. The other half falls into $N=2$.

⁵This notation is short hand for a transition from $N=0$ to an excited state where total angular momentum J (excluding the nuclear component I) is conserved, $\Delta J=0$.

⁶As calculated from Clebsch-Gordon coefficients. Reference [80] provides an excellent analysis of the coefficients of the molecule CaF, which apply equally to YbF.

⁷The laser population cannot decay into the $N=1$ level as this state has the opposite parity to $N=0$.

Chapter 2

The YbF Interferometer

2.1 Interferometer Overview




FIGURE 2.1: An illustration of spin precession, demonstrating the torques due to external electric and magnetic fields.

The essence of our EDM measurement is quite straightforward to understand. Magnetic and electric dipoles experience a torque when placed in static magnetic and electric fields respectively ($\vec{\mu} \times \vec{B}$ & $\vec{d} \times \vec{E}$). In general, these torques rotate the dipoles towards the field axis. However, if the dipoles are in an atom or molecule with angular momentum, $\vec{F}$, rather than rotating towards the field axis, the angular momentum and dipole moments actually precess about the field, as illustrated in figure 2.1. The angle of precession, ϕ , is directly proportional to the products $-\vec{\mu} \cdot \vec{B}$ & $-\vec{d} \cdot \vec{E}$. In our experiment we prepare a cloud of YbF molecules to have their total angular momentum perpendicular to the axis of aligned electric and magnetic fields. We then allow the molecules to precess about the fields over a certain period of time and measure the change in precession angle, when we reverse the orientation of $\vec{E}$ relative to $\vec{B}$. As the electric field is well defined, the change in angle provides a direct measure of the electric dipole moment, d .

Another way to view spin precession is as a phase difference that accumulates between angular momentum quantum states, due to an energy splitting induced by applied electric

and magnetic fields. We prepare YbF molecules into their electronic, vibrational and rotational ground state, $X^2\Sigma^+ (v=0, N=0)$, and measure the phase difference between the $F=1, m_F=\pm 1$ hyperfine states that accumulates as a result of the EDM and Zeeman interactions with parallel applied electric and magnetic fields. To measure the phase difference that is solely due to the EDM interaction, we repeatedly reverse the directions of the electric and magnetic fields. The electron EDM is then calculated by comparing the accumulated phase difference when the electric magnetic fields are aligned to when they are opposed.

To precisely measure the phase difference we use a quantum state interferometer that interferes the spin wave functions of the $F=1, m_F=\pm 1$ states. The interferometer utilises a pulsed beam of cold but fast moving molecules. We call each pulse of the molecular source a ‘shot’. In principle each shot provides a measure of the phase difference between the $m_F=\pm 1$ states. We can divide a shot into seven stages, defined by various times, t , after an initial trigger signal. The near uniform molecular velocity maps these times to positions along the beam line. The positions and their associated molecular state are illustrated in figure 2.2. Figure 2.2 also defines the experimental axes that will be referred to in the following sections.




FIGURE 2.2: Simplified overview of the EDM experiment annotated with the YbF molecular state at various positions within the machine.

Stage 1: Parameter Selection

At $t=0$ the physical state of all the experimental parameters are set. In the idealised scenario this is the direction of the applied electric field, and the direction and amplitude of the applied magnetic field. In reality many more parameters are adjusted.

Stage 2: Molecule Production

A small molecular packet is then produced with a forward velocity of approximately 600 m/s. At this speed the molecules take only 2 ms to fly the length of the machine. The source valve is pulsed at 25 Hz, so only one molecular packet is in the machine at any moment in time.

Stage 3: Optical Pumping

570 μ s after production the molecules pass through a 553 nm green laser beam propagating along x , that has a randomly chosen linear polarisation in the yz plane. This laser beam prepares the molecules into the $F = 0, m_F = 0$ hyperfine state of $X^2\Sigma^+ (v = 0, N = 0)$ by driving population in the $F = 1$ triplet up to the first excited electronic state, $A^2\Pi_{1/2} (v = 0, N = 0)$ (see figure 1.3). As previously stated, this state decays almost instantly into the $N = 0$ and $N = 2$ levels of the vibrational and electronic ground state. Any population which returns to the $N = 0, F = 1$ state repeats the cycle, so all population in this level becomes optically pumped out before the molecules pass through the full width of the laser beam. As the excited molecules decay they fluoresce with an intensity that is proportional to the number of molecules in the cloud. We measure this signal with a photo multiplier tube (PMT) and use it to normalise our probe signal against fluctuations in molecular flux.

After the molecular packet has passed through the pump laser beam, the molecules can be described by the quantum state:

$$|\psi_0\rangle = |F = 0, m_F = 0\rangle. \quad (2.1)$$

Stage 4: Split π -Pulse

The molecules then pass between two parallel metal plates which both apply an electric field and serve as a TEM transmission line for ≈ 170 MHz rf. The rf (magnetic field) is polarised along x (figure 2.2), so couples $|0, 0\rangle$ to $|1, \pm 1\rangle$, as defined in the z axis, which is parallel to the applied electric field. We apply an 18 μ s pulse of rf with its amplitude and frequency tuned to place the molecules in an equal superposition of the $|1, \pm 1\rangle$ states — a π -pulse¹.

$$|\psi_0\rangle \rightarrow |\psi_1\rangle = \frac{ie^{-i\omega\tau}}{\sqrt{2}} (|1, +1\rangle - |1, -1\rangle). \quad (2.2)$$

τ and ω are the rf pulse length and angular frequency respectively.

¹The significance of the angle π will become apparent in section 2.2 where Rabi flopping theory will be presented.

Stage 5: Phase Evolution

The molecules are then left to evolve in the static electric and magnetic fields, $(E_0, B_0) \hat{z}$, for a period, $T = 642 \mu\text{s}$. Typically $E_0 = \pm 10 \text{ kV/cm}$ and $B_0 = \pm 13.6 \text{ nT}$. During the evolution period the $|1, \pm 1\rangle$ states gain a phase difference, $\Delta\phi = 2\phi$, as a result of the Zeeman² and EDM interactions:

$$|\psi_1\rangle \rightarrow |\psi_2\rangle = \frac{i e^{-i(\omega_{11} T + \omega \tau)}}{\sqrt{2}} \left(e^{-i\phi} |1, -1\rangle - e^{+i\phi} |1, +1\rangle \right), \quad (2.3)$$

where ω_{11} is the mean separation between $|0, 0\rangle$ and $|1, \pm 1\rangle$, and the ‘interferometer phase’

$$\phi = (\mu_B B_0 - d_e E_{\text{eff}}) T / \hbar. \quad (2.4)$$

Stage 6: Recombination π -Pulse

The phase evolution period ends when a second $18 \mu\text{s}$ π -pulse is applied to the molecular cloud. This pulse is applied just before the molecules leave the parallel plates. As before, the rf pulse couples the state ψ_1 with ψ_0 but, due the rotation of the molecular angular momentum about the z axis, the molecular state, ψ_2 , now contains a component, $\frac{i}{\sqrt{2}} (|1, +1\rangle + |1, -1\rangle)$, which doesn’t couple to the rf. Therefore, after the second π -pulse, the molecules reside in a coherent superposition of $F = 0$ and $F = 1$ states:

$$|\psi_2\rangle \rightarrow |\psi_3\rangle = -e^{iT(\omega - \omega_{11})} \cos \phi |0, 0\rangle - \frac{e^{-i\omega_{11}(T + \tau) - i\omega \tau} \sin \phi}{\sqrt{2}} (|1, +1\rangle + |1, -1\rangle). \quad (2.5)$$

Stage 7: Population Readout

Approximately 2.2 ms after initial production, the molecules pass through a linearly polarised probe laser beam propagating along x . As in stage 3, this laser drives the $Q(0)$ electronic transition, but here the laser is tuned to measure population in the $F = 0$ hyperfine level. A second (probe) PMT is positioned to collect the laser induced fluorescence (LIF) and detects a signal proportional to

$$|\langle 0, 0 | \psi_3 \rangle|^2 = \cos^2 [(\mu_B B_0 - d_e E_{\text{eff}}) T / \hbar]. \quad (2.6)$$

This PMT signal is recorded, and the experiment returns to stage 1 to make another measurement with different experimental parameter settings.

²In our system the Landé g-factor, $g_F = 1$.

2.1.1 Extracting an EDM




FIGURE 2.3: An illustration of the sampled points which are used to extract an EDM value. The two curves represent the two orientations of the electric field relative to the z axis. An EDM causes a phase shift between these lines, which is vastly exaggerated here. Four magnetic field values are applied in each electric field switch state.

Figure 2.3 shows the dependence of the probe PMT signal on the applied magnetic field, B_0 . As shown, reversing the direction of the applied electric field, E_0 , shifts the phase of the fringes by a small amount,

$$\delta\phi_{\text{EDM}} = 2d_e E_{\text{eff}} T/\hbar. \quad (2.7)$$

This shift causes a small change in detector signal $\delta S_{\text{EDM}} = (\partial S/\partial\phi)\delta\phi_{\text{EDM}}$ that is correlated with the sign of $-\vec{E} \cdot \vec{B}$. Note that the change in signal is opposite on either side of the central fringe. We apply the magnetic field $B_0 = \pm 13.6 \text{ nT}$ to operate near to $\phi = \pm\pi/4$ where the interferometer is most sensitive to phase changes, and where signal changes are linear in phase. For each orientation of the applied electric and magnetic fields we calibrate the gradient $(\partial S/\partial\phi)$ by making a small phase step, $\delta\phi_{\text{CAL}} = 2\pi/32$, with the magnetic field, $B_{\delta B} = \pm 1.7 \text{ nT}$. We measure the corresponding change in detector signal, δS_{CAL} . As illustrated in figure 2.3, the eEDM phase is measured by combining the (normalised) probe signal at points $(A, B, C, D, A', B', C', D')$ according to:

$$\delta\phi_{\text{EDM}} = \delta\phi_{\text{CAL}} \frac{\delta S_{\text{EDM}}}{\delta S_{\text{CAL}}} = \delta\phi_{\text{CAL}} \frac{(A' + B' + C + D) - (A + B + C' + D')}{(B + B' + C + C') - (A + A' + D + D')}. \quad (2.8)$$

In addition to reversing the direction of the applied electric and magnetic fields, and modulating the magnitude of the magnetic field, we also modulate six other experimental parameters between shots. We modulate: the laser frequency, the amplitude and frequency of each rf pulse, as well as the phase difference between the two rf pulses. The additional parameter modulations allow us to measure how the signal correlates with any switched parameter, or combination thereof. Full details will be given in section 3.1.

With nine parameter modulations there are 512 unique machine states. An individual EDM measurement, or ‘block’, is made from 4096 shots over which each unique machine state is visited eight times. Each block takes approximately six minutes to acquire. Typically ~ 150 blocks are collected before acquisition is interrupted for laser or target maintenance, or to manually reverse electrical connections. We call an uninterrupted collection of blocks a ‘cluster’ and we fold all the clusters together to make a complete ‘dataset’. We require $\sim 5,000$ blocks to reduce the EDM uncertainty to the 5×10^{-28} e cm level.

2.2 Interferometer Theory

I will now provide a more rigorous mathematical account of the interference signal, so that we may have a better understanding of how it depends on all the various experimental parameters. This will allow expressions for physical values to be written in terms of measured detector signals (see appendix A). It will also give an insight into how some systematic effects are accounted for in the final result. Note, however, that the model presented here will still be limited by a number of approximations. In section 3 focus will switch to a much more empirical approach. To establish systematic dependencies we exhaustively test how our final result depends on each experimental parameter. There is a balance between practical empirical tests and simple theoretical models, one in which we lean more heavily towards the former rather than the latter.

2.2.1 Evolution In Static Fields




FIGURE 2.4: Relative Stark shifts of the ground state hyperfine levels. Note the very different vertical scales.

If left unperturbed, the magnetic sub levels of the $F = 1$ hyperfine state remain degenerate and sit 170.254 MHz above the $F = 0$ level. The predominant effect of applying a static electric field is to polarise the YbF molecule, aligning the internuclear axis to lie along

the applied field. In effect this alignment mixes the molecular rotational levels, causing the $N = 0$ level to shift downward, as was measured by Sauer *et al.* [36, 78]. On top of the rotational level shift there is a relatively small shift between the $F = 0$ and $F = 1$ hyperfine states, as well as a tensor splitting between the $m_F = 0$ and $m_F = \pm 1$ states of the $F = 1$ triplet, as shown in figure 2.4.

Fortunately F and m_F are still good quantum numbers of the system in an electric field [78], so we may write a diagonal effective Stark shifted Hamiltonian for the z basis states $\{|F = 0, m_F = 0\rangle, |1, -1\rangle, |1, 0\rangle, |1, +1\rangle\}$, introducing the hyperfine dependent Stark shifts as phenomenological parameters. We can ignore the higher lying molecular states as they play no role in the interferometer except during the state preparation and readout phases, which are incoherent processes driven by laser beams.

The whole aim of the experiment is to measure the energy shift due the permanent electric dipole moment of the electron. To account for this we add the tiny perturbation $\hat{H}_d = -d_e E_{\text{eff}} \hat{F}_z$. Sections 1.2 and 1.3.1 provide a thorough discussion of this interaction Hamiltonian.

We also apply a magnetic field along the same axis as the electric field which perturbs the spin states according the Hamiltonian $\hat{H}_Z = -\vec{\mu} \cdot \vec{B} = g\mu_B B \hat{F}_z$. The g -factor is very nearly 1 and μ_B is, as usual, the Bohr magneton. Writing this interaction in terms of the diagonal Pauli matrix $\hat{F}_z$ represents the weak field limit where we neglect off diagonal couplings. We may do this as the shift $\mu_B B$ is a million times smaller than the $F = 1$ to $F = 0$ level separation.

Combining the above three interactions we write the following Hamiltonian for the rovibrational ground state hyperfine levels $|F, m_F\rangle$ in static electric and magnetic fields $(E_0, B_0)\hat{z}$:

$$\hat{H}_{\text{sf}} = \omega_{10}\hbar|1, 0\rangle\langle 1, 0| + (\omega_{11} - \Delta)\hbar|1, -1\rangle\langle 1, -1| + (\omega_{11} + \Delta)\hbar|1, +1\rangle\langle 1, +1|. \quad (2.9)$$

ω_{10} and ω_{11} are the (quadratic) Stark shifted separations³ as defined in figure 2.5. $\Delta = g\mu_B B_0 - d_e E_{\text{eff}}$ represents the Zeeman and EDM energy shift combined.

For the general molecular state $|\psi(t)\rangle_z$, written in matrix form in the z basis with increasing F and m_F , the static field Hamiltonian produces the following time evolution operator:

$$\hat{\Pi}_{\text{sf}}(T, \Delta) = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & e^{-i(\omega_{11}-\Delta)T} & 0 & 0 \\ 0 & 0 & e^{-i\omega_{10}T} & 0 \\ 0 & 0 & 0 & e^{-i(\omega_{11}+\Delta)T} \end{pmatrix}, \quad (2.10)$$

³The ‘tensor Stark shift’ is thus defined as $\omega_{11} - \omega_{10}$.




FIGURE 2.5: The ro-vibrational ground state hyperfine level shifts due to static electric and magnetic fields aligned along z . The states are represented in the z basis.

defined such that $|\psi(t+T)\rangle_z = \hat{\Pi}_{\text{sf}}(T, \Delta) |\psi(t)\rangle_z$, where T is the interaction time.

2.2.2 RF Pulse Theory

The interferometer begins and ends with a carefully tuned pulse of rf. The rf magnetic field of each pulse is linearly polarised along the x axis, so both pulses drive a $\Delta m_F = 0$ transition in the x basis. It is therefore convenient to transform the z basis set into the $m_F = 0$ states of the x , y and z bases. Using the standard rotation matrix defined by Weissbluth [81] it can be shown that the $m_F = 0$ component of the x and y bases, labelled $|x\rangle$ and $|y\rangle$ respectively, can be related to the z basis states as $|x\rangle = \frac{1}{\sqrt{2}}(|1, +1\rangle - |1, -1\rangle)$ and $|y\rangle = \frac{i}{\sqrt{2}}(|1, +1\rangle + |1, -1\rangle)$. In accordance, we define the transformation matrix

$$\hat{U} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & -\frac{1}{\sqrt{2}} & 0 & \frac{1}{\sqrt{2}} \\ 0 & 0 & 1 & 0 \\ 0 & \frac{i}{\sqrt{2}} & 0 & \frac{i}{\sqrt{2}} \end{pmatrix}, \quad (2.11)$$

which rotates the z basis set, written with increasing F and m_F , into the ‘Cartesian’ basis⁴ $\{|0, 0\rangle, |x\rangle, |1, 0\rangle, |y\rangle\}$, as defined above. Working in this basis effectively reduces the rf interaction to a coupling between two quantum states, which allows the rf interaction Hamiltonian to be written as $\hat{H}_{\text{rf}} = g\mu_B B_x^{(\text{rf})} \cos(\omega t + \Phi_{\text{rf}}) (|0, 0\rangle \langle x| + |x\rangle \langle 0, 0|)$. The coupling between $|1, 0\rangle$ and $|x\rangle$ is small enough to ignore.

In this basis we may describe the hyperfine system with the wave function

$$|\psi(t)\rangle_c = (a_{00}(t), a_x(t), a_{10}(t), a_y(t)), \quad (2.12)$$

⁴This nomenclature was first used in Hudson’s thesis [71].

where $a_i(t)$ define the time dependent amplitudes of each basis state. This molecular state evolves according to the time dependent Schrödinger equation with the Hamiltonian $\hat{H} = \hat{U}\hat{H}_{\text{sf}}\hat{U}^\dagger + \hat{H}_{\text{rf}}$. Ideally the rf is tuned close to resonance so that $|\omega_{11} - \omega| \ll \omega_{11}$. In this case we may make the standard rotating wave approximation which neglects exponentials that evolve at the relatively rapid rate $\omega_{11} + \omega \approx 2\omega_{11}$ and essentially integrate to zero. Explicitly, the Schrödinger equation becomes

$$i\hbar \frac{d}{dt} |\psi(t)\rangle_c = \begin{pmatrix} 0 & \frac{\Omega\hbar}{2} e^{i(\omega t + \Phi_{\text{rf}})} & 0 & 0 \\ \frac{\Omega\hbar}{2} e^{-i(\omega t + \Phi_{\text{rf}})} & \omega_{11}\hbar & 0 & -i\Delta\hbar \\ 0 & 0 & \omega_{10}\hbar & 0 \\ 0 & i\Delta\hbar & 0 & \omega_{11}\hbar \end{pmatrix} |\psi(t)\rangle_c, \quad (2.13)$$

where the Rabi frequency $\Omega = g\mu_B B_x^{(\text{rf})}/\hbar$. Under typical running conditions $g\mu_B B_x^{(\text{rf})} \gg \Delta$, so we may make the approximation that $\Delta = 0$. Substituting equation 2.12 into equation 2.13 leads to the following equations for amplitudes $a_i(t)$:

$$\ddot{a}_{00}(t) - i\delta \dot{a}_{00}(t) + \frac{1}{4}\Omega^2 a_{00}(t) = 0, \quad (2.14a)$$

$$\dot{a}_{00}(t) + \frac{i}{2}\Omega a_x(t) e^{i(\omega t + \Phi_{\text{rf}})} = 0, \quad (2.14b)$$

$$\dot{a}_{10}(t) + i\omega_{10} a_{10} = 0, \quad (2.14c)$$

$$\dot{a}_y(t) + i\omega_{11} a_y = 0, \quad (2.14d)$$

where $\delta = \omega - \omega_{11}$ is the rf detuning from resonance. Equation 2.14a has solutions of the form $a_{00}(t) = e^{\frac{1}{2}i\delta t} (p_1 \cos(\frac{Wt}{2}) + p_2 \sin(\frac{Wt}{2}))$. p_1 and p_2 are constants which can be related to the initial populations $a_{00}(0)$ and $a_x(0)$, and $W = \sqrt{\Omega^2 + \delta^2}$ is the generalised Rabi frequency.

With some not-so-trivial algebra, equations 2.14a–2.14d may be integrated to give the time evolution operator for an rf pulse of length τ and initial phase Φ_{rf} :

$$\hat{\Pi}_{\text{rf}}(\tau, \Phi_{\text{rf}}) = \begin{pmatrix} (c - i\frac{\delta}{W}s) e^{\frac{1}{2}i\delta\tau} & -i\frac{\Omega}{W}s e^{\frac{1}{2}i\delta\tau} e^{i\Phi_{\text{rf}}} & 0 & 0 \\ -i\frac{\Omega}{W}s e^{-i\Phi_{\text{rf}}} e^{i\tau(\frac{1}{2}\delta - \omega)} & (c + i\frac{\delta}{W}s) e^{\frac{1}{2}i\tau\delta} e^{-i\omega\tau} & 0 & 0 \\ 0 & 0 & e^{-i\tau\omega_{10}} & 0 \\ 0 & 0 & 0 & e^{-i\tau\omega_0} \end{pmatrix}, \quad (2.15)$$

where $c = \cos(\frac{1}{2}\tau W)$ and $s = \sin(\frac{1}{2}\tau W)$, following the nomenclature of Ramsey [55]. This operator acts to give $|\psi(\tau + t)\rangle_c = \hat{\Pi}_{\text{rf}}(\tau, \Phi_{\text{rf}}) |\psi(t)\rangle_c$.

If at $t = 0$, $|\psi(0)\rangle_c = |0, 0\rangle$, it is straightforward to show that the populations $|a_{00}(\tau)|^2$ and $|a_x(\tau)|^2$ oscillate in both pulse length (τ) and effective amplitude (W):

$$|a_{00}(\tau)|^2 = 1 - \frac{\Omega^2}{W^2} \sin^2 \left(\frac{1}{2} \tau W \right), \quad (2.16a)$$

$$|a_x(\tau)|^2 = \frac{\Omega^2}{W^2} \sin^2 \left(\frac{1}{2} \tau W \right). \quad (2.16b)$$

This is where the π -pulse nomenclature arises — when $\delta = 0$ and $\Omega \tau = \pi$ there is a 100% coherent population transfer from $|0, 0\rangle$ into $|x\rangle$. When $\delta \neq 0$, the transfer still peaks when $W\tau = \pi$, but now the transfer is incomplete as $\Omega^2/W^2 \neq 1$. If we set $\Omega \tau = \pi$, as a function of detuning the transition lineshape is a sinc^2 function with a linewidth of $\approx 1/\tau$ (FWHM). Under usual running conditions $\tau \approx 18 \mu\text{s}$, so the transition linewidth is $\approx 30 \text{ kHz}$. Given that the 170 MHz rf is far detuned from the $\approx 5 \text{ MHz}$ resonance between $|1, 0\rangle$ and $|x\rangle$, the relatively narrow linewidth justifies why we may neglect any couplings between these two states.

2.2.3 Interference Lineshape

By combining the static field and rf time evolution operators, we derive the interferometer equation as follows:

$$|\psi(\tau_2 + T + \tau_1)\rangle_c = \hat{\Pi}_{\text{rf}}(\tau_2, \omega \tau_1 + \omega T + \Phi_{\text{rf}}) \hat{U} \hat{\Pi}_{\text{sf}}(T, \Delta) \hat{U}^\dagger \hat{\Pi}_{\text{rf}}(\tau_1, 0) |\psi(0)\rangle_c. \quad (2.17)$$

The second rf pulse begins with the phase $\omega(\tau_1 + T)$, as in the lab the generated rf is phase continuous. The additional factor Φ_{rf} accounts for an arbitrary phase difference between the two rf pulses that we are able to introduce. The reason for this phase will become clear shortly.

Let us assume that the optical pumping stage is perfect so that $|\psi(0)\rangle_c = |0, 0\rangle$. The possibility that $F = 1$ might remain slightly populated is addressed in section 5.2.5. It can be shown that the final population in $F = 0$ is given by

$$S = |\langle 0, 0 | \psi(\tau_2 + T + \tau_1) \rangle_c|^2 = S_C + S_I + S_R, \quad (2.18)$$

where

$$S_C = \left(c_1^2 + \frac{\delta_1^2}{W_1^2} s_1^2 \right) \left(c_2^2 + \frac{\delta_2^2}{W_2^2} s_2^2 \right), \quad (2.19a)$$

$$S_I = \left(\frac{\Omega_1 \Omega_2}{W_1 W_2} s_1 s_2 \cos(\phi) \right)^2, \quad (2.19b)$$

$$S_R = 2 \frac{\Omega_1 \Omega_2}{W_1 W_2} s_1 s_2 \cos(\phi) \left[\left(\frac{\delta_1}{W_1} s_1 c_2 + \frac{\delta_2}{W_2} s_2 c_1 \right) \sin(\vartheta) + \left(\frac{\delta_1 \delta_2}{W_1 W_2} s_1 s_2 - c_1 c_2 \right) \cos(\vartheta) \right]. \quad (2.19c)$$

$$c_i = \cos\left(\frac{1}{2}\tau_i W_i\right), \quad s_i = \sin\left(\frac{1}{2}\tau_i W_i\right), \quad \vartheta = \delta_{\text{evo}} T + \Phi_{\text{rf}}, \quad \phi = (g\mu_B B_0 - d_e E_{\text{eff}}) T / \hbar.$$

The subscript, i , refers to each rf pulse (1 or 2), and $\delta_{\text{evo}} = \frac{1}{T} \int_{\tau_1}^{\tau_1+T} (\omega(t) - \omega_{11}(t)) dt$ is the average rf detuning during the free evolution period. The first term, S_C , depends only on the tuning of the rf pulse parameters and not on interferometer phases ϕ and ϑ . S_I is the part of the signal which we want to measure, containing the combined EDM and Zeeman phase, ϕ . Under ideal conditions ($W_i \tau_i = \Omega_i \tau_i = \pi$) the interferometer signal, S , depends only on $\cos^2(\phi)$. As the rf pulses deviate from ideal tuning (in either power or frequency) the $\cos^2(\phi)$ term is attenuated and the interferometer signal becomes dependent on a second term S_R which is a function of both ϑ and ϕ . We call this the Ramsey interference term as it depends upon the rf phase and measures the splitting between $|0, 0\rangle$ and $|x\rangle$, as is usual for a Ramsey interferometer. Ramsey interference is maximised when $\Omega\tau = \pi/2$, i.e. with $\pi/2$ -pulses.

The Ramsey term is a potential source of systematic error as it depends upon the applied electric field through the Stark shift. We ensure this component averages to zero by adding a phase difference $\Phi_{\text{rf}} = \Phi_{\text{rf}} \pm \pi/2$ between the two rf pulses that switches according to a shot-shot switch sequence. We also randomise the constant phase offset, Φ_{rf} , in between blocks as a second line of defence.

A number of approximations were made in deriving the expression for the interferometer signal. All incoherent processes such as spontaneous decay were neglected, as excited state lifetimes are much longer than the time it takes the molecular cloud to fly through the beam machine. The rotating wave approximation was also applied as it was assumed that the rf is tuned close to resonance. We also neglected the influence of magnetic fields during each rf pulse. It is relatively straightforward to modify this theory so that some or all of these approximations are not made. We can also extend this model to account for other imperfections such as rf frequency chirps and polarisation rotations, and introduce possible perpendicular static field components. We generally run these simulations numerically rather than derive equations which would no doubt be considerably more complex than equations 2.19a-2.19c. The details of these simulations will be given in chapter 4 when relevant.

2.3 Apparatus

I will now describe the experimental apparatus in more detail as it was during acquisition of the dataset that was published in May 2011 [75].

2.3.1 Vacuum Chamber




FIGURE 2.6: A schematic of the YbF molecular beam machine. Not drawn to scale.

As illustrated in figure 2.6, the interferometer resides in a vertically orientated vacuum chamber which is approximately 2 m tall and 25 cm in diameter. The chamber is

separated into two sections with a gate valve to isolate the upper chamber from the lower source chamber during maintenance of the molecular source. This isolation is necessary to keep the upper chamber clean in order to prevent electric field break down.

Inside the upper chamber an ‘inner’ cylindrical magnetic shield surrounds a pair of parallel plates for applying electric and rf fields as well as a set of coils for producing controlled magnetic fields. Also within the upper chamber, yet outside the ‘inner’ shield, reside the optics for collecting laser induced fluorescence onto a pair of photomultiplier tubes (one for each laser beam). The PMTs themselves sit outside the main vacuum chamber. A schematic of the upper chamber layout is shown in figure 2.7.




FIGURE 2.7: A schematic of the upper vacuum chamber. Not drawn to scale.

The source chamber houses the pulsed gas inlet valve and the target assembly. Molecules propagate upwards from here, through a skimmer, into the upper chamber. Two turbomolecular pumps keep the pressure within this chamber typically below 5×10^{-5} Pa as is necessary to limit collisional decoherence [70]. A further two turbo pumps evacuate gas introduced to the source chamber by the gas inlet valve, maintaining a pressure of approximately 10^{-2} Pa whilst the beam machine is running.

2.3.2 YbF Beam Source

A sensitive measurement of the EDM requires a stable source that can produce a large flux of YbF molecules in their ground state. For this reason we continue to use the pulsed supersonic source as described extensively in the thesis of Condylis [72].

YbF is produced through the laser ablation of Yb metal in the presence of a supersonically expanding mixture of 98% Ar and 2% SF₆ gas. As the gas isentropically⁵ expands into the vacuum chamber, its thermal energy is converted into forward kinetic energy, resulting in a fast moving jet with a narrow velocity spread. A focused Nd:YAG laser creates a hot plume of Yb metal which reacts with the SF₆ to form, amongst other things, YbF radicals. The YbF becomes entrained in the Ar carrier gas and thermalises with it, cooling the YbF molecules to around 4 K [82]. We measure the YbF molecules to have a forward (vertical) molecular velocity of 590 ms^{-1} , with a velocity spread of 40 ms^{-1} (FWHM). A skimmer with a 2 mm aperture placed 8 cm from the target collimates the YbF beam before it passes into the upper chamber.

A schematic diagram of the source chamber is given in figure 2.8. The molecular beam is pulsed at 25 Hz via a commercially available solenoid valve with a 4×10^5 Pa carrier gas backing pressure. Whilst we could run the source at 50 Hz we find that, at the more rapid rate, the increased gas pressure reduces interference contrast so there is no overall gain in experimental sensitivity. The valve is typically pulsed with a $180 \mu\text{s}$ high voltage (380 V) drive triggered by a digital signal from the computer. After a short $350 \mu\text{s}$ delay, 1064 nm light from a Q-switched Nd:YAG laser is fired at the Yb target. This is a 1 mm thick, 3 mm wide ribbon of Yb metal mounted with epoxy on the edge of a rotatable 12 cm diameter stainless steel disc. The focussed laser spot is ~ 5 mm in diameter and the target disc is positioned a few mm above and behind the exit aperture of the gas valve; we find molecular signal to be critically dependent on this exact positioning⁶. Laser pulse length is fixed at 12 ns, however the maximum pulse energy can be controlled by varying

⁵No heat is added to the system during expansion $\Delta Q = 0$.

⁶Optimising this position is a painful process as we must break vacuum in the source chamber for every adjustment.





FIGURE 2.8: A schematic of the inlet valve and the rotating ytterbium target assembly. Not drawn to scale.

the time delay between the flashlamp and Q-switch digital triggers. Typically the flash-to-Q interval is set to $\sim 280 \mu\text{s}$, but the optimal delay depends upon the age of the laser flashlamps and the condition of the target surface, which degrades fairly rapidly. After ~ 5000 YAG shots the (initially silver) target turns completely black⁷ and the molecular signal dies to almost nothing. To rejuvenate the signal a stepper motor rotates the target disc by a few degrees when the signal drops below a certain value to expose fresh Yb. This gives the target a lifetime of approximately 100 hours before we must break vacuum in the lower chamber in order to scrape it clean with wire wool or a razor blade.


Typically molecular pulses contain $\sim 10^5$ molecules, however we find the intensity and stability of the YbF beam to be quite dependent on a number of tunable parameters including flash-to-Q time, valve-to-Q time and valve-pulse-length. We optimise the molecular signal measured by the probe PMT every week by scanning the tunable parameters as shown in figure 2.9. Each parameter is roughly independent so we can optimise each individually whilst holding the other parameters constant. As in figure 2.9(a), often the most appropriate value is chosen to be just below the maximum signal.

⁷We assume the black residue to be a sulphurous compound but its true composition is unknown.


This is because we find relatively intense molecular signals tend to decay more rapidly and become noisy which reduces EDM sensitivity in the long run.



(a) The Flash-to-Q interval is a measure of the Nd:YAG laser power. The output power decreases as the interval increases. Note there is an optimum setting just before a rapid drop-off. The position of the rapid fall moves to shorter times as the laser flashlamps age.



(b) The valve-to-Q time measures the delay between firing the Nd:YAG laser and opening the gas inlet valve. The optimum interval $\sim 355 \mu\text{s}$ doesn't tend to change much.



(c) Valve pulse length measures the time the gas inlet valve is held open. The optimum value also depends on the valve drive voltage. This scan was taken with a 380 V drive, which is the usual running value. A large drive voltage lets more gas through.

FIGURE 2.9: Example set of source parameter scans. Note the smooth lines serve only to guide the eye.

2.3.3 Pump and Probe Laser System

Stabilisation and Control of the Coherent 899 Dye Laser

The pump and probe laser beams are derived from a single Coherent 899 ring cavity dye laser pumped by an Ar^+ laser lasing on all of its visible lines. We pump Rhodamine 110 dye with 2–6 W of power (depending on the age of the dye) to give 120 mW output at 553 nm. The pump light (resonant with the $Q(0)$ transition from $F = 1$) is generated by double passing part of the laser output through an acousto-optic modulator (AOM) that red shifts the light by roughly 170 MHz. We call this the ‘pump AOM’.

Whilst the laser linewidth is less than 1 MHz, its output frequency can drift by up to 10 MHz/h. It therefore could drift, in one hour, by an amount comparable to the

(~ 20 MHz) natural linewidth of the $Q(0)$ transition. We require the laser to be locked to within ~ 2 MHz for periods of up to 12 hours at a time, so we lock the dye laser reference cavity to an iodine (I_2) spectral feature using saturated absorption spectroscopy as described in full in the thesis of Hudson [71]. This requires approximately 70 mW of laser power which is split off from the main laser beam before it is divided to form the pump and probe beams.

The iodine spectral features lie 120 MHz blue detuned from the $Q(0)(F = 0)$ transition, so we shift the so called ‘strong pump beam’⁸ of the I_2 lock by double passing it through an AOM driven at 120 MHz, with a 12 MHz deep frequency modulation superimposed (which oscillates at 100 kHz). We modulate the frequency to provide an error signal proportional to the derivative of the I_2 spectral features.

We can reliably scan the laser linearly over a narrow ~ 20 MHz range by supplying the voltage controlled oscillator (VCO) that drives the 120 MHz I_2 lock AOM with a 0–5 V analogue offset. We servo this frequency to optimise molecular signal and interference contrast.

The I_2 lock is quite stable but is susceptible to jump spectral features if it is jolted by other activity in the lab. On a good day the dye laser only needs human attention roughly every 12 hours as laser power drops with ageing dye. Normally the dye will last 80 hours before it needs replacing with a fresh mix.

Pump and Probe Beam Optics




FIGURE 2.10: A schematic of the pump/probe laser optics train for arbitrary polarisation rotation about the x axis. The red arrows indicate the (green) laser polarisation.

⁸Not to be confused with the pump laser beam used to optically pump the YbF molecules out of $F = 1$.

Before the pump and probe laser beams are sent into the vacuum chamber, each beam is passed through a short optics train which linearly polarises each laser beam and allows arbitrary, independent polarisation rotation about the x axis. Under ideal circumstances the molecular signal should not depend on the polarisation angle, but rotating the polarisation provides us with a valuable test for systematic effects.

The optical setup is shown in figure 2.10. This arrangement is replicated for both the pump and probe laser beams. First, a polarising beam splitter (PBS) ensures the laser is, to a good degree, linearly polarised. This component also deflects a small amount of light onto a photodiode to monitor the power in each beam; its output signal is logged by the main computer. The laser beam then passes through a quarter wave plate (QWP) oriented with its fast axis at 45° to the initial beam cube axis. This circularly polarises the beam, ensuring uniform power in all axes about x . The beam is then linearly polarised again during passage through a second, high quality *Thorlabs GT10-A* Glan-Taylor cube mounted on a computer controlled, motorised rotation stage.

Unfortunately as the QWP is not perfect, the light becomes slightly elliptically polarised before it encounters the Glan-Taylor cube. The Glan-Taylor cube converts this ellipticity into a power fluctuation which has been measured to be no more than a 5% variation. This should not pose a problem as we ensure that there is more than enough power in the pump beam to deplete the $F = 1$ population. Furthermore, we only rotate laser polarisation in between blocks; recall from section 2.1.1 that probe PMT signal changes are calibrated for each individual EDM measurement.

Both laser beams enter the vacuum chamber through a flat anti-reflection coated window. It is possible that these windows could introduce a certain amount of circular polarisation. By placing a second Glan-Taylor cube after each window and rotating it through 360° , we have determined that, in both cases, the light remains linear to better than 5 parts in 10^4 .

2.3.4 Laser Induced Fluorescence Detection

Figure 2.7 illustrates the optical setup for the detection of laser induced fluorescence (LIF). The pump and probe laser beams pass through the vacuum chamber 0.35 m and 1.35 m above the source valve respectively. As previously stated, these beams drive the $Q(0)$ transition from the $F = 1$ and $F = 0$ levels of the rovibrational ground state, causing the molecules to fluoresce isotropically as they pass through each laser beam. The signal intensity is proportional to the population in each state being probed.

The molecular cloud expands slightly as it moves through the vacuum chamber. In the pump region the cloud is only a few mm wide, so we shape the pump laser beam to have a 4 mm wide Gaussian profile to ensure it intersects all molecules. At the top of the chamber the cloud is considerably more spread out. An aperture at the top of the inner shield limits its dimensions to be 40 mm x 15 mm in the x and z axes respectively. We shape the probe laser beam into a top hat profile, spanning 5 mm along y and 10 mm along z , to maximise signal whilst minimising background laser scatter.

In both the pump and probe regions, we collect fluorescence onto a single PMT: a *Hamamatsu R5070A* collects pump fluorescence whereas a *Hamamatsu R669* PMT collects the probe fluorescence. The quantum efficiency of these detectors is, at best, 10% and unfortunately ray tracing simulations have shown that the probe collection optics have an overall collection efficiency of just under 5% [83]. We, therefore, have much room for improvement with regards to detection efficiency.

We use the pump signal to normalise the probe signal against variations in molecular flux. This does not require a particularly large signal from the pump PMT. However, EDM sensitivity is critically dependent on the probe signal, which not only depends upon probe laser beam alignment, but also optical pumping efficiency. For this reason the power and profile of the pump laser beam are chosen to ensure all molecules in $F = 1$ are pumped out before they pass through the full width of the laser beam. We also tweak the pump laser beam alignment and frequency to maximise the probe PMT signal, rather than pump PMT signal⁹. Typically, optical pumping increases the probe signal by a factor of two.

The molecular cloud has a 40 m/s longitudinal velocity spread, so it disperses along the y axis as it travels up the machine. We operate the PMTs in current mode and connect their output to an amplifier which converts the photo-current into a voltage. This gives a temporal resolution of $\sim 1 \mu\text{s}$ which enables us to observe near Gaussian shaped time of flight (TOF) profiles from each PMT as the molecular cloud passes through the laser beams. An example probe PMT TOF profile is shown in figure 2.11.

We can analyse the TOF profiles in a number of ways. Most often we simply gate out the central region, as shown, and integrate the sample data to produce a single value for each molecular pulse. However, we can retain the temporal information should we want to investigate possible velocity or, more likely, position dependent effects. See sections 3.1–4.6 for a full description of how we analyse the PMT signals.

⁹In an ideal world the maximum of the pump and probe PMTs would coincide, however the optical arrangements are not perfect so it is possible to increase pump PMT signal at the cost of probe fluorescence detected




FIGURE 2.11: Fluorescence measured by the probe PMT as a function of arrival time relative to the YAG Q-switch. We call this a time of flight (TOF) profile. The points represent the average signal from 100 shots. The red line shows a fit to a Gaussian curve. When analysing EDM data we typically integrate only the signal within the shaded region.

2.3.5 Applied RF Fields

Once the pump laser beam has depleted population in $F = 1$, we are able to coherently manipulate the quantum state of the YbF molecules by driving transitions between the $F = 0$ and $F = 1$ hyperfine manifolds with short pulses of ~ 170 MHz rf.

The parallel plate structure housed within the upper chamber is used both to produce the static electric field, and to serve as a $37\,\Omega$ transmission line to carry rf in a TEM mode, with the oscillating magnetic field linearly polarised along the x axis. This gives us the ability to apply rf to the molecules at any position along the plate length. The rf transmission line is symmetric, which allows us to propagate rf in either direction along the beam line to test for possible systematic effects.

The parallel plate transmission line is coupled to $50\,\Omega$ semi-rigid coaxial cables at either end through $220\,\text{pF}$ DC isolating capacitors which connect to the plates in opposite corners as shown in figure 2.12. Field mapping studies have shown that this unfortunately causes the rf polarisation to rotate towards the coupling point at either end (see section 5.1.1 and [74]). We therefore only use the central 40 cm of plate length under normal running conditions.

The $37\,\Omega$ – $50\,\Omega$ impedance jumps are bridged with a pair of $10\,\text{pF}$ trimming capacitors, which were empirically chosen to improve rf transmission through the coupling points at 170 MHz. Impedance matching is not quite perfect so a small amount of rf is reflected at either end. Using the molecules as probes (see below), we measure a voltage standing wave ratio (VSWR) of 1.5. This does not cause a major problem as we can independently tune the power of each rf pulse to satisfy the π -pulse condition.

A schematic of the rf generation system is also shown in figure 2.12. We protect this system from high voltage and discharges with a pair of bias-tees and lightning arrestors. As shown, both the ‘split’ and ‘recombination’ rf pulses are generated by the same *HP 8657A* rf synthesizer. The synthesizer is programmed by the computer over GPIB. This process takes a few milliseconds, which is longer than the $640\ \mu\text{s}$ interferometer phase evolution time. To create two short pulses (typically $18\ \mu\text{s}$ in length), we pass the rf through two fast *Mini-Circuits ZASWA-2-50DR+* switches which can switch in $<1\ \mu\text{s}$ and give a combined isolation of $> 100\ \text{dB}$. When the rf is not being pulsed, these switches redirect the synthesizer output into a *HP 53131A* frequency counter and a *HP 438A* power meter, which are logged once per block by the main computer.

A slight warping of the parallel plates (see sections 2.3.6 and 5.1.2) and the rf standing wave require that the frequency and amplitude of each pulse be controlled independently. To address this problem, we use the internal DCFM of the synthesizer to independently tune the frequency of each pulse. In addition we pass the rf through a *Mini-Circuits ZX73-2500* rf attenuator to tune pulse amplitude. Both of these components can respond on a $10\ \mu\text{s}$ time scale.

We control the DCFM and the attenuator with analogue voltages which are generated by the computer before each shot. A set of SPDT switches, triggered by digital signals,




FIGURE 2.12: Schematic of rf system.

then select the appropriate voltage as the molecules fly through the beam machine. The DCFM is tuned over 20 kHz with a 0–1.2 V drive ($V_{\text{RF}i\text{F}}$), and the attenuator decreases rf amplitude by up to 40 dB as its control voltage ($V_{\text{RF}i\text{A}}$) is decreased from 17 V down to 0 V. As in the interferometer theory section, the subscript i indexes the two rf pulses. The first pulse ($i = 1$) is always the pulse which the molecules experience first. It is independent of the rf propagation direction.




FIGURE 2.13: Example set of rf parameter scans. The red lines are fits to the physical models presented in section 2.2.

We initially tune the frequency and amplitude of each rf pulse by performing simple pump-probe experiments with single rf pulses. In these experiments the centre time of each rf pulse is set to their usual running value, which is 1120 μs and 1780 μs after the initial value pulse respectively. As shown in figures 2.13(a) and 2.13(b), we scan out each parameter and choose the point where the integrated probe signal, and hence $F = 0$ population, is minimised.

Note that the probe signal oscillates in the amplitude scan; this phenomenon is known as Rabi flopping and is fully described by equation 2.16a. The amplitude minima correspond to odd integer π -pulses, however population transfer is only truly maximised when the rf frequency is accurately set on resonance, as rf frequency detunings modify the generalised Rabi frequency slightly. For this reason the rf frequency scan is always performed before

the amplitude scan. The rf scan must, however, be performed at relatively low rf power (below the π -pulse requirement) so that the transition lineshape does not split into more than one peak.

Both the pitch of the Rabi flopping scan and the transition linewidth are inversely proportional to pulse length. For our EDM measurement we would like this to be as short as possible so that the interferometer is less sensitive to rf detuning. Unfortunately the gain of the *ar worldwide KAW1050* rf amplifier limits pulse length to be not much smaller than $18\text{ }\mu\text{s}$. This value allows a little room above the π -pulse condition for amplitude tuning.

It is sometimes useful to run pump-probe experiments with a longer rf pulse length for a more precise measurement of Zeeman splittings or the tensor Stark shift. Of course the precision of these measurements is still relatively low in comparison to two pulse interference measurements. However, we can use these experiments to map the amplitude of the electric, magnetic and rf fields along the beam line by incrementing the rf pulse time in between parameter scans to hit the molecules at different positions within the machine. This technique even allows for a measurement of the rf polarisation, by comparing the relative heights of the Zeeman split $F = 1$ magnetic sub levels on various frequency scans. The details of these field mapping measurements are given in references [74, 84] and are summarised in the following sections.

When taking EDM data (with two rf pulses) we actively lock rf frequency and amplitude close to their ideal settings by independently switching the frequency and amplitude of each pulse by 1.6 kHz and 0.25 dB between shots and measuring the corresponding change in PMT signal. At the end of each block the computer automatically adjusts these rf parameters to maximise interference contrast.

Ideally we would like to apply perfect π -pulses for 100% population transfer between $F = 0$ and $F = 1$. If the rf pulse parameters are not perfect we observe Ramsey interference which contaminates the pure $\cos^2(\phi)$ interference fringes. The Ramsey component, described by equation 2.19c, is a potential source of systematic error as it depends upon the applied electric field. To cancel this error we reverse the sign of the Ramsey component between shots¹⁰ so that it averages to zero over a whole block. We do this by applying a π phase flip to the second rf pulse with a *Mini-Circuits PAS-3* bi-phase.

The bi-phase can flip the phase to better than 1% accuracy. We control it with two computer generated digital signals $\varepsilon_{\text{flip}}$ and ε_{π} . As demonstrated in figure 2.19 $\varepsilon_{\text{flip}}$ goes high in between the two rf pulses to apply the π phase shift to the second rf pulse. This

¹⁰Specifically we do this according a shot-shot switch sequence as explained in section 3.1.

happens on each and every shot but, as shown in figure 2.12, it doesn't necessarily always trigger the bi-phase. The second digital signal, ε_π , must be in the logic state high to allow the π -flip to actually occur. We switch ε_π between shots so that the π phase reversal is applied on alternate shots cancelling the Ramsey component over a block.

As a second line of defence against error originating from the Ramsey component we also randomise the phase offset between the two rf pulses on a block-by-block basis. This ensures that it averages to zero over the entire dataset. As the DCFM output is phase continuous we are able to arbitrarily change the phase difference between the two rf pulses by driving the DCFM with a voltage V_ϑ for a certain amount of time in between the two rf pulses. This is demonstrated by the blue line labelled with V_{RF1F} , V_ϑ and V_{RF2F} in figure 2.19. When running we scramble the phase offset over $(14.00 \pm 0.02)\pi$ by driving the DCFM with a voltage, V_ϑ , in the range 0–0.7248 V, for 500 μs . The voltage is randomly chosen by the computer and, like V_{RF1F} , is selected with a SPDT switch as the molecules fly through the machine. The sequence of switches that control which voltage is sent to the DCFM at any moment in time is demonstrated in figure 2.12.

We tune the range of V_ϑ by running two pulse interference experiments with the rf amplitudes set to apply $\pi/2$ -pulses in order to maximise Ramsey interference. We select a suitable duration over which phase scrambling will occur and scan V_ϑ over the range 0–1 V to obtain a Ramsey interference pattern, as shown in figure 2.13(c). We then select the voltage which corresponds to the maximum possible number of complete fringe cycles. For a phase scrambling period of 500 μs , the maximum number of cycles is seven.




FIGURE 2.14: The Ramsey phase shift across the molecular cloud, taken at 2 kV. The red points are calculated for various narrow integrated slices of the probe TOF profile. These points show good agreement with the blue points, which were generated by integrating an interpolated electric field map for each molecular velocity class.

We have found that Ramsey fringes are only visible when scans are taken with a relatively low applied electric field (<2 kV)¹¹. The reason for this is because inhomogeneities in the electric field rapidly scramble the Ramsey phase across the molecular pulse. Since the two

¹¹You could run Ramsey interference scans at high voltage but with a short interferometer length. However, you would have to reduce the phase scrambler length accordingly, which would not be appropriate for setting the range of V_ϑ . An example where it might be more appropriate to use a

rf pulses are defined in time not position, the fast molecules integrate a slightly different electric field to the slow molecules and hence accumulate a difference phase. This effect is illustrated in figure 2.14, where the red points show the phase of Ramsey fringes, taken at 2 kV, for various narrow integrated slices of the probe TOF profile. These points show good agreement with the blue points, which were generated by integrating an interpolated electric field map for each molecular velocity class. Although not intentional, this effect provides us with a third method of scrambling the Ramsey phase, further reducing its impact on our final EDM measurement.

2.3.6 Applied Electric Field

The static electric field applied to the molecules during the phase evolution period is produced by two parallel plates that are separated by 12 mm and are typically charged to ± 6 kV with respect to ground. We continue to use the gold coated aluminium plates, as detailed in the thesis of Ashworth [73], which are 750 mm long and 70 mm wide. The plates are gold coated to prevent a dielectric oxide surface layer forming. This reduces the possibility for the build up of patch potentials, which improves electric field reversibility.

Field mapping studies have shown that the plates are actually slightly bowed in the middle, deviating from the mean separation by up to $100\text{ }\mu\text{m}$ [74, 84]. Electrical contact is made by a pair of spring loaded contacts which push onto the plate edges in the middle of the vacuum chamber, as shown in figure 2.15(a). It was previously thought that pressure from these springs caused the plates to warp. Recent tests, which involved softer springs and improved plate support have since shown this not to be correct. The true cause of this warping is still as yet unknown.

short machine length would be when measuring the asymmetry of electric field reversals, which would be preferable to be taken at high voltage. In this case interferometer length is not important.




FIGURE 2.15: Schematic diagrams of the HV coupling and the HV control system. Red arrows indicate an optical line to the computer.

The high voltage (HV) supplies, that charge the electric field plates, sit 5 m away from the vacuum chamber to attenuate any generated magnetic fields that switch with the field polarity. Each plate is connected to a *Bertan 602C* supply through a *Crydom D-series* solid state DPDT relay, which is used to set the electric field polarity, as shown in figure 2.15(b). The state of the relay is set by a digital signal, ε_E , sent from the computer. This signal is carried across the lab in an optical fibre to reduce the risk of generating magnetic fields which switch with the electric field. The HV supplies are controlled and monitored by a data acquisition (DAQ) card, which is also connected to the computer via a fibre optic USB link. A 0–5 V control voltage programs the supply over a 15 kV range and the HV is read out through $1 : 10^4$ potential dividers. For additional isolation and protection from HV discharges, the individual DAQ I/O terminals are each isolated with an *Analog Devices AD215BY* analogue isolation amplifier.

Leakage currents which reverse with the direction of the applied electric field are a potential source of systematic error, as they could generate an E -state dependent magnetic field that would split the $F = 1$ manifold and hence produce an EDM masquerading signal. We monitor the leakage currents on each plate independently with two home built floating nanoammeters as described in [85]. Each ammeter outputs a train of optical pulses down a fibre optic link at a frequency proportional to the measured current. This frequency is converted to a voltage, at the computer, by a circuit which combines a monostable chip with a low pass filter. The calibration factor is 5 mV/nA. We monitor leakage currents in real time to check that the electric field is not discharging and to

ensure that it actually reverses. We expect $<5\text{ nA}$ to flow onto each plate during data acquisition which, crucially, should reverse polarity following an electric field reversal.

We pause data acquisition during the electric field reversals as charging and discharging currents become reasonably large and plate voltage is not stable. It is possible that such large current could magnetise the inner magnetic shield and generate a magnetic field which depends on the polarity of the electric field. As illustrated in figure 2.15(a), we minimise this effect by arranging the feedthroughs to pass side-by-side through a single hole in the inner shield. To keep current below $5\text{ }\mu\text{A}$ and to ensure a stable plate voltage, we switch the electric field over approximately 14s, during which time data acquisition is paused.

The switching procedure was empirically optimised by Kara, as detailed in his thesis [74]. In brief the procedure is as follows: Over 2s the supplies are ramped down to zero from high voltage; The plates are then grounded for 2s, to bleed off any charge remaining on the plates, and the relay state is given 1s to switch; The supplies are then ramped up to 1.15 times their operating voltage over 2s and are held at this voltage for 5s before being reduced to the running voltage. We allow a further 2s delay to allow transient currents to settle before acquisition is resumed. The overshoot procedure is implemented to reduce the time required to charge the plates to their operating voltage. Following a field reversal, the electric field has been shown to be stable to better than 1 V/cm . Without the overshoot, an electric field reversal would take 4s longer to obtain this stability.

2.3.7 Ambient Magnetic Fields

As shown in figures 2.6 and 2.7, the interferometer interaction region is enclosed within two cylindrical mu-metal magnetic shields: one inside the vacuum chamber with dimensions $\varnothing 0.17\text{ m} \times 0.79\text{ m}$; and one situated outside the chamber which has the dimensions $\varnothing 0.6\text{ m} \times 1.4\text{ m}$. The inner shield is capped at either end with flat mu-metal lids to maximise the shielding efficiency. The lids have small holes in their centre to allow the molecular cloud to pass through. The inner and outer shields have a radial shielding factors of approximately 75 and 50. We find that, together, they reduce the ambient radial magnetic field within the machine to $<2\text{ nT}$. Unfortunately field maps have shown that this field does increase to 200 nT at either end of the HV plates, presumably this is due to poor demagnetisation of the end caps [84]. This is another reason why we can only use the central 55cm of the HV plates for phase evolution. The longitudinal shielding is quite poor, however we are not very sensitive to perturbing magnetic fields along y .

In general magnetic field noise reduces experimental sensitivity, and components which switch with the electric field are of particular concern as they can bias the final result. We monitor magnetic fields within the lab with three *Stefan Mayer FLC100* fluxgate magnetometers: one placed near the computer, another positioned above the optics table and one mounted on the HV supply rack. We also place a more sensitive *Bartington Mag-03MCL100* magnetometer at a point near the molecular beam (see figure 2.6) to measure magnetic field along the z axis. We reduce magnetic field noise by phase locking the (25 Hz) data acquisition cycle to the mains, which ensures the molecular packet always samples the same part of the 50 Hz lab field.

2.3.8 Applied Magnetic Fields

As explained in section 2.1, we apply a magnetic field along z to operate the interferometer where it is most sensitive to phase changes. We also have the ability to apply a static magnetic field along the x axis, which is particularly useful for performing systematic checks to investigate the sensitivity of our EDM measurements to perpendicular magnetic fields. Unfortunately we can only apply fields along y with limited precision.

Both the x and z static fields are produced from a pair of rectangular current loops which are glued with vacuum compatible epoxy onto the inside of the inner shield as shown in figure 2.7. The loops are 15.5 cm wide, are separated by 7 cm and extend the full length of the inner magnetic shield (79 cm in length), looping round the outside of the shielded region through small holes in the end caps. There are no coils inside the vacuum chamber which are aligned to apply a magnetic field along y . When we would like to apply a vertical magnetic field, we usually wind coils around the vacuum chamber. This is not ideal as the applied field is then attenuated by the inner magnetic shield.

The fields along x and z are quite homogeneous — mapping studies have shown that they vary by <2% over the central section of the interferometer region [84]. Unfortunately, these fields do rotate slightly at either end of the vacuum chamber to satisfy boundary conditions with the end caps, which is a further reason why interferometer length is slightly limited.

We have two current supplies for driving the magnetic field coils: one which steps the current, and one which can continuously scan the current over 10 mA, allowing for interference scans such as figure 2.16 to be acquired. The scanning




FIGURE 2.16: A typical interference scan taken by ramping the applied magnetic field over 10 mA. The red line shows a fit to the cosine squared model.

box is particularly important for measuring the magnetic field current coefficient^{12,13} $\beta = (16.790 \pm 0.003) \text{ nT/mA}$. However, it is also useful for general diagnostics which include checking parameters are well tuned and measuring ambient magnetic fields. The scanning box is controlled by a 0–5 V drive which is output from the main computer.

Under normal running conditions, the magnetic field coils are connected solely to the stepping box¹⁴. This coil driver produces a magnetic field within the machine of

$$B = \beta \varepsilon_{B_0} |I_{B_0} - \varepsilon_{\delta B} I_{\delta B}| + \beta I_{\text{bias}}, \quad (2.20)$$

where the bias current I_{bias} (usually $< 0.15 \text{ mA}$) is controlled by an analogue voltage, V_{bias} , in the range $-5 - +5 \text{ V}$. This ensures the current steps symmetrically about the central fringe — it cancels the average residual ambient field inside the inner chamber. $I_{B_0} = 0.812 \text{ mA}$ and $I_{\delta B} = 0.099 \text{ mA}$ are fixed currents that are set manually with two potentiometers which divide a stable voltage reference. $\varepsilon_B, \varepsilon_{\delta B}$ represent the state of two digital signals (either -1 or +1).

As with the electric field supply, both the digital signals are optical and the bias voltage is supplied by a DAQ card which is connected to the computer via a USB optical link. We use optical fibres for electrical isolation and to reduce magnetic fields around the lab.

Figure 2.17 illustrates the essence of the current stepping box circuitry. A low voltage noise instrumentation amplifier measures the voltage drop across a high stability metal

¹²This number is actually calculated from the average of two interference scans in which the phase difference between the two rf pulses differs by π radians. This is to cancel out the Ramsey component of the interference signal, which has the ability to change the pitch of the interference fringes.

¹³Note this is the current coefficient which corresponds to the average field experienced by the molecules during the phase evolution period. Due to minor field inhomogeneities, this value may vary slightly with pulse timings, as the molecules integrate a slightly different magnetic field.

¹⁴There is no switch here. We physically unplug the coils from the scanning box and plug them into the stepping box.




FIGURE 2.17: Schematic of the magnetic field current supply.

film resistor. This voltage is then fed back to the inverting input of a precision op-amp to ensure a constant current output that is proportional to the voltage on the non-inverting input, and is independent of the load placed upon it.

By placing a magnetometer at the centre of the inner shield we have observed that the magnetic field can switch in $< 350 \mu\text{s}$ (the current actually switches in $< 100 \mu\text{s}$). The current running through the magnetic field coils is monitored with a *HP 34401A* six digit multimeter. When running we log the current output in each of the four magnetic field switch states once per block to ensure the supply is working correctly and to make sure the phase calibration step ($\delta\phi_{\text{CAL}}$) is calculated correctly.

2.3.9 Computer

Hardware

The central hub of the EDM experiment is a *National Instruments PXI* computer system, that drives the experimental apparatus and logs detector signals. A schematic diagram of the computer–hardware interface is shown in figure 2.18.

Recall that an individual block is made up from 4096 shots in which a series of experimental parameters¹⁵ are switched in between each shot (the so called ‘shot–shot switch sequence’). Each shot, however, is made up of a complicated sequence of events that occur as the molecules fly through the beam machine. This relatively rapid sequence of events that occurs as the molecules travel through the machine is demonstrated in figure 2.19.

¹⁵Nine parameters to be precise. See section 3.1.




FIGURE 2.18: Schematic diagram of the computer–hardware interface.

The *NI PXI 6535* pattern generator (PG) provides both the (fast) sub-shot and (slow) shot-shot digital signals. It not only drives all the hardware switches but also triggers the analogue input cards to begin acquisition. Note that this card generates the inverted triggers ($!\varepsilon_{E_0}, !\varepsilon_{B_0}, !\varepsilon_{\delta B}, !\varepsilon_{\pi}$) in addition to $(\varepsilon_{E_0}, \varepsilon_{B_0}, \varepsilon_{\delta B}, \varepsilon_{\pi})$, as were defined in the previous sections. This is to balance the currents which drive the LEDs used for optical isolation, in order to reduce any magnetic fields associated with these switches. The inverted switches drive dummy LEDs which do not trigger anything.




FIGURE 2.19: The sequence of events that occur as the molecules fly through the beam machine. All times are taken relative to the YAG laser Q-switch trigger. The attenuator and DCFM voltages are set in between each shot.

We provide the PG with a clock reference from a 1 MHz synthesizer to time events with $1 \mu\text{s}$ resolution. As was mentioned previously, this clock is phase locked to the lab mains to reduce shot-shot magnetic field noise. We first digitise the mains with a Schmitt trigger. The digitised mains and the 1 MHz signal is then fed into a *NI PXI 6602* rf counter. The synthesizer is phase locked by actively servoing its frequency to ensure the counter measures the same number of cycles within a period, which is gated by the lab mains half cycle. We also feed the digitised mains into the PG to provide a start trigger for each block. This ensures each EDM measurement is taken on the same part of the mains cycle. Typically the 1 MHz synth frequency is varied by a few parts in 10^4 .

Many analogue voltages are modulated in addition to the digital modulations produced by the PG. The DCFM and laser control voltages are examples of such. These analogue voltages are produced by a selection of USB data acquisition (DAQ) cards. We also use

the GPIB and RS232 protocols to interface with external hardware, such as the 170 MHz rf synthesizer¹⁶. These protocols are bidirectional so are used to both program hardware and transfer measurements from the stand alone meters. We update the USB, GPIB and RS232 devices between shots as their update rate is relatively slow.

We use two separate DAQ cards for collecting data: A *NI PXI 6220* card is used for single point voltage measurements that are taken once per shot, whereas a *NI PXI 6133* is used to continuously sample certain detectors as the molecules fly through the beam machine. These are the detectors which we analyse to look for an EDM signal. The detectors connected to the *NI PXI 6133* are the two PMTs, the *Bartington* magnetometer and two dummy inputs: a 9 V battery and a $50\ \Omega$ terminator. We use the dummy inputs as tests for systematic effects, for example if the 9 V battery measures a non-zero EDM there clearly is something wrong with our experiment.

The *NI PXI 6133* DAQ card was carefully chosen to provide isolated inputs, each with a dedicated analogue to digital converter, that can be read simultaneously at a high sample rate. This limits the level of cross talk between channels and prevents interference from ground. When triggered by the PG, the *NI PXI 6133* begins sampling each input at 100 kHz, however we bin data from the $50\ \Omega$ terminator and magnetometer into blocks of $200\ \mu\text{s}$ which are averaged in order to reduce the memory overhead per shot. The sample time periods of each detector are shown in table 3.1.

Software

The EDM experiment is almost completely automated. It is fully capable of operating without human intervention for over 12 hours at a time. The experiment is operated by a small collection of programs written in C# on the .NET framework. The following items provide a brief overview of the main function of each program in turn.

HardwareController As the name suggests, this program is used to give the user direct control of experimental apparatus. All changes to the hardware go through this program.

ScanMaster The fundamental idea of this program is to provide the ability to scan any one experimental parameter over a number of shots, while holding all other parameters constant. This software is also used to program the pattern generator.

¹⁶We find USB to be the least reliable protocol. It fails at least once a week, requiring acquisition to be halted and the computer restarted. This doesn't cause a major problem, but can lead to the loss of a whole night's worth of data ~ 150 blocks.

BlockHead This program runs the experiment during the acquisition of EDM data. It runs the pattern generator and modulates the various experimental parameters in between shots. It also logs the detectors as specified above. After a full block shot sequence is completed, BlockHead saves all the raw data to disk in compressed XML format. The XML format adds tags to the data to make it human readable. This is really important should anyone like to re-analyse the raw data in the future. BlockHead also runs a simple analysis for feedback in between blocks and to provide real-time diagnostics — we can see how far detuned parameters are and monitor the noise on important switch channels. This real-time analysis report is particularly useful for telling us if the dye laser has unlocked.

ScriptingConsole This runs an instance of IronPython, which is a version of the Python computer language that is tightly integrated with Microsoft’s .NET Framework. We use this program to run scripts to automate the experiment. For example the script `MapLoop` is used to control ScanMaster. This script changes the value of one of the fixed parameters in between individual scans.

PhaseLock Runs constantly in the background, servoing the 1 MHz PG reference clock as described above.

2.4 EDM Data Acquisition Cycle

The EDM data acquisition cycle is almost completely computer automated by the `EDMLoop` script. At the beginning of a block, this script instructs `HardwareController` to randomise the laser polarisations and the rf phase scrambler voltage, and measure the various experimental parameters as listed in table D.2. It then generates a set of switch sequences (waveforms) for the shot–shot modulation of the experimental parameters and logs these along with the parameter measurements and the current pattern generator settings. The script then requests BlockHead begin the pattern and start the shot–shot parameter modulation sequence. At the end of a full parameter switch sequence (4096 successful shots), BlockHead runs its quick analysis and saves all the raw block data to disk. Before beginning a new block, `EDMLoop` requests the analysis values from BlockHead. It then uses these values to servo various experimental parameters closer to their ideal settings¹⁷ (see section 3.3). If the interference signal drops below a certain value the script rotates the target to expose fresh Yb. Once the new centre values are set, the script then loops

¹⁷The computer is much better at optimising the parameter space than we are using individual parameter scans.

round and starts a new block; this procedure continues indefinitely until it is interrupted. We typically leave the script running for ~ 12 hours to collect ~ 150 data blocks. We then manually reverse one or more of the electrical connections between the beam machine and the electric field, magnetic field and rf supplies. This reverses the direction of each field providing a valuable handle upon the origin of systematic effects. We try to keep the order in which ‘manual reversals’ are performed as random as possible, and we ensure an equal number of blocks are taken in each of the eight possible configurations. Data acquisition runs 24 hours a day, five days a week and is only interrupted for a manual reversal, if the dye laser needs maintaining, or if the Yb target wheel needs refreshing. When everything is running well we obtain approximately 200 data blocks per day.

Chapter 3

Analysing Detector Signals

This chapter covers the details of how we actually make an EDM measurement. To understand how we extract a value for d_e from the raw detector signals, I first discuss the shot–shot modulation of various experimental parameters over the course of a block. Since the signals from many detectors vary as the molecules pass through the beam machine we must be especially careful as to how we combine the recorded signal changes and integrate them to a single value. This is discussed in section 3.2. Section 3.3 then lists the most important signal changes, or ‘analysis channels’, that we derive from the shot–shot parameter modulations. In section 3.4 I explain how we use many of these channels to lock experimental parameters to their ideal values during data acquisition. In section 3.5 I then discuss how we manually reverse the electrical connections between the beam machine and the E , B and rf supplies. We perform these reversals every ~ 150 blocks to verify the origin of the important analysis channels. Section 3.6 brings sections 3.1–3.5 together to demonstrate in detail how we combine the analysis channels to extract an EDM value for each block. These values are hidden from us until the analysis is completed. In section 3.7 I explain that we do this so we do not unintentionally bias the measured results. Finally sections 3.8–3.9 explain how we average all the individual block values together and calculate an average value for the electron EDM.

3.1 Switch Channels and Waveforms

In the previous chapter I often stated that, during the data acquisition sequence, various experimental parameters are modulated (or switched) on a shot–shot basis according to predefined switch sequences. First and foremost we do this as a means of phase-sensitive detection — it allows us to measure the small signal changes which corresponds to each particular parameter modulation in the presence of many sources of noise. However, the

benefits of this technique are more far reaching than this alone: At the end of a block it provides us with error signals which we use to servo parameters to their optimum values; It lets us understand how the machine behaves when parameters are not set perfectly, and enables an EDM measurement to be calculated even when they are not; Finally it provides us with a simple method to search for systematic errors by allowing correlation analysis to be performed between any number of switch channels.

Our parameter switching technique is very similar to that published in [86]. This technique allows us to analyse detector signals that depend upon several experimental parameters. In essence our technique is just a generalisation of phase sensitive detection to the concurrent modulation of many experimental parameters. During the acquisition of a data block each experimental parameter is modulated according to a unique ‘waveform’. A waveform which is just a list of +1s and -1s which define the state of a parameter for each shot in the block; it precisely defines the frequencies according to which the parameter is switched (or modulated). Crucially the waveforms all contain equal numbers of +1s and -1s (they are balanced), and are orthogonal. This means we may extract a detector signal change that is unique to each switch waveform.

For K switched experimental parameters there are 2^K unique experiment states. We therefore require waveforms of length $m 2^K$, where m is an integer ≥ 1 . We generate these waveforms through element-wise multiplication of simple basis waveforms. The basis waveforms are taken from a set of waveforms which describe regular square waves of period 2^k , where k is an integer in the range $1 \leq k \leq K$. By combining these waveforms we generate complex waveforms that contain frequency components at the fundamental and harmonic¹ frequencies of all the constituent basis waveforms. This does not generate all possible waveforms that are balanced and orthogonal [87], but the number of waveforms that this procedure can generate far exceeds the requirements that are described below. This way of generating waveforms also enables them to be precisely specified with a simple, K -bit ‘wavecode’, which makes recording and comparing waveforms much more straightforward. For the basis waveforms, the wavecodes are simply their period (2^k) written in binary notation. The wavecodes of the more complex waveforms are the bitwise OR of the basis component wavecodes.

During data analysis we use the waveforms to extract the signal correlated with each experimental parameter. We denote a ‘channel’, $\{X\}$, as the signal change that corresponds to the change in parameter or parameter combination (see below), X . This parameter/parameter combination is associated with the waveform W_X . We evaluate the

¹The basis waveforms are square not sinusoidal, so each contains many harmonics at integer multiples of the fundamental frequency.

value of $\{X\}$ for each block by applying the formula:

$$\{X\}_d = \frac{1}{m 2^K} \sum_{i=1}^{m 2^K} W_X(i) S_d(i), \quad (3.1)$$

where $S_d(i)$ is the signal recorded by detector d on the i th shot of the block, and $W_X(i)$ is the i th element of the waveform W_X . In essence equation 3.1 calculates the difference in signal between two switch states of X , and averages over the states of all other experimental parameters. A list of detectors to which this equation applies is provided in the first column of table 3.1. For simplicity we often drop the detector subscript. If absent, it should be assumed that the channel is as applied to the normalised probe PMT signal.

TABLE 3.1: The various detectors which may be analysed using equation 3.1.

Data Source	Start Time (μ s)	Resolution (μ s)	Sample Points
probe PMT	1800	10	80
pump PMT	400	10	40
Bartington magnetometer	400	200	11
9 V battery	1800	10	80
50 Ω terminator	400	200	11
fluxgate magnetometers	~ 2600	—	1
leakage monitors (x2)	~ 2600	—	1
laser photodiodes (x2)	~ 2600	—	1
phase lock frequency	~ 2600	—	1
phase lock error	~ 2600	—	1

Modulating multiple switch channels concurrently, as we do, gives a significant benefit as it allows correlation analysis to be performed. During data acquisition each waveform specifies the shot-shot modulation of only one experimental parameter. However, the waveform W_X in equation 3.1 need not correspond to the modulation of only one parameter. It can be formed from any combination of the waveforms that were used to modulate the experimental parameters over the course of a block. We generate the combined waveforms via element-wise multiplication of the component waveforms. Consider, for example, two arbitrary experimental parameters: Y and Z , that are respectively modulated with the waveforms W_Y and W_Z during data acquisition. We can substitute $W_X(i)$, in equation 3.1, with $(W_Y \cdot W_Z)(i)$ to extract the channel $\{Y \cdot Z\}$. This channel compares the signal when Y and Z are in the same switch state to that when they are in opposing switch states. It is important to note, however, that the channel $\{Y\}\{Z\}$ is completely different to $\{Y \cdot Z\}$. We typically refer to compound channels as “X dot Y”.

Analysing the data with correlated waveforms is how we extract an EDM measurement. Let us consider this important example to make the above abstract mathematics more




FIGURE 3.1: The 8 bit switch sequences that make up the simplest EDM measurement. These switch states are illustrated by the points (A,B,C,D,A',B',C',D') in figure 2.3.

transparent. As stated in the last chapter, currently a block constitutes 4096 individual shots which affords the possibility of $K = 12$ parameter modulations ($2^{12} = 4096$). The most crude measurement of the EDM, however, requires only three modulations: the electric field reversal (W_E), the magnetic field reversal (W_B), and the calibration magnetic field step ($W_{\delta B}$). This requires a block length of (at least) $2^3 = 8$ shots. The basis waveforms of a block of this length are W_{100} , W_{010} , W_{001} , which are illustrated in figure 3.1. We could assign each of W_E , W_B , $W_{\delta B}$ to any combination of these waveforms (so long as the combinations differ), such as $W_{011} = W_{010} \cdot W_{001}$, and for the noise response arguments described below it would actually be beneficial to do this, but to keep this demonstration simple let us assign $W_E = W_{100}$, $W_B = W_{010}$ and $W_{\delta B} = W_{001}$.

By itself, the channel $\{B\}$, calculated from the waveform W_B , measures the average interferometer phase, as demonstrated in figure 3.2. We aim to measure the change in phase associated with reversing the electric field, $\delta\phi_E$. So the EDM signal is extracted with the waveform $W_{\text{EDM}} = W_E \cdot W_B$ which, in the case described above, is $W_{100} \cdot W_{010} = W_{110}$. As shown in figure 3.1, this waveform subtracts the signals recorded when E and B are opposed from the signals recorded when they are aligned.

The parameter switching technique allows us to extract a small signal change in the presence of many sources of noise. It has previously been observed [73] that the EDM experiment suffers from two clear noise components: a $1/f$ component, mainly from stray magnetic fields; and a white noise component which dominates above ≈ 4 Hz. There are also small noise peaks which appear at 5.5 Hz, 10.75 Hz and 11.75 Hz [73]. In reference [86] Harrison *et al* decompose noise into two forms: ‘statistically stationary noise’ — noise whose features are time independent; and drift — signal changes that have no clear frequency, which are slow in comparison to the block time period (roughly six minutes). In the first case our switching technique is most effective when experimental parameters are switched as frequently as possible, to avoid $1/f$ dominated noise. It is also advantageous to avoid frequencies which coincide with peaks in the noise spectrum.




FIGURE 3.2: The channel $\{B\}$ that is associated with the modulation of the applied magnetic field direction measures the average interferometer phase. In (a) there is no ambient magnetic field so the interference fringes are centred and $\{B\} = 0$. In (b) an ambient magnetic field shifts the fringe pattern and $\{B\} \neq 0$.

Drifts are eliminated by adding many component waveforms into the analysis waveform. In fact with every extra frequency component added, we eliminate one extra derivative in the signal drift Taylor expansion [86]. In particular, we require an analysis waveform composed from at least two frequencies to eliminate a linear change in signal.

TABLE 3.2: Applied experimental parameter modulations.

Experimental Parameter	Modulation	Step Size
electric field direction	W_E	—
magnetic field direction	W_B	—
magnetic field calibration step	$W_{\delta B}$	± 1.7 nT
laser frequency	W_{LF1}	± 340 kHz
rf pulse frequency	W_{RFiF}	± 1.6 kHz
rf pulse amplitude	W_{RFiA}	± 0.25 dB
rf phase offset	W_π	$\pm \pi/2$ rads

At the beginning of every block BlockHead generates a set of waveforms for the nine experimental parameters² listed in table 3.2, according to the conditions specified in the paragraph above. With the exception of W_E and W_{LF1} , the program generates these waveforms randomly to prevent specific noise frequencies from leaking into our experiment. It also checks and ensures that each of the $2^9 = 512$ possible analysis waveforms are unique. Furthermore, the waveform generator inverts a random selection of the nine parameter waveforms (including W_E and W_{LF1}) to average out systematics due to repeatable backgrounds that are the same each block. Unfortunately we cannot reverse the electric field rapidly as this would generate large currents which would magnetise the inner shield. We must impose 14s of dead time during an E reversal in order to give charging/discharging currents time to settle. It is therefore impractical to frequently reverse the electric field, so we constrain W_E to be relatively slow with the wavecode

²There are two rf pulses!

111100000000, which reverses the field roughly every 20 s, with 10 switches per block. The magnetic field can, however, switch rapidly, with a settling time that is much faster than the 40 ms shot–shot interval. By switching the magnetic field rapidly, we ensure that the EDM waveform, $W_{E.B}$, satisfies the required conditions. The dye laser also cannot be modulated rapidly else the pump and probe beams become frequency unstable. We therefore assign W_{LF1} the wavecode 00011111000.

Note that with only nine switch parameters, the machine visits each unique parameter state $m = 2^{12}/2^9 = 4096/512 = 8$ times within each block. This replication allows more parameter modulations to be incorporated into the block as the experiment developments over time. It also allows a rough estimate of error/uncertainty to be calculated for each channel at the end of every block. This is particularly useful for checking the experiment is running properly with no anomalously noisy channels.

When setting the amplitudes of the parameter modulations listed in table 3.2, we empirically chose values that would not perturb the experiment far from ideal conditions, yet would be large enough to be seen over experimental noise within a reasonable time frame. We also chose to keep parameter modulations small so that we may use a second order Taylor expansion (see appendix A) to easily convert the analysis channel values into more useful physical units. Our analysis assumes that the probe PMT signal is linear in phase, so the calibration magnetic field step, $\pm\phi_{\delta B}$, must not be too large. We set $\pm\phi_{\delta B} = \pm\pi/32$ which is 1/8 th of the larger magnetic field step. We set the laser and rf frequency modulations $\delta\nu_{LF1}$ and $\delta\nu_{rf}^{(i)}$ to be 2% and 3% of the respective optical and rf transition linewidths. We also set the rf amplitude step, $\delta a_{rf}^{(i)}$, to be 5% of the π -pulse power requirement.

3.2 Non-linear Channels

As explained in section 2.3.9, we record the temporal variation of many detectors over the course of each shot, as the molecules fly the length of the beam machine. We therefore need to modify equation 3.1 to account for a possible temporal variation in the analysis channels:

$$\{X\}_d(t) = \frac{1}{m 2^K} \sum_{i=1}^{m 2^K} W_X(i) S_d(i, t). \quad (3.2)$$

Now $S_d(i, t)$ is the signal from detector d , at time t after ablation, on the i th shot of the block. Table 3.1 lists the start time, the resolution, and the number of samples of each detector. We most often apply equation 3.2 to the normalised probe PMT signal which is calculated as $S_{PMT}(i, t) = \mathcal{S}(i, t)/\mathcal{N}(i)$, where $\mathcal{S}(i, t)$ is time and shot dependent probe PMT signal, and $\mathcal{N}(i)$ is the integrated pump PMT signal for the i th shot. We normalise




FIGURE 3.3: An illustrative example of equation 3.4. The shaded regions indicate a typical integration gate. Integrating channels (a) $\{X\}(t)$ and (b) $\{Y\}(t)$ separately returns the product $\langle\{X\}\rangle\langle\{Y\}\rangle = 0$ which is incorrect. The two channels must be multiplied before integrating, as illustrated in (c). The integration of the TOF profile for $\{X\}(t)\{Y\}(t)$ returns a non-zero value in this case.

the probe PMT signal to correct for shot–shot variations of the molecular flux. Note that we do not retain time dependence in the pump PMT signal as this solely measures the number of molecules produced per shot.

We check the temporal variation of all the important analysis channels for each detector. Many channels do exhibit a significant variation when calculated using the normalised probe PMT signal. This is because no field is perfectly homogeneous and the faster molecules inevitably integrate slightly different electric, magnetic and rf fields to the slower molecules. For many channels such as the rf frequency modulation channel, $\{RFiF\}$, this is not a major concern so long as the variation is not too large. But for other channels, such as the electric field switch channel, $\{E\}$, any variation over the TOF profile would indicate a significant problem. If nothing untoward is observed in any detector we calculate a single average value for each channel by integrating over a some chosen time interval, or ‘gate’:

$$\langle\{X\}\rangle = \int_{t_i}^{t_f} \{X\}(t) dt. \quad (3.3)$$

We are often interested in calculating the product or ratio of two or more analysis channels of the normalised probe PMT. As stated, many of these channels do vary significantly across the TOF profile, often crossing zero at the centre of the profile, as illustrated in figures 3.3(a) and 3.3(b). If we blindly multiply and divide the integrated channels calculated using equation 3.3 we are likely to make a mistake, this is because

$$\int \frac{\{X\}(t)\{Y\}(t)dt}{\{Z\}(t)} \neq \frac{\int\{X\}(t)dt \int\{Y\}(t)dt}{\int\{Z\}(t)dt}, \quad (3.4)$$

as demonstrated in figure 3.3. For this reason, when calculating the product or ratios of channels, we always perform the integration last, i.e. we calculate $\langle\{X\}\{Y\}/\{Z\}\rangle$ as opposed to $\langle\{X\}\rangle\langle\{Y\}\rangle/\langle\{Z\}\rangle$.

In the following analysis I will almost always refer to integrated channels so for simplicity from here onwards I will drop the angle bracket notation. It should be assumed that the various channel combinations are integrated over a time window, as detailed above. Details regarding the range of the interval $t_i \rightarrow t_f$ will be given when it is important.

3.3 Important Channels

There are a number of analysis channels which provide an insight into the state of the experiment during each block. Not only do we use their values during the analysis process, but we also monitor a number of them in real time to diagnose the experiment and optimise hardware settings in between each block. The physical meanings of the most important analysis channels are summarised in table 3.3. Unless otherwise stated, it should be assumed that the descriptions refer to the application of each analysis waveform to the normalised probe PMT signal. This is an important point as each detector responds differently to changes in each experiment parameter, so the meaning of an analysis channel changes accordingly. As an example, the probe PMT signal varies sinusoidally with magnetic field, whereas the magnetometer signal varies linearly. An E -switch correlated magnetic field would therefore be associated with $W_{E.B}$ in the probe PMT and W_E in the magnetometer.

TABLE 3.3: The physical meaning of the most important analysis channels.

Channel	Description
$\{SIG\}$	The average detector signal. This channel is associated with the waveform that contains no switches.
$\{\delta B\}$	A measure of interference contrast. Specifically, it measures the gradient $\partial S / \partial \phi$. Recall that the δB switch ($\varepsilon_{\delta B}$) modulates the absolute size of the applied magnetic field. This channel is, therefore, always positive and should be linearly related to $\{SIG\}$. We use this channel almost universally to calibrate all other analysis channels, as it is proportional to the molecular signal and corresponds to a well known step in interferometer phase.
$\{X \cdot \delta B\}$	A measure of the change in interference contrast that correlates to a change in parameter X .
Continued on the next page...	

Table 3.3 – continued from previous page

Channel	Description
$\{E\}$	The change in signal when the electric field is reversed. Ideally no detector should report a statistically significant value for this channel. If either of the PMTs or the <i>Bartington</i> magnetometer do return a non-zero value then something is significantly wrong with the electric field. It is inevitable that the <i>Stefan Mayer</i> fluxgate magnetometer positioned close to HV relays will always detect the electric field reversals. This is not a problem so long as none of the other magnetometers, closer to the beam machine, also detect the E -switch correlated magnetic field.
$\{B\}$	A measure of the average interferometer phase, as demonstrated in figure 3.2. On its own this channel measures the z component of uncanceled stray magnetic fields. This channel is zero when the interferometer steps equally either side of the central interference fringe.
$\{X \cdot B\}$	The correlation between interferometer phase and a change in parameter X . The most important example of this is $\{E \cdot B\}$, which measures the change in phase correlated to an electric field reversal, which is an EDM!
$\{RFiF\}$ $\{RFiA\}$	These channels correspond to the rf frequency and amplitude modulations. A Taylor expansion of equations 2.19a and 2.19b reveals that the probe PMT signal is quadratic for small deviations of both rf parameters about their ideal values (see appendix A). These channels, therefore, respond linearly to small detunings, and are zero when each pulse is perfectly tuned to the π -pulse condition. Note the letter ‘F’ stands for frequency and ‘A’ for amplitude. As always, the index i labels each rf pulse.
$\{RFiF \cdot E\}$	The change in rf detuning when the electric field is reversed. If non-zero, this channel indicates that the magnitude of the electric field changes upon reversal.
$\{\pi\}$	A measure of the size of the Ramsey component. This should be zero.
Continued on the next page...	

Table 3.3 – continued from previous page

Channel	Description
$\{LF1 \cdot \delta B\}$	The change in interference contrast when the laser frequency is stepped. This channel measures the detuning of the probe laser from the $Q(0)(F = 0)$ transition. The channel is zero when the probe laser is on resonance. $\{LF1 \cdot \delta B\}$ provides a better measure of laser detuning than $\{LF1\}$, as $\{LF1\}$ is zero when the probe laser is tuned to a peak in the optical spectrum. The frequency of this peak does not necessarily correspond to $Q(0)(F = 0)$ as there may be underlying spectral lines. However, interference contrast is maximised when the probe laser is tuned to $F = 0$.

3.4 Lock Channels

In between each block we servo hardware settings in order to optimise experimental sensitivity. We do this to account for slow drifts and imperfect initial settings. We wish to maximise interference contrast as EDM uncertainty is proportional to $1/\{\delta B\}$. In general we calculate the required adjustment of parameter X by using the channel ratio $\{X \cdot \delta B\}/\{\delta B\}$ since this measures the change in interference fringe contrast when X is stepped. $\{X \cdot \delta B\}$ is therefore zero when interference contrast is maximised (the denominator, $\{\delta B\}$, is essentially a constant that never changes sign). The `EDMLoop` script implements an integral type lock, adjusting experimental hardware by an amount

$$\Delta X = \begin{cases} +\Delta X_{\max} & \text{if } G \frac{\{X \cdot \delta B\}}{\{\delta B\}} > \Delta X_{\max}, \\ +G \frac{\{X \cdot \delta B\}}{\{\delta B\}} & \text{if } \Delta X_{\max} < G \frac{\{X \cdot \delta B\}}{\{\delta B\}} < \Delta X_{\max}, \\ -\Delta X_{\max} & \text{if } G \frac{\{X \cdot \delta B\}}{\{\delta B\}} < -\Delta X_{\max}, \end{cases} \quad (3.5)$$

where $\Delta X_{\max}$ is the maximum permissible change in parameter X between blocks, and G is an empirically chosen gain value. Table 3.4 lists the lock channel ratios used during normal running. The values listed for G and $\Delta X_{\max}$ have little meaning outside the lab. They are quoted only for future reference. Figure 3.4 demonstrates an example of one particular parameter locking in over approximately 10 blocks.

We monitor the values of $\{\delta B\}/\{SIG\}$ and $\{\delta B\}$ whilst the experiment is running. Typically $\{\delta B\}/\{SIG\}$ should lie in the range 0.11-0.13. A low $\{\delta B\}/\{SIG\}$ value in

the range $\sim 0.08 - 0.09$ is indicative of poor laser beam alignment. If $\{\delta B\}/\{SIG\}$ suddenly drops below 0.07, this most often implies that the dye laser has unlocked.

TABLE 3.4: The channels used to optimise experimental hardware settings.

Parameter	Lock Channel	Gain (V)	Max. Change (V)	Comments
Laser frequency	$\frac{\{LF1 \cdot \delta B\}}{\{\delta B\}}$	-1.25	0.1	Corrects any Doppler shifts which occur when tweaking laser beam alignment.
rf frequency	$\frac{\{RFiF \cdot \delta B\}}{\{\delta B\}}$	-2	0.1	Corrects for drifts in synthesizer output and electric field amplitude.
rf amplitude	$\frac{\{RFiF \cdot \delta B\}}{\{\delta B\}}$	-0.73	0.1	Corrects for drifts in synthesizer output and amplifier gain.
magnetic field current bias I_{bias}	$\widehat{M}_B \frac{\{B\}}{\{\delta B\}}$	-0.0125	0.05	Corrects for the slow day-to-day changes in the ambient magnetic field. It ensures PMT signal remains linear in phase.




FIGURE 3.4: Demonstration of parameter locking. The cluster began with an offset in the frequency of the second rf pulse. Error bars indicate variation within each block.

This data was taken from the cluster labelled “02Dec1001”.

3.5 Manual Reversals

Approximately once a day we manually reverse at least one of the connections between: the electric field plates and the HV relays; the magnetic field coils and the current supply; the rf transmission line and the rf amplifier. These reversals provide us with a valuable tool to test the origin of non-zero analysis channels.

The electric and magnetic field manual reversals change assignment of field direction to waveform switch state. These reversals allow us to determine whether the channels derived from W_E and W_B result from the influence of fields produced inside the machine, or whether they are the result of some effect associated with the switching hardware. A real EDM, for example, will generate a signal in $\{E \cdot B\}$ that will change sign when the connections to either (but not both) the electric field plates or magnetic field coil are manually reversed. A non-zero $\{E \cdot B\}$ would also result if the HV supplies generated a magnetic field that depended upon the state of the HV relays, for example. However, this EDM masquerading signal would not change sign following an electric field manual reversal, as reversing the connections to the HV plates would have no effect on the E -switch correlated magnetic field. We label the ‘manual state’ of the electric and magnetic field connections with the boolean values: M_E and M_B . When both are true the electric and magnetic fields are aligned along $+\hat{z}$, when each are in their waveform switch state -1 . In manual state false, these orientations are reversed. Most often we would like to cancel the systematic effects due to the field switching hardware³ and measure the influence of the electric and magnetic fields generated inside the beam machine. Therefore, when averaging channels that involve W_E or W_B over a set of blocks, we usually sign each channel value according to the state of M_E and M_B (+1 for ‘true’ and -1 for ‘false’) and average over all manual states.

The rf manual reversal changes the direction of the rf propagation. We label the manual state of the rf connections with M_{RF} . When M_{RF} is ‘true’ the rf field propagates up the transmission line, and when M_{RF} is ‘false’ the rf propagates down the transmission line. Note that changing the rf manual state does not affect any parameter modulations. So, we only sign the channel values for a collection of blocks by M_{RF} when we would like to measure the average change in a channel value between the two rf manual states.

³Note that manual reversals *do not* eliminate the systematic effects generated inside the beam machine

3.6 EDM Calculation and the Lineshape Correction

Ultimately we must express a value for the EDM in physical units. As was discussed in section 2.1.1, we take measurements on the interference curve at $\phi = \pm\pi/4$, where the interferometer is most sensitive to phase changes. Here a small step in phase, $\delta\phi$, will cause a linear change in the probe PMT signal, $\delta S|_{\pm\pi/4} = (\partial S/\partial\phi)|_{\pm\pi/4} \delta\phi$. We calibrate the EDM signal change, δS_{EDM} , by measuring the signal change, δS_{CAL} , that results from making a small, known step in phase, $\delta\phi_{CAL}$, with the magnetic field, $B_{\delta B}$. So long as the ambient background field is perfectly cancelled within in the machine, and providing that the interference contrast does not depend upon the switch state of the applied electric field, then:

$$\frac{\delta\phi_{EDM}}{\delta\phi_{CAL}} = \frac{\delta S_{EDM}}{\delta S_{CAL}} = \frac{\{E \cdot B\}}{\{\delta B\}}. \quad (3.6)$$

However, if these conditions are not met then we require two other analysis channels, namely $\{B\}$ and $\{E \cdot \delta B\}$, to correctly extract the EDM phase shift, $\delta\phi_{EDM}$.

As demonstrated in [76], to derive an expression for the EDM phase shift that is robust against experiment imperfections, we approximate the interference curve with a linear model. We define a state dependent probe PMT signal, S , which is the product of an amplitude term, A , and a phase term $f(\phi)$:

$$S(\widehat{E}, \widehat{B}, \widehat{\delta B}, \widehat{X}) = Af(\phi). \quad (3.7)$$

The arguments $(\widehat{E}, \widehat{B}, \widehat{\delta B}, \widehat{X}) \in \{+1, -1\}$ represent the switch state of the applied electric and magnetic fields, as well as the state of an arbitrary experimental parameter, X . X may be any one of the other experimental parameters that are modulated over the course of a block, for example the frequency of the first rf pulse.

We assume that the PMT signal is linearly dependent on a number of phases, as shown in figure 3.5:

$$f(\phi) = \alpha \left(\phi_{bg} \widehat{B} + \phi_{EDM} \widehat{E} \widehat{B} + \phi_X \widehat{X} \widehat{B} - \phi_{\delta B} \widehat{\delta B} \right) + \beta. \quad (3.8)$$

$\phi_{\delta B} = g \mu_B B_{\delta B} T/\hbar = \pi/32$ and $\phi_{bg} = g \mu_B B_{bg} T/\hbar$ are the magnetic phases due to the calibration field step $B_{\delta B}$, and an uncanceled background field B_{bg} . $\phi_{EDM} = -d_e E_{\text{eff}} T/\hbar$ is the EDM phase, and ϕ_X is the phase induced⁴ by parameter X . Note that each parameter modulation steps the phase by $\pm\phi$, so that the total change $\delta\phi = 2\phi$. Note also that change in signal depends upon which side of the central interference fringe is

⁴The origin of ϕ_X need not be well understood.




FIGURE 3.5: A linear model for the phase dependent part of probe PMT signal, $f(\phi)$, as expressed in equation 3.8. $\hat{B}$ and $\hat{E}$ denote the states of the E and B fields, $\hat{\delta B}$ denotes the state of the calibration step. ϕ_{bg} is the phase due to an uncanceled background magnetic field, $\phi_{\delta B}$ is the phase due to the calibration field step and ϕ_{EDM} is the EDM phase. β is average probe PMT signal. The dependency of $f(\phi)$ on arbitrary switch parameter X is exactly analogous to $\hat{E}$.

being sampled. $\phi_{\delta B}$ is the exception to this rule, as $W_{\delta B}$ modulates the absolute size of the applied magnetic field. $\alpha = (\partial S/d\phi)|_{+\pi/4} = -(\partial S/d\phi)|_{-\pi/4} < 0$, and β represents the average signal.

We allow the amplitude of the interference curve, A , to change by an amount $2\Lambda_Q$ for each switch parameter combination Q , with the exception of the (small) δB step.

$$A = 1 + \Lambda_E \hat{E} + \Lambda_B \hat{B} + \Lambda_X \hat{X} + \Lambda_{E \cdot B} \hat{E} \hat{B} + \Lambda_{E \cdot X} \hat{E} \hat{X} + \Lambda_{B \cdot X} \hat{B} \hat{X} + \Lambda_{E \cdot B \cdot X} \hat{E} \hat{B} \hat{X}. \quad (3.9)$$

By comparing the signal when $\hat{E}$ and $\hat{B}$ are in the same switch state to the signal when $\hat{E}$ and $\hat{B}$ are in opposite states, and averaging over all other parameter switch states, we find that

$$\{E \cdot B\} = \beta \Lambda_{E \cdot B} + \alpha (\phi_{EDM} + \phi_{back} \Lambda_E + \phi_X \Lambda_{E \cdot X}). \quad (3.10)$$

The all-important $\{E \cdot B\}$ channel is now not only dependent on the EDM phase, but also a number of other phases and undesirable contrast modulations. Fortunately, however, we can derive an exact expression for the EDM phase. By multiplying equation 3.8 with equation 3.9, and collecting terms that multiply ϕ_{EDM} , it can be shown [76] that:

$$\frac{\phi_{EDM}}{\phi_{\delta B}} = \frac{-1}{C_6} (C_1 + C_2 + C_3 - \beta (C_4 + C_5)) \quad (3.11)$$

where

$$C_1 = \{E \cdot B\} (\{\delta B\}^2 - \{X \cdot \delta B\}^2), \quad (3.12a)$$

$$C_2 = \{E \cdot \delta B\} (\{X \cdot \delta B\} \{X \cdot B\} - \{B\} \{\delta B\}), \quad (3.12b)$$

$$C_3 = \{E \cdot X \cdot \delta B\} (\{B\} \{X \cdot \delta B\} - \{\delta B\} \{X \cdot B\}), \quad (3.12c)$$

$$C_4 = \{\delta B\} \{E \cdot B \cdot \delta B\} - \{B \cdot \delta B\} \{E \cdot \delta B\} - \{X \cdot B \cdot \delta B\} \{E \cdot X \cdot \delta B\} \quad (3.12d)$$

$$C_5 = \frac{\{X \cdot \delta B\}}{\{\delta B\}} \left(\{E \cdot \delta B\} \{X \cdot B \cdot \delta B\} - \{X \cdot \delta B\} \{E \cdot B \cdot \delta B\} + \{B \cdot \delta B\} \{E \cdot X \cdot \delta B\} \right), \quad (3.12e)$$

$$C_6 = \{\delta B\}^3 + 2\{X \cdot \delta B\} \{E \cdot \delta B\} \{E \cdot X \cdot \delta B\} - \{\delta B\} (\{X \cdot \delta B\}^2 + \{E \cdot \delta B\}^2 + \{E \cdot X \cdot \delta B\}^2). \quad (3.12f)$$

We call any term beyond $\{E \cdot B\}/\{\delta B\}$ a ‘lineshape correction’. In practice most of the terms above are negligible. In fact $\{B\}\{E \cdot \delta B\}/\{\delta B\}^2$ is the only term which is actually significant, so we approximate equation 3.11 with the equation:

$$\frac{\phi_{EDM}}{\phi_{\delta B}} = \left(\frac{\{E \cdot B\}}{\{\delta B\}} - \frac{\{B\}\{E \cdot \delta B\}}{\{\delta B\}^2} \right). \quad (3.13)$$

The correction in equation 3.13 neglects terms of $\mathcal{O}(\Lambda^2)$ as well as the correlated amplitude modulations $\Lambda_{E \cdot B}$ and $\Lambda_{E \cdot X}$, which are negligible anyhow. As $f(\phi)$ was defined to be linear we neglected all higher order terms in ϕ . However, for our published result [75], we found that the $\{B\}\{E \cdot \delta B\}/\{\delta B\}^2$ term was four times smaller than the statistical uncertainty in $\{E \cdot B\}/\{\delta B\}$. This means that any further correction terms that would result from a non-linear model for $f(\phi)$ would be completely negligible.

We calculate a value for d_e by averaging equation 3.13 over many blocks. We aim to take equal numbers of blocks in each of the eight possible manual reversal states. Recall that the electric and magnetic field manual reversals change the assignment between field direction and waveform switch state. So these reversals flip the sign of the channels containing E and B . As both terms in equation 3.13 contain E and B , we must therefore sign each according to the manual states M_E and M_B so that the channel values do not cancel when averaged over an entire dataset. To calculate an EDM value we apply the equation

$$d_e = \widehat{M}_B \widehat{M}_E \frac{g \mu_B B_{\delta B}}{E_{\text{eff}}} \frac{\phi_{EDM}}{\phi_{\delta B}} \quad (3.14)$$

to the channel values for each block. We define the coefficients $(\widehat{M}_E, \widehat{M}_B) \in \{1, -1\}$ to be positive when their corresponding manual state M_E/M_B is labelled as true.

During the acquisition and analysis of EDM data we temporarily add an unknown, constant offset to each value of d_e in order to blind ourselves from the measured values.

We do this so that we cannot unintentionally influence the final result. Only after a full dataset has been acquired and all systematic tests have been completed do we remove the offset and reveal a final value for d_e . This will be discussed further in the next section.

Before we continue, it is worth noting that equation 3.11 was derived for a phase that depended upon the applied electric field, E , which we call an EDM. In order to calculate the phase that results from some other experimental parameter Y , we simply substitute E with Y in the equations 3.11–3.13.

3.7 Blind Analysis

Scientists have long known that their preconceived notions as to the outcome of an experiment can unintentionally influence the final result. The many subtle ways in which this bias may occur is discussed in references [88] and [89]. The general idea is that if we are fully aware of our measurements during the data acquisition and analysis processes, we might modify our experimental technique to amplify or suppress any number of effects. If, for example, a dataset is filtered, or cut, with full knowledge of what events/measurements are included or not, the experimenter’s choice may suppress a signal, or add significance to one which may only be a statistical fluctuation. A bias might also occur when applying corrections for systematic errors. Without realising we might find only corrections which would shift the measured value in the desired direction. Furthermore, it is an easy mistake to stop searching for sources of systematic error only when the value measured agrees with our preconceptions. If agreement or disagreement is used as justification for completion of experimentation, it is possible that a number of experimental problems could go unnoticed. Therefore we must prevent against a potential experimenter bias by performing our analysis blind to the final result. However, we must not be blind to how changes in experimental parameters and the data analysis affect the final result, if we are to quantify potential systematic uncertainties.

Blind analysis is now routine for most nuclear, particle and precision measurement experiments (see references in [88]). There are several approaches to blind analysis. We use a ‘hidden offset’ blind, that does not affect parameter fitting or any applied corrections. We never directly look at the integrated $\{E \cdot B\}$ channel, and to all calculated EDM values the analysis software adds an unknown constant that was randomly chosen from a Gaussian probability distribution centred on zero, with a standard deviation of 5×10^{-27} e cm.

As discussed in the next section, we collect many thousands of blocks in order to precisely calculate an average value and uncertainty for d_e . Only after all the data has been

collected, and we are satisfied that all possible sources of systematic error have been thoroughly investigated do we reveal the blind offset. Once the blind is removed, no further data analysis is permitted.

3.8 Trimmed Mean and Statistical Uncertainty Calculations




FIGURE 3.6: The distribution of 6194 d_e values that make up the 2011 dataset. For the most part, the distribution is fit well by a Gaussian distribution (red line) with a standard deviation given by the 68.3% statistical uncertainty of the bootstrapped trimmed mean multiplied by $\sqrt{6194}$. The distribution does, however, deviate from the Gaussian fit out in the wings, as demonstrated by the quantile-quantile plot on the right.

Before calculating an average EDM value, we check the distribution of d_e values measured by each block. Figure 3.6(a) plots the distribution of the 6194 d_e values that make up the 2011 published dataset. In general the distribution is fit well by a Gaussian distribution. However, the distribution does deviate from the Gaussian model in the wings of the distribution. Figure 3.6(b) shows a quantile-quantile plot⁵ that compares the measured EDM distribution to a normal distribution. The departures from the $y = x$ line (the solid straight red line) show that there are a number of outliers which do not fit the Gaussian model. The distribution has long wings which might be fitted better by a mixture distribution formed by adding a small Gaussian component to the predominant Gaussian model, which is centred on the same mean but has a much larger standard deviation. Interestingly the rough error bars of these outliers, as calculated from the variance of $\{E \cdot B\}$ within each block⁶, are of a similar size to the error bars of those points which fit the Gaussian model. This would indicate that there is nothing untoward

⁵Quantiles are regular points taken from a cumulative distribution function of a random variable.

⁶Recall that each unique machine switch state is visited eight times within each block, as there are nine modulated experimental parameters that are each modulated between two states over the course of 4096 shots.

about these outliers⁷. One possible explanation is that the noise observed within each block is different to the noise observed across the whole dataset. This agrees with previous observations that a significant $1/f$ component is present in the noise spectrum of stray magnetic fields around the lab [73].

A common technique to estimate the centre of a distribution is to calculate the sample mean and apply the central limit theorem (CLT) to calculate an uncertainty on the mean from the variance of the dataset. Unfortunately outliers present a few problems to this approach: the first is that a single outlying measurement can shift the sample mean by an arbitrarily large amount; the second is that the sample variance is particularly sensitive to outliers, even a small departure from normality can grossly inflate error bar calculations⁸ [90]; the third problem, which is related to the second, stems from the assumption that the distribution of means is Gaussian. This approximation is only satisfied in the limit of large N . Long tails in the sample distribution will slow the convergence of the distribution of means to a Gaussian.

We can combat the first problem by using an estimator which is more robust to outliers than the mean. One option would be to use the sample median. However, this does not make the most efficient use of the data, in fact we would need at least 50% more measurements to achieve similar uncertainty [91]. We compromise efficiency with robustness by calculating the trimmed mean, in which the sample mean is calculated on a truncated dataset, where a small fraction of values is removed from the either end of the sample distribution. Because the deviation from normality is relatively small, we find that 5% is sufficient for calculated mean values to converge. The second and third problems are solved by avoiding the CLT all together. We apply the non-parametric bootstrap method, as was developed by Efron [92], which substitutes complex theoretical analysis with a considerable amount of numerical computation⁹. The fundamental principle is that our sample dataset provides us with the best estimate of the underlying parent distribution; so we may use it as a surrogate distribution and calculate any statistic of interest via a simple Monte Carlo simulation.

We implement the non-parametric bootstrap method as follows: (i) Take N samples with replacement from the measured sample dataset, which contains N independent and identically distributed random samples. (ii) Calculate the 5% trimmed mean of the

⁷Reasons for anomalous points may be due to the laser unlocking, electric field discharges, temporarily blocked laser beams etc. However, these occurrences are normally obvious. Not only would they dramatically inflate the rough error bar calculated for each block, but also other channels such as $\{\delta B\}/\{SIG\}$ applied to the probe PMT, and $\{E\}$ applied the leakage monitors or magnetometer would show as stark discontinuities.

⁸If we apply the CLT to the 2011 dataset, 4% more measurements would be required to calculate the same size uncertainty as that published in [75].

⁹The amount of processing required nowadays is considered relatively small. We can run our bootstrap analysis on the whole dataset in under 30 s.

replicate dataset. (iii) Repeat steps (i) and (ii) a large number of times to generate a probability distribution for this statistic. We find that 5000 repeats is sufficient for our calculations of the centre value and uncertainty to converge. From the distribution of bootstrapped trimmed means we generate a cumulative distribution function. As the generated cumulative distribution is very nearly symmetric, we calculate an average value and statistical uncertainty for the EDM according to:

$$\bar{d}_e \pm \sigma_{\text{stat}} = \frac{d_e(0.5 - p/2) + d_e(0.5 + p/2)}{2} \pm \frac{d_e(0.5 + p/2) - d_e(0.5 - p/2)}{2}, \quad (3.15)$$

where $d_e(x)$ is the value of d_e at percentile x , and p is the size of the confidence interval. Typical values of p are 0.683, 0.90, 0.958 and 0.995, which correspond to the 1–4 σ deviations of a Gaussian distribution. Note that this method does not depend directly on the sample variance which is inflated by outlying values.

3.9 Non-zero Channel Finder

Before removing the blind we also inspect all the analysis channels in each of the detectors to check for unexpected values. With 9 parameter modulations, including $\{SIG\}$, there are $2^9 = 512$ possible analysis channels. We use a Mathematica notebook called ‘Non-zero channel finder.nb’ to list the bootstrapped trimmed mean with its associated error for each channel, on each detector, for each manual state. This notebook also combines the channel averages together, signing the values according to the 8 possible manual state combinations, to test for the influence of external systematic effects. The program ranks the channels in order of statistical significance. We typically choose 3σ as the criterion for significance. Note that a non-zero value is not necessarily a bad thing. You would, for example, expect the unsigned averages of $\{SIG\}$ and $\{\delta B\}$ to be non-zero.

In the next chapter this analysis method is applied to the 2011 dataset. The final result is broken down by source of error and I demonstrate how each error is accounted for in the final uncertainty calculation.

Chapter 4

The 2011 Dataset

When I joined the team working on the EDM experiment at the end of the 2009, it was known that the experiment suffered from a major systematic error — namely a variation of d_e across the TOF profile. Overcoming this error formed a large part of my work during the first year of my PhD. Like many subtle processes, this effect would only show itself after a week’s worth of data acquisition. What was particularly frustrating was that it didn’t seem to be dependent on any experimental parameter. We would detune one parameter, take data for a week, then see if the result was statistically different from the previous week. Eventually we tracked down this effect to be linked to a voltage offset on the HV plates, which created a correlation between d_e and $\{RF1F\}$ (see section 4.3.2). I tell this anecdote to highlight the fact that measuring the EDM is a rather slow process — not only must we run the machine continually for three months to get enough data to set a world-leading limit, we must also check that our result is not dependent on any systematic influences. In fact we spent approximately four times as much time investigating systematic effects as we did taking ‘normal’ EDM data. To check for a dependence we cycle through all our experimental parameters, detuning each in turn from its ideal setting, then take data until we can be sure of the absence (or presence) of an effect. It is worth emphasising again that our approach to measuring the EDM is highly empirical. When we observe an effect we must remain pragmatic: How does it affect the result? Can we correct for it using some combination of analysis channels? What is the worst case scenario and what systematic uncertainty must we associate with this effect? If we tried to develop an in depth physical model for every effect, simulating and measuring the intricacies of each, we would never make any progress.

By September 2010 we felt that all effects were sufficiently under control that we could begin taking data which would directly contribute to a publishable result. In early 2011 we removed the blind and published our result [75]. In this chapter I present a

thorough analysis of the 2011 result, addressing each systematic in turn before revealing the final value. I do this partly for dramatic effect, but also to illustrate our process of blind analysis. Many of the details presented in this chapter were recently published in reference [76].

We calculate an average value for d_e from 6194 blocks that were taken under normal running conditions. As described in section 4.2, we make two ‘empirical corrections’ to reduce the impact of magnetic field noise and correct for an unexplained, yet well defined phase that depends upon the frequency of each rf pulse. These corrections are applied to each individual d_e value so affect the statistical uncertainty of our measurement. We account for the possibility of residual effects and other systematic dependencies with a second systematic uncertainty value. This systematic uncertainty is calculated from a series of supplementary tests that were performed with experimental parameters detuned from their ideal settings. The details of each experiment, and the uncertainties assigned to each potential source of systematic error are discussed, in turn, in section 4.3.

4.1 Defining the PMT Gates

Before we can evaluate any analysis channel, we must first define the interval over which the PMT signals are integrated. For the probe PMT signal we choose a $130\,\mu\text{s}$ wide gate centred $2198\,\mu\text{s}$ after the initial YAG laser Q-switch trigger. This gate is illustrated by the shaded region in figure 2.11. We define the pump PMT gate to be a constant fraction of the probe PMT gate times. We set this fraction to the ratio of the pump and probe laser distances from the molecular source ($= 3.842$). This ensures that both PMTs sample the same molecules. As will be discussed later in section 4.6, these gate times maximise the EDM sensitivity. But to fully understand the uncertainty calculation we must first consider the corrections and systematic errors that affect the final result.

4.2 Empirical Corrections

As stated, we make two empirical corrections to account for E -switch correlated magnetic fields, and an unexplained, yet well defined phase correlated with frequency modulation of each rf pulse. These corrections are described in the following two sections.

4.2.1 Ambient Magnetic Field Correction

Stray magnetic fields that change when the electric field direction is reversed are a major concern, as both the EDM interaction and the Zeeman effect perturb the $F = 1$ manifold in exactly the same way. While random magnetic field fluctuations add noise to our measurement, a systematic correlation between the stray field and the E switch will systematically shift our value for d_e .

We can split E -switch correlated stray magnetic fields into two components: a component that is changed by a manual reversal of the HV connections, and a component that is not. The component that does not change can be isolated by averaging $\{E\}_{\text{mag}}$ over both electric field manual states¹. The component that does change with the manual reversal is isolated by signing the values of $\{E\}_{\text{mag}}$ according to the manual state (i.e. by $\widehat{M}_E$) and then averaging. The detector label ‘mag’, may refer to any one of the four magnetometers positioned about the lab.

Only the fluxgate magnetometer placed next to the high voltage relays measures a non-zero magnetic field in $\{E\}$. The 5 m separation between the relays and beam machine, and the magnetic shielding together attenuate this magnetic field considerably. The *Bartington* magnetometer, placed in between the two layers of magnetic shields, measures the average value of $\widehat{M}_E\{E\}_{\text{Bart}}$ to be (-0.6 ± 3.0) pT and the average value of $\{E\}_{\text{Bart}}$ to be (-2.3 ± 3.1) pT, which are both consistent with zero².




FIGURE 4.1: The observed correlation between d_e and E -switch correlated stray magnetic field as measured by the *Bartington* magnetometer. The average shift due to the stray magnetic field is consistent with zero.

¹The magnetometers respond linearly to magnetic field (not sinusoidally), so we should consider $\{E\}_{\text{mag}}$ not $\{E \cdot B\}_{\text{mag}}$.

²As the *Bartington* magnetometer is positioned outside the inner shield you might worry whether it can detect an E -switch correlated magnetic field generated inside the machine. This is a genuine concern, but we argue that internally generated stray magnetic fields are most likely to result from shield magnetisation or leakage currents. We specifically test for these sources (section 4.3.3) and assign each a systematic error bar.

Our measurement of d_e is particularly sensitive to stray magnetic fields that correlate with E , and change when the electric field connections are manually reversed. While the average value of $\widehat{M}_E\{E\}_{\text{Bart}}$ is not statistically significant, we do observe a significant correlation between d_e and $\widehat{M}_E\{E\}_{\text{Bart}}$, as illustrated in figure 4.1. This is not surprising as both are sensitive to the component of the magnetic field noise that correlates with the switching of the electric field direction. A linear fit to the data gives the gradient $(\partial d_e / \partial B_{\text{mag}}) = (4.51 \pm 0.16) \times 10^{-26} \text{ e cm/nT}$. However, we need to apply a correction to this gradient as the error bars in the two variables are comparable. When this is the case ordinary least squares regression underestimates the gradient by constant factor, λ , known as the ‘reliability ratio’. In [93] Carroll states that the reliability ratio is given by

$$\lambda = \frac{\sigma_w^2 - \sigma_u^2}{\sigma_w^2}, \quad (4.1)$$

where, in our case, σ_w is the standard deviation of the $\widehat{M}_E\{E\}_{\text{Bart}}$ measurements, and σ_u is the uncertainty on each individual $\widehat{M}_E\{E\}_{\text{Bart}}$ value. For the data presented here $\lambda = 0.83$. The straight line fit in the figure illustrates the observed correlation, accounting for the factor λ .

We correct for the observed correlation by subtracting the value

$$d_e^{\text{mag}} = \frac{1}{\lambda} \left(\frac{\partial d_e}{\partial B_{\text{mag}}} \right) \widehat{M}_E\{E\}_{\text{Bart}} \quad (4.2)$$

from d_e (equation 3.14), on a block-by-block basis. Overall, this correction reduces the statistical uncertainty in d_e by 3.5%. The average size of the correction is $(-0.3 \pm 1.7) \times 10^{-28} \text{ e cm}$, which is much smaller than our EDM measurement uncertainty. This is of course just a restatement of the average value for $\widehat{M}_E\{E\}_{\text{Bart}}$ as quoted above, but in units of EDM.

4.2.2 RF Detuning Phase Correction

The second correction that we apply is for an interferometer phase change that occurs when the frequency of either rf pulse is stepped. This phase dependency is potentially very dangerous when coupled with an imperfect electric field reversal, as is explained below.

We calculate the phase change that results from the frequency step of each rf pulse in exactly the same way that we extract the EDM phase change that results from reversing

the electric field direction. Explicitly we apply the equation

$$\delta\phi_{\text{rf}}^{(i)} = \delta\phi_{\delta B} \left(\frac{\{RFiF \cdot B\}}{\{\delta B\}} - \frac{\{B\}\{RFiF \cdot \delta B\}}{\{\delta B\}^2} \right) \quad (4.3)$$

to the channel values for each block. As before, $\delta\phi_{\delta B} = 2\pi/32$ is the phase change that results from the calibration step with the magnetic field. From equation 4.3 we can calculate the rate of change of phase with rf detuning, by dividing through by the size of the rf frequency step, $\delta\nu_{\text{rf}}^{(i)} \simeq 3 \text{ kHz}$:

$$\frac{\partial\phi}{\partial\nu_{\text{rf}}^{(i)}} = \frac{\delta\phi_{\text{rf}}^{(i)}}{\delta\nu_{\text{rf}}^{(i)}}. \quad (4.4)$$

When equation 4.4 is applied to the main dataset we find that the average rate of change of phase is different for each rf pulse, and depends upon the direction of rf propagation along the transmission line. We measure the gradients: $(316 \pm 8) \text{ nrad/Hz}$ and $(-139 \pm 7) \text{ nrad/Hz}$ for the first and second rf pulses respectively when the rf is travelling up through the machine ($M_{\text{rf}} = \text{true}$); and $(260 \pm 8) \text{ nrad/Hz}$ and $(-42 \pm 8) \text{ nrad/Hz}$ respectively when the rf is travelling downward ($M_{\text{rf}} = \text{false}$).

Unfortunately we cannot explain the causes of these phase gradients. We have run numerical simulations which included imperfect settings for the rf amplitudes and frequencies, and allowed for magnetic fields to be applied during the rf pulses [74]. We have also tried incorporating realistic changes to the rf frequency, amplitude and polarisation, electric field and magnetic field as the molecules travel through the machine. No model has been able to provide the answer, although we have not yet tried including the oscillating rf electric field, nor the possibility that the optical pumping could set up coherent state. While the rf frequency induced phase is not well understood, it is at least well defined.

The rf detuning is intimately related to the magnitude of the electric field through the (quadratic) Stark shift. We now consider a situation where the magnitude of the electric field changes by an amount, δE , when its direction is reversed. The origin and size of δE in the experiment will be discussed in section 4.3.1. Through the Stark shift, this magnitude step will change the rf frequency by an amount

$$\delta\nu_{\delta E} = \delta E \left(\frac{\partial\nu}{\partial E} \right) \Big|_{E_0}, \quad (4.5)$$

and hence introduce a phase dependence on the electric field orientation. At our operating electric field the Stark gradient $(\partial\nu/\partial E)|_{10 \text{ kV/cm}} = 285 \text{ Hz/(V/cm)}$. This value was

calculated from studies by Sauer, in which the Stark shift of the hyperfine levels in YbF were measured to high precision [36] (see also figure 2.4).

The change in rf frequency correlated with the electric field reversal is directly measured by the channel $\{E \cdot RFiF\}/\{\delta B\}$. We assume this channel has the linear dependence:

$$\frac{\{E \cdot RFiF\}}{\{\delta B\}} = \xi_i \delta E. \quad (4.6)$$

As shown in figure 4.2, we calibrate the constants of proportion, ξ_i , with a separate experiment where the magnitude of the electric field is intentionally stepped when the electric field is reversed. For $18 \mu\text{s}$ rf pulses, $\delta E = \pm 20 \text{ V/cm}$ is more than adequate to measure ξ_i . We find $\xi_1 = (-4.11 \pm 0.04) \times 10^{-3} (\text{V/cm})^{-1}$ and $\xi_2 = (-4.02 \pm 0.06) \times 10^{-3} (\text{V/cm})^{-1}$, which are consistent with each other within error.




FIGURE 4.2: Calibration the constants of proportion, ξ_i , between $\{RFiF \cdot E\}/\{\delta B\}$ and electric field asymmetry, δE . The solid red line is for rf1 and the dashed blue line for rf2.

Substituting equation 4.6 into 4.5 we write the following expression for the frequency shift of each pulse:

$$\delta\nu_{\delta E}^{(i)} = \frac{1}{\xi_i} \frac{\{E \cdot RFiF\}}{\{\delta B\}} \frac{\partial \nu}{\partial E} \Big|_{E_0}. \quad (4.7)$$

Now combining equations 4.7 and 4.4 with the definition $\delta\phi_{\delta E}^{(i)} = (\partial\phi/\partial\nu_{\text{rf}}) \delta\nu_{\delta E}^{(i)}$ we thus derive an expression for the change in phase between the two electric field switch states:

$$\frac{\delta\phi_{\delta E}^{(i)}}{\delta\phi_{\delta B}} = \frac{1}{\xi_i \delta\nu_{\text{rf}}^{(i)}} \frac{\{E \cdot RFiF\}}{\{\delta B\}} \left(\frac{\{RFiF \cdot B\}}{\{\delta B\}} - \frac{\{B\}\{RFiF \cdot \delta B\}}{\{\delta B\}^2} \right) \frac{\partial \nu}{\partial E} \Big|_{E_0}. \quad (4.8)$$

This is, of course, the signature of an EDM. To calculate the false EDM, and apply a correction, we must account for the electric and magnetic field manual reversals as we did in section 3.6. Recall that these manual reversals change the assignment between field direction and waveform switch state. We must, therefore, sign all channels containing E with $\widehat{M}_E$, and channels containing B with $\widehat{M}_B$, so that they do not cancel when averaged

over the entire dataset. Explicitly, we calculate a false EDM for each rf pulse according to:

$$d_e^{\text{rf}} = \widehat{M}_E \widehat{M}_B \frac{g \mu_B B_{\delta B}}{E_{\text{eff}}} \frac{\delta \phi_{\delta E}^{(i)}}{\delta \phi_{\delta B}} \quad (4.9)$$

Like the magnetic field correction, this quantity is subtracted from d_e (equation 3.14) on a block-by-block basis using the measured channels and the values for $B_{\delta B}$, E_{eff} and $\delta \nu_{\text{rf}}^{(i)}$ as recorded for each block. When applied to the main dataset we find that, on average, the rf1 correction is $(5.0 \pm 0.9) \times 10^{28}$ e cm and the rf2 correction is $(0.5 \pm 0.7) \times 10^{28}$ e cm. Both are smaller the statistical uncertainty of our final result.




FIGURE 4.3: The observed correlation between the measured (blinded) value of d_e and applied electric field asymmetry, δE . The solid red line is before the rf detuning phase correlation correction is applied, the dashed blue line is after applying the correction. The solid red line shows a clear dependence on δE whereas the dashed blue line is consistent with a flat line. This shows that the correction appears to remove all dependence on electric field asymmetry.

Figure 4.3 plots measured EDM against δE using data that was taken when the change in electric field magnitude between E states was made intentionally large. The two lines shown represent linear fits to the data before and after the correction is applied. It is apparent that the correction does indeed remove all dependence on δE . However, it is possible that there are other δE effects that are not corrected by equation 4.9. In the next section (4.3.1) I discuss the sensitivity of our measurement to any residual δE effects, and place an upper bound on the systematic shift that would result from the typical value of δE that is present under normal running conditions.

4.3 Systematic Uncertainties

In this section I discuss the experiment imperfections that could introduce sources of systematic error. Where possible, we measured the dependence of d_e on each imperfection by taking supplementary EDM measurements for a range of non-ideal experiment conditions. Combined with conservative estimates of the typical conditions, we use these measurements to place limits on the possible systematic shifts.

This is a particularly powerful method as we don't need to know the exact mechanism for each source of systematic error. By exaggerating an experimental imperfection, we measure the shift that results from all sources of systematic error that associate with that particular imperfection.

There are, however, a couple of imperfections that cannot (easily) be exaggerated, namely leakage currents and electric field rotations that result from bent field plates and patch potentials. To assign a systematic uncertainty to these imperfections, we formulate a likely model for how these imperfections could generate a systematic shift, and consider a worst case scenario. Fortunately neither effect generates a significant systematic shift.

Before considering each imperfection in turn, it should be noted that the EDM measurements presented in this section have been corrected using the expressions derived in section 4.2 so that we can look for further sources of systematic error.

4.3.1 Uncancelled Asymmetric Electric Field Reversals

When the electric field reverses, its magnitude may change by a small amount, δE . We call this an 'asymmetric electric field reversal'. In section 4.2.2 it was demonstrated that asymmetric electric field reversals could generate an EDM through an unexplained, yet well defined phase that depends on rf frequency detuning. We correct for this effect with equation 4.4, but we must account for the possibility of a residual systematic shift. There are many other mechanisms, in addition to the rf phase correlation, through which asymmetric field reversals might generate a false EDM. A straightforward example is that the detection efficiency of the PMTs could vary with the electric field potentials inside the vacuum chamber. A more subtle example, that will be discussed in detail in section 4.3.3, is that the sensitivity of the molecules to stray magnetic fields might depend upon the magnitude of the applied electric field. By intentionally making δE large, we measure the EDM dependence on all possible mechanisms associated with asymmetric electric field reversals.




FIGURE 4.4: The observed correlation between measured value of d_e and applied electric field asymmetry when the plate voltages are offset from ground by $+102.5$ V (see section 4.3.2). The dotted red line is a linear fit to the data taken with the rf is propagating up the machine ($M_{RF}=\text{true}$). The dashed blue is a linear fit to data taken with the rf propagating downwards ($M_{RF}=\text{false}$). The solid green line is a fit to the aggregated data from both rf directions.

To measure the dependence of d_e on δE we made a slight modification to the electric field control software (HardwareController) to boost the voltages applied to the plates for one of the electric field switch states. As shown by the dashed blue line in figure 4.3, under ideal conditions we find no evidence for residual δE shifts. Unfortunately, however, we do observe a slight dependence on δE if the electric field plate voltages are offset from ground by a small amount, $\bar{V}$. We define $\bar{V}$ such that the potential on each plate is $\pm V + \bar{V}$ with respect to ground.

For a number of months we accidentally ran the experiment with the HV supplies set to apply a voltage offset $\bar{V} = 102.5$ V. From the supplementary tests on δE that were performed during this period we measure a gradient of $(-8.0 \pm 3.6) \times 10^{-28}$ e cm/(V/cm). This is shown by the solid green line in figure 4.4. Figure 4.4 also plots the $\bar{V} = 102.5$ V test data separated into each rf manual state. For the data taken with the rf propagating up through the machine the gradient is $(-1.9 \pm 4.4) \times 10^{-28}$ e cm/(V/cm). This value is entirely consistent with no dependence on δE . However, the data taken with the rf propagating downward produces a gradient $(-21.1 \pm 6.5) \times 10^{-28}$ e cm/(V/cm), which differs from zero by 3.2 standard deviations.

The final EDM measurement was eventually taken with the voltage offset carefully zeroed using a high voltage probe. With the offset zeroed we measure the gradient $(-0.5 \pm 8.8) \times 10^{-28}$ e cm/(V/cm). As this value is consistent with no dependence on δE , we apply no further correction to the dataset. However, it is possible that the slight dependence observed with $\bar{V} = 102.5$ V might also apply to the main dataset taken with $\bar{V} = 0$ V. So to set an upper limit on the systematic shift we use half the non-zero gradient measured from the $\bar{V} = 102.5$ V data taken with the rf propagating downward:

$(-11.0 \pm 3.3) \times 10^{-28} \text{ e cm}/(\text{V}/\text{cm})$. We use half the value because the main dataset contains approximately equal numbers of blocks taken in each rf manual state³.

To calculate an uncertainty for all uncorrected δE systematic effects, we multiply this gradient with an estimate of the asymmetry in the main dataset, taken under normal running conditions. We measure this asymmetry by applying equation 4.6 to the measured average value of $\{E \cdot RFiF\}/\{\delta B\}$.

We can divide sources of asymmetry into two categories: those which are internal to the beam machine, and those which are external. Patch potentials on the electric field plates are an example of an internal source, whereas a change in HV supply output when the electric field direction reverses is an example of an external source. For an internal source of asymmetry, a manual reversal of the HV connections changes the electric field switch state of the larger electric field⁴ (the asymmetry, δE , reverses). In contrast, for an external source, the switch state of the larger electric field does not change with the electric manual reversal (the asymmetry, δE , remains unchanged).

Recall that the channel $\{E \cdot RFiF\}$ measures the change in rf detuning between electric field switch states. The Stark shift is quadratic in E_0 , so rf detuning depends only upon the magnitude of the electric field, not direction. As a consequence, for an internal source of asymmetry, the sign of $\{E \cdot RFiF\}$ changes when the electric field connections are manually reversed. For an external source of asymmetry the sign remains unchanged. We can, therefore, isolate internal and external sources of asymmetry by respectively choosing whether or not to sign $\{E \cdot RFiF\}$ according to $\widehat{M}_E$, and averaging over both electric field manual states.

When the blocks are averaged without signing, we find that the external asymmetry measured in both rf regions are consistent with each other, as would be expected. The weighted mean of the two values measured gives $\delta E_{\text{ext}} = (0.463 \pm 0.015) \text{ V}/\text{cm}$. When the blocks are signed according to $\widehat{M}_E$ and averaged, to measure internal sources of asymmetry, we find $\delta E_{\text{int}}^{(1)} = (0.21 \pm 0.02) \text{ V}/\text{cm}$ when measured with the first rf pulse and $\delta E_{\text{int}}^{(2)} = (-0.21 \pm 0.02) \text{ V}/\text{cm}$ when measured with the second. Note the change in sign, the two values are not equal. This most likely indicates patch potentials which are local to each rf pulse region.

The systematic phase shift, measured predominantly by the channel $\{E \cdot B\}$ (equation 3.13), depends only upon the change in magnitude of the electric field. So, as above, for an internal source of asymmetry $\{E \cdot B\}$ will change sign following a manual reversal of

³We could have assumed the average gradient $(-8.0 \pm 3.6) \times 10^{-28} \text{ e cm}/(\text{V}/\text{cm})$, but half the gradient for the downward propagating rf data produces a more conservative value of $(-11.0 \pm 3.3) \times 10^{-28} \text{ e cm}/(\text{V}/\text{cm})$.

⁴Note that we are only interested in changes in magnitude not direction.

the electric field connections, whereas for an external source of asymmetry the sign will remain unchanged. Recall that to extract a value for d_e we sign $\{E \cdot B\}$ by $\widehat{M}_E$ and $\widehat{M}_B$ to account for the fact that the electric and magnetic field manual reversals change the assignment between waveform state and field direction (equation 3.14). So when $d_e \propto \widehat{M}_E \widehat{M}_B \{E \cdot B\}$ is averaged over both electric field manual states, the false EDM due to external sources of asymmetry cancels. However, the false EDM due to internal sources does not.

Following the above arguments, to limit the possible shift on d_e , we take the measurements of δE_{int} as typical values of δE . The regions spanned by the molecular cloud during each rf pulse make up a significant fraction of the total interaction region ($\sim 1/8$ and $\sim 1/4$ respectively), and the asymmetries measured by each pulse have opposite sign, so it is likely that the effects of these sources will cancel to some extent. We, therefore, estimate that $|\delta E_{\text{int}}| < 0.1 \text{ V/cm}$ when averaged over whole interaction region. Combining this value with the gradient, $(-11.0 \pm 3.3) \times 10^{-28} \text{ e cm/(V/cm)}$, as detailed above, we limit the total shift due to all δE effects to be:

$$\sigma_{\delta E} < 1.1 \times 10^{-28} \text{ e cm.} \quad (4.10)$$

4.3.2 Electric Field Voltage Offset From Ground

In the previous section, we accounted for the possibility that the magnitude of the electric field might change when the electric field is switched. It is also possible that the direction of the applied field may not reverse perfectly.

The electric field plates are surrounded by a grounded support structure and magnetic shield. Ideally we would like to charge the electric field plates symmetrically to the potentials $\pm V$ relative to ground. However, if the HV settings are set imperfectly, the plates become charged to the potentials $\pm V + \bar{V}$. We call the mean plate voltage, $\bar{V}$, the ‘offset voltage’. Near the edges of the HV plates the electric field lines bend towards the grounded structures. If $\bar{V}$ is not carefully zeroed, the distribution of these fringe field lines change when the HV relays change state.

There are a few ways that a change in the field distribution could introduce a systematic shift. The most obvious mechanism is through the magnetic field. As will be discussed in the next section, the molecules are only sensitive to the component of the magnetic field parallel to the local electric field. So if the electric field distribution were to change between switch states, the Zeeman interaction phase, $g \mu_B B T / \hbar$, would become E -state dependent and hence generate a false EDM signal. Similarly, the imparted rf phase depends upon the projection of the oscillating magnetic field in the plane perpendicular

to the local electric field. This phase would also masquerade as an EDM phase if the field distribution were to change with the E -switch. The geometric phase (discussed in section 4.3.4) also depends upon the exact field distribution.

We make sure that the molecules are well within the parallel plates before subjecting them to rf. This ensures that the interferometer phase is accrued over a region of the plates where the electric field is not significantly affected by the surrounding grounded structures. Consequently, no numerical simulations (which account for imperfections within the machine) have been able to produce a significant EDM as a result of realistic changes in the electric field distribution⁵. However, it must be noted that these simulations are not complete, as they do not account for the optical pumping stage in the interferometer.

To empirically test for a $\bar{V}$ dependence we acquired data with the large applied offset voltages: $\bar{V} = -1000.5$ V, $\bar{V} = +102.5$ V and $\bar{V} = +1015.0$ V. The offset voltages were generated by biasing the power supply outputs. For example, to set $\bar{V} = +1$ kV, we adjusted the power supplies to give $V_{E+} = +7$ kV and $V_{E-} = -5$ kV. The voltages were measured at the HV feedthroughs using a sensitive HV probe, with the plates connected.

Unfortunately we observe two distinct effects that result from a non-zero offset voltage: 1) a direct correlation between d_e and the detuning of the first rf pulse, $\delta_{\text{rf}}^{(1)}$, 2) an overall shift in d_e , that is a function of only $\bar{V}$. We model the shift from these two effects with the equation

$$d_e^{\bar{V}} = f_1(\bar{V}) \delta_{\text{rf}}^{(1)} + f_0(\bar{V}). \quad (4.11)$$

I will consider each term separately, and limit the impact of each on the main dataset, that was taken with $\bar{V} = 0.0$ V.

Incidentally, it was hoped that we could save the data taken with $\bar{V} = 102.5$ V by applying a correction of the form given in equation 4.11. The first term, $f_1(102.5) \delta_{\text{rf}}^{(1)}$, is easily calculated from the $\bar{V} = 102.5$ V dataset itself (see below). However, to calculate the second term, $f_0(102.5)$, we require the exact functional form of $f_0(\bar{V})$. There are many models for $f_0(\bar{V})$ which would fit our data, so we cannot reliably apply a $\bar{V}$ correction to the $\bar{V} = 102.5$ V data. We therefore cannot use the $\bar{V} = 102.5$ V data in our final measurement of d_e .

Correlation with RF Detuning

When $\bar{V} \neq 0$ we see a systematic variation in d_e over the TOF profile. Figure 4.5 shows the variation we observe when $\bar{V} = 102.5$ V. We find that the variation is directly proportional to the detuning of the first rf pulse, $\delta_{\text{rf}}^{(1)}$. In section 4.2.2 I discussed an

⁵The field distributions were taken from simulations in *Comsol* by Tarbutt.

interferometer phase that results from an rf detuning. This phase mimics an EDM induced phase through a change in the rf detuning between electric field switch states. The correlation $d_e^{\text{rf}} \propto \{E \cdot RFiF\}\{B \cdot RFiF\}$ (equation 4.9), results due to a change in magnitude of the electric field between switch states. The effect discussed in this section results from an imperfect reversal of the direction of the electric field. This causes a systematic shift in d_e that depends only on the average detuning of the first rf pulse. When $\bar{V} \neq 0$ we find that $d_e \propto \{RF1F\}$ but not $\{E \cdot RF1F\}$. We find no statistically significant correlation between d_e and the detuning of the second rf pulse.




FIGURE 4.5: Variation in d_e across the TOF profile observed when $\bar{V} = +102.5 \text{ V}$. The range of the horizontal axis is set to the interval of the analysis gate specified in section 4.1.

When the molecules are subjected to the first rf pulse the molecular cloud spans approximately 10 cm of the HV plate length. Over this region the electric field varies by approximately 50 V/cm due to a slight bend in the field plates. Through the Stark shift, the electric field variation results in a 15 kHz rf detuning over the molecular cloud. As shown in appendix A, the detuning of the first rf pulse, $\delta_{\text{rf}}^{(1)}$, is linearly proportional to $\{RF1F\}$ according to the relation

$$\frac{\{RF1F\}}{\{\delta B\}} = \frac{-2 \delta_{\text{rf}}^{(1)} \delta \nu_{\text{rf}}^{(1)} \tau^2}{\pi^3} \frac{B_0}{B_{\delta B}}, \quad (4.12)$$

where τ and $\delta \nu_{\text{rf}}^{(1)}$ are the pulse length and frequency step respectively. As before, B_0 and $B_{\delta B}$ are the amplitudes of the two magnetic field steps. As $\delta_{\text{rf}}^{(1)}$ varies over the molecular cloud we observe a variation in $\{RF1F\}$ over the probe PMT TOF profile.

By dividing the probe PMT TOF profile into many narrow gates and analysing each gate individually, we obtain d_e measurements for a range $\{RF1F\}/\{\delta B\}$ values. We convert the $\{RF1F\}/\{\delta B\}$ values to rf detunings using equation 4.12. Figure 4.6 shows the correlations between d_e and $\delta_{\text{rf}}^{(1)}$ for the four values of $\bar{V}$. Fitting the data for each offset voltage with a straight line provides the gradients, $f_1(\bar{V})$, as quoted in the figure caption.




FIGURE 4.6: The direct proportionality between d_e and the detuning of the first rf pulse, as calculated via subdivision of the TOF profile. A linear fit to each dataset gives the gradients: (a) (43.7 ± 1.8) , (b) (-43.7 ± 2.0) , (c) (-14.3 ± 2.0) , (d) (0.3 ± 1.1) . Each are quoted in units of $10^{-28} \text{ e cm/kHz}$.




FIGURE 4.7: Variation in d_e across TOF profile observed when $\bar{V} = 0.0 \text{ V}$. The range of the horizontal axis is set to the interval of the analysis gate specified in section 4.1.

For the $\bar{V} = \pm 1 \text{ kV}$ data we measure the values of $f_1(\bar{V})$ to be opposite yet equal in magnitude, but the increase in slope is non-linear. We find no dependence of $f_1(\bar{V})$ on the manual state of the machine. Fortunately, when the offset voltage is carefully zeroed we find no evidence for a dependence on $\delta_{\text{rf}}^{(1)}$, as the gradient $f_1(0.0) = (0.3 \pm 1.1) \times 10^{-28} \text{ e cm/kHz}$ is consistent with zero within error. As we expect, with no clear dependence on $\delta_{\text{rf}}^{(1)}$, we find no variation in d_e across the TOF profile, as

shown in figure 4.7.

As we observe no dependence of d_e on $\delta_{\text{rf}}^{(1)}$ when $\bar{V} = 0.0 \text{ V}$, we do not apply a correction for this effect. However, $f_1(\bar{V})$ depends quite strongly on $\bar{V}$, so we must place a limit on the systematic shift due to a possible residual correlation. To calculate this we require the average detuning of the first rf pulse, $\bar{\delta}_{\text{rf}}^{(1)}$. We calculate this by integrating $\{RF1F\}(t)/\{\delta B\}(t)$ over the (wide) analysis quoted in section 4.1. From the main dataset (taken with $\bar{V} = 0.0 \text{ V}$) we measure $\bar{\delta}_{\text{rf}}^{(1)}$ to be⁶ $(1.217 \pm 0.025) \text{ kHz}$. Multiplying this value with the measured correlation for the $\bar{V} = 0.0 \text{ V}$ data, $f_1(0.0) = (0.3 \pm 1.1) \times 10^{-28} \text{ e cm/kHz}$, we limit the systematic shift from the rf detuning correlation to be⁷:

$$\sigma_{f_1} < 1.3 \times 10^{-28} \text{ e cm.} \quad (4.13)$$

Note that the value for $\bar{\delta}_{\text{rf}}^{(1)}$ is significantly non-zero. This is because $\{RF1F\}(t)/\{\delta B\}(t)$ varies quite significantly over the probe PMT TOF profile, and the experimental parameter locks used a slightly different analysis gate to that quoted in section 4.1.

Overall Shift with $\bar{V}$

We can correct the data where $\bar{V} \neq 0$ for the detuning correlation described above. This lets us observe the remaining dependence on $\bar{V}$, as described by the function $f_0(\bar{V})$ in equation 4.11. To correct each dataset we subtract the product $f_1(\bar{V}) \bar{\delta}_{\text{rf}}^{(1)}$ from the measured average value of d_e . As above, the average detuning of the first pulse, $\bar{\delta}_{\text{rf}}^{(1)}$, is calculated from $\{RF1F\}/\{\delta B\}$ using the (wide) analysis gate specified in section 4.1. The measured gradients, $f_1(\bar{V})$, are as quoted in the caption of figure 4.6.

Figure 4.8 shows the remaining dependence on $\bar{V}$ that we observe after applying the correction for the rf1 detuning correlation. As can be seen, the data is consistent with a linear dependence on $\bar{V}$. A linear fit to the data, shown by the solid red line in the figure, measures a gradient of $(-0.099 \pm 0.016) \times 10^{-28} \text{ e cm/V}$.

With the HV probe we are able to zero $\bar{V}$ to well within 1 V. So for the main dataset we can limit systematic shift due to the remaining $\bar{V}$ dependence to be:

$$\sigma_{f_0} < 0.1 \times 10^{-28} \text{ e cm.} \quad (4.14)$$

⁶The uncertainty here is limited solely by the uncertainty in $\{RF1F\}/\{\delta B\}$. We ignore the uncertainty in the conversion factor between $\{RF1F\}/\{\delta B\}$ and $\bar{\delta}_{\text{rf}}^{(1)}$ as this coefficient cancels in the product $f_1(0.0) \bar{\delta}_{\text{rf}}^{(1)}$.

⁷The value for σ_{f_1} quoted here is slightly different to the value published in [75], in which there is a minor error of no consequence.




FIGURE 4.8: Variation in the average value of d_e in addition to the variation across the TOF profile. The gradient of the red line is $(-0.099 \pm 0.016) \times 10^{-28}$ e cm/V.

Combining this uncertainty with the systematic uncertainty due to the rf1 detuning correlation, we limit the total systematic uncertainty on all $\bar{V}$ effects to be:

$$\sigma_{\bar{V}} = \sqrt{(\sigma_{f_1})^2 + (\sigma_{f_0})^2} < \sqrt{1.3^2 + 0.1^2} \times 10^{-28} \text{ e cm} < 1.30... \times 10^{-28} \text{ e cm}. \quad (4.15)$$

4.3.3 Magnetic Fields

Magnetic fields that change⁸ when the electric field is switched are a potential source of systematic error. Fortunately, the molecules are only sensitive to the z component of such magnetic fields, because the tensor Stark shift dominates over any Zeeman interactions. This can be seen from a second order perturbation expansion (see [70]), where the energy splitting between the $|1, \pm 1\rangle$ states, $\Delta U_{\pm}$, is given by:

$$\Delta U_{\pm} \simeq 2 d_e E_{\text{eff}} + 2 g \mu_B B_z - \frac{(g \mu_B B_{\perp})^2 g \mu_B B_z}{\hbar^2 (\omega_{11} - \omega_{10})^2}. \quad (4.16)$$

Note that the third term is suppressed by the square of the Stark splitting between $|1, \pm 1\rangle$ and $|1, 0\rangle$, $(\omega_{11} - \omega_{10})$. Recall that an electric field shifts $|1, \pm 1\rangle$ equally; only a magnetic field and an EDM can break the symmetry. Unfortunately equation 4.16 is not a good approximation for the typical magnetic fields present inside the machine. In particular $B_{\perp} \sim B_z$. To accurately calculate the splitting due to perpendicular magnetic fields we

⁸Note that the magnetic fields do not necessarily need to reverse. In fact, magnetic fields that are perpendicular to the applied electric field do not change the splitting between the $|1, \pm 1\rangle$ states if reversed perfectly.

must numerically diagonalise the full Hamiltonian:

$$\frac{\hbar}{\sqrt{2}} \begin{pmatrix} 0 & -\Delta_x & \sqrt{2} \Delta_z & \Delta_x \\ -\Delta_x & \sqrt{2}(\omega_{11} - \Delta_z) & \Delta_x & 0 \\ \sqrt{2} \Delta_z & \Delta_x & \sqrt{2} \omega_{10} & \Delta_x \\ \Delta_x & 0 & \Delta_x & \sqrt{2}(\omega_{11} + \Delta_z) \end{pmatrix} \quad (4.17)$$

as written in the z -basis for a state of increasing F and m_F . This Hamiltonian accounts for the magnetic fields $B_x \hat{x}$ and $B_z \hat{z}$ with the the respective shifts $\Delta_x = g \mu_B B_x / \hbar$ and $\Delta_z = g \mu_B B_z / \hbar$. The matrix is slightly different if $B_x \hat{x}$ is replaced with $B_y \hat{y}$, but its overall effect is the same, as there is nothing special about either perpendicular axis. As illustrated in figure 2.5, ω_{11} and ω_{10} are the respective Stark shifted separations of $|1, \pm 1\rangle$ and $|1, 0\rangle$ from $|0, 0\rangle$. For a 10 kV/cm field, as normally applied to the molecules, $\omega_{11} = 2\pi \times 172.954$ MHz and $\omega_{10} = 2\pi \times 164.754$ MHz.

The $\phi = \pi/4$ magnetic field step, $B_0 = 13.6$ nT, induces a $|1, +1\rangle \leftrightarrow |1, -1\rangle$ splitting of 380 Hz. As shown in figure 4.9, adding a 100 nT B_x field only increases this splitting by a further 0.06 mHz. For comparison, an EDM of 1×10^{-28} e cm, an order of magnitude smaller than the current experimental upper limit, would induce a splitting of 0.7 mHz. This demonstrates how insensitive our measurement is to $B_\perp$. As perpendicular magnetic fields inside the machine are typically < 2 nT in magnitude, the splitting that they induce is completely negligible in comparison to the splitting of even a relatively small EDM. It is worth emphasising, however, that a B_z of only 25 fT, that reverses with $\hat{E}$, would imitate a 1×10^{-28} e cm EDM, furthermore this false EDM would have a linear dependence on B_z . The splitting due to $B_\perp$ is highly non-linear as can be seen in figure 4.9.




FIGURE 4.9: Induced splitting between $(F = 1, m_F = \pm 1)$ due to perpendicular magnetic fields *in addition* to the splitting due to a 13.6 nT magnetic field along z .

This plot was generated using a 10 kV/cm electric field along z .

E-switch Correlated Magnetic Fields Along $\hat{z}$

As discussed in section 4.2.1, averaged over the main dataset, the *Bartington* magnetometer measures no statistically significant magnetic field correlated with the switching of the electric field direction. However, we should be cautious of internal sources of *E*-switch correlated magnetic, as they are not cancelled by the manual field reversal. Also the magnetometer does not measure internal fields accurately since they are attenuated by the inner magnetic shield⁹. There are two likely internal sources of *E*-switch correlated magnetic field: leakage current between the HV plates, and magnetic shield magnetisation.




FIGURE 4.10: An illustration of a conservative model for leakage current flow. The purple arrows indicate the leakage current path.

Averaged over the main dataset, the leakage monitors measure a leakage current of less than 1 nA. We can consider a model where all this current flows along the path illustrated in figure 4.10. In this model, 1 nA of leakage would generate a 5 fT switching magnetic field along z , when averaged over the whole interaction region. This model is very conservative as the z -component of the magnetic field would cancel if current flow was restricted to the zy -plane. It is also unlikely that all the current would flow all the way around the edges as shown. However, with this worst case scenario, we limit the systematic shift due to leakage currents to be:

$$\sigma_{\text{leak}} < 0.2 \times 10^{-28} \text{ e cm.} \quad (4.18)$$

To test the extent to which the shields are magnetised by the currents which charge the HV plates, we used a spare set of magnetic shields with the same radius as the inner

⁹Recall that the magnetometer is situated in between the two layers of magnetic shields.

magnetic shield¹⁰. A small loop of wire was fed through the aperture in the middle of the shields in place of the HV plates. For 7 s we drove DC current through the loop then measured the magnetic field along z , in the centre of the shield, for a subsequent 20 s. The current driven was alternated between $\pm 500 \mu\text{A}$ to simulate the switching of the electric field direction. The change in magnetic field correlated with the current reversal was measured to be $(39 \pm 42) \text{ pT}$. Scaling this down to the normal change in current, which is kept below $5 \mu\text{A}$ with the 14 s long E -switch period, shield magnetisation would generate a systematic error:

$$\sigma_{\text{shield}} < 0.25 \times 10^{-28} \text{ e cm.} \quad (4.19)$$

Perpendicular Magnetic Fields

It is clear that neither source of magnetic field stated above could change $B_{\perp}$ enough to generate a significant systematic shift to the EDM value. There is another source of magnetic field which is worth considering. As the molecules move through the applied electric field, $E_0 \hat{z}$, with a velocity, $v \hat{y}$, a magnetic field $\vec{E} \times \vec{v} / c^2 = B_v \hat{x}$ is generated in the rest frame of the molecules. When combined with a static stray field, $B_s \hat{x}$, the molecules experience a magnetic field $B_x = B_s + B_v$ in $\hat{x}$ that changes size when the electric field is reversed. Fortunately the effects of this motional magnetic field are suppressed by the tensor Stark shift. Even with a 30 nT stray magnetic field aligned perfectly along x the systematic shift is

$$\sigma_{\text{mot}} < 5 \times 10^{-32} \text{ e cm.} \quad (4.20)$$

This was calculated with a molecular velocity of 590 m/s, and the electric and magnetic fields: $E_0 = \pm 10 \text{ kV/cm}$ and $B_z = B_0 = 13.6 \text{ nT}$.

It is possible that a completely static perpendicular magnetic field may generate an EDM when combined with an asymmetric electric field reversal. This is because the Stark suppression would change with the magnitude of the applied electric field. However, even with conservative estimates of δE and $B_{\perp}$, the effect is completely negligible. Furthermore, we account for this mechanism with the asymmetric field reversal tests that were discussed in section 4.3.1. We see no evidence for this effect.

To be sure that perpendicular magnetic fields did not influence our measurements, we took EDM data with large magnetic fields applied, in turn, along x and y . These tests were taken with $E_0 = 2.5 \text{ kV/cm}$, which is four times smaller than the usual running field, so the Stark suppression was considerably reduced (the $|1, \pm 1\rangle \leftrightarrow |1, 0\rangle$ splitting

¹⁰These shields were purchased in preparation for a machine upgrade to allow a longer interaction period.

is eight times smaller at this field). No electric field asymmetry was applied, as we did not want to stimulate some other systematic effect. We took data with $B_y \sim \pm 100$ nT and $B_x = \pm 500$ nT. The value for B_y is only an estimate as it was calculated by combining a magnetometer measurement with an estimate of the inner shield longitudinal shielding factor. To estimate B_x we assumed the same current coefficient as for the B_z coils, which are wound in exactly the same way as the B_x coils. From the B_x test we obtain the gradient $(0.3 \pm 0.8) \times 10^{-28}$ e cm/nT, and from the B_y test the gradient $(-0.2 \pm 0.3) \times 10^{-28}$ e cm/nT. Neither are statistically significant. From the typical bias fields required to centre the interference fringes we infer that stray fields within the machine are < 2 nT in amplitude. We verified this with offline magnetometer measurements taken inside the upper chamber through the pump PMT vacuum chamber port. Combining the gradients with the typical offset, we limit the EDM shifts due to perpendicular magnetic fields to be $< 0.6 \times 10^{-28}$ e cm for B_x and $< 0.4 \times 10^{-28}$ e cm for B_y . However, these limits are very conservative as increasing the electric field to 10 kV/cm should make us at least an order of magnitude less sensitive to perpendicular magnetic fields. The δE tests already accounted for this systematic, so we do not separately include these limits in the systematic uncertainty of the final result.

Driving RF Transitions in Magnetic Fields

The equations in section 2.2.3 assume the rf transitions are driven in regions that are free of static magnetic fields. However, in practice, both rf transitions are driven in the presence of the applied magnetic field, $(B_0 + B_{\delta B}) \hat{z}$, as well as a possible stray magnetic field, $B_s \hat{z}$. In his thesis [74], Kara derived a set of equations that account for this, which I have verified with numerical simulations. The effects are twofold: 1) Fringe contrast is reduced, as the applied magnetic field splits the $|1, \pm 1\rangle$ states, causing the rf to become detuned from each level; 2) The acquired phase is increased fractionally, as the split $|1, \pm 1\rangle$ states gain a phase difference during each rf pulse in addition to the phase difference gained over the free evolution period.

Driving rf transitions within the applied magnetic field does not produce a fake EDM unless the magnitude of the electric field changes upon reversal, and the rf is detuned from the mean separation between $|0, 0\rangle$ and $|1, \pm 1\rangle$. However, even for considerable rf detunings, the fake EDM is many orders of magnitude smaller than 10^{-28} e cm. A much larger fake EDM is generated if an uncanceled stray magnetic field, $B_s \hat{z}$, is present. In his thesis [74], Kara shows that the systematic error is approximately linearly proportional to the rf detuning, electric field asymmetry, and stray magnetic field. For 20 μ s rf pulses:

$$d_e^{\text{rf}} = 8 (\delta_{\text{rf}}^{(1)} + \delta_{\text{rf}}^{(2)}) \delta\nu B_s \times 10^{-29} \text{ e cm kHz}^{-2} \text{ nT}^{-1}, \quad (4.21)$$

where $\delta\nu$ is the change in the mean separation between $|1, \pm 1\rangle$ and $|0, 0\rangle$ that results from an imperfect electric field reversal, and $\delta_{\text{rf}}^{(i)}$ is the detuning of the i th rf pulse. Both $\delta\nu$ and $\delta_{\text{rf}}^{(i)}$ are measured in kHz.

If we take the values $\delta_{\text{rf}}^{(i)} = 1$ kHz, $\delta\nu = 0.014$ kHz and $B_s = 2$ nT as are typical for the experiment, then $d_e^{\text{rf}} = 0.04 \times 10^{-28}$ e cm. This value is small enough to ignore. As with other magnetic field effects, this systematic effect is accounted for with the electric field asymmetry test, so no further uncertainty is added to the final result.

Electric Field Dependence of the Landé g -Factor

Applying a magnetic field, $B_z \hat{z}$, breaks the degeneracy of the $|1, \pm 1\rangle$ states, causing them to acquire a phase difference $\Delta\phi_B = 2g\mu_B B_z T/\hbar$ over the time interval, T . g is the Landé g -factor that accounts for the total angular momentum of the molecules. For the unperturbed $F = 1$ states, $g = 1.001\dots$. Its exact value is not particularly important, as we adjust the applied magnetic field to apply the $2\pi/4$ and $2\pi/32$ phase steps. It is, however, important that the g -factor does not change with the electric field, as this also could generate a fake EDM when coupled with an asymmetric electric field reversal, δE , and an uncanceled stray magnetic field, $B_s \hat{z}$. Electric fields mix angular momentum states, so it is at least plausible that g might change if the magnitude of the electric field were to change when its direction is reversed. By considering the difference in phase acquired in each E -state, it is straightforward to show that the false EDM would be

$$d_e^g = \frac{\gamma_g \mu_B B_s \delta E}{E_{\text{eff}}}, \quad (4.22)$$

where $\gamma_g = (\partial g / \partial |E|)|_{E_0}$ is the gradient of g with respect to electric field asymmetry.

As documented in his thesis [74], Kara measured the variation of g by recording interference curves similar to figure 2.16 at various applied electric fields in the range (3–14) kV/cm. He found no systematic dependence of the g -factor on E_0 . The maximum possible variation that is consistent with the data is $\gamma_g^{\text{max}} = 3 \times 10^{-6} (\text{V/cm})^{-1}$. We can measure the uncanceled stray magnetic field with the expression $\widehat{M}_B B_{\delta B} \{B\} / \{\delta B\}$. Averaged over the whole dataset this gives $B_s = (-136 \pm 4)$ pT. Recall from section 4.3.1, $\delta E \sim 0.1$ V/cm. By substituting these values into equation 4.22 we limit the shift due to g -factor variation to be:

$$\sigma_g < 1.7 \times 10^{-31} \text{ e cm}, \quad (4.23)$$

which is small enough to ignore.

The systematic shift due to perpendicular magnetic fields is a little less straightforward. If we assume that the g -factor dependence on δE in $\hat{x}$ and $\hat{y}$ is no larger than in $\hat{z}$, we can estimate the potential size of the systematic shift by calculating how the eigenvalues of equation 4.17 change for a small change in the electric field. If we assume that $\delta E = 0.1 \text{ V/cm}$ and $E_0 = 10 \text{ kV/cm}$, even for a 30 nT perpendicular stray magnetic field in $\hat{x}$, the resulting systematic shift is only $8 \times 10^{-35} \text{ e cm}$, which is completely negligible.

4.3.4 Geometric Phase

In section 2.1, it was shown that, during the free evolution period, the $|1, \pm 1\rangle$ states gain a phase difference

$$\Delta\phi = \frac{2}{\hbar} (g \mu_B B_z - d_e E_{\text{eff}}) T, \quad (4.24)$$

due to electric and magnetic fields oriented along $\hat{z}$. The magnetic field, $B_z \hat{z}$, includes both the applied magnetic field, $(B_0 + B_{\delta B}) \hat{z}$, as well a possible uncanceled stray magnetic field, $B_s \hat{z}$. If the magnitude of the electric and magnetic fields vary over the length of the interferometer, B_z and E_{eff} become average amplitudes, as integrated over the region traversed by the molecules during the free evolution period. As was discussed in the previous sections, there are many mechanisms which would cause a systematic shift through a change in the first term, $g \mu_B B$, when the electric field is reversed.

We are also concerned about field rotations along the beam line, as they can generate a ‘geometric’ phase difference, $\Delta\phi^{(g)}$, between the $|1, \pm 1\rangle$ states in addition to the ‘dynamic’ phase difference, $\Delta\phi$, given in equation 4.24. If the geometric phase changes when the electric field is reversed, there would be a systematic shift in the measured value of d_e .

Geometric phase was first demonstrated by Berry in 1984 using a rotating magnetic field [94]. We need only consider rotations of the electric field, as the molecules are only sensitive to the component of the magnetic field in the axis of the local electric field (section 4.3.3). For a spin-one system corresponding to the $F = 1$ state of the YbF molecules, Tarbutt *et al.* show [95] that in the adiabatic limit (where the rate of rotation of the electric field is small in comparison to the Stark induced $|1, \pm 1\rangle \leftrightarrow |1, 0\rangle$ splitting) the $|1, \pm 1\rangle$ states gain a geometric phase difference

$$\Delta\phi^{(g)} = 2\Omega, \quad (4.25)$$

where Ω is the solid angle swept out by the electric field vector in the reference frame of the molecules. This approximation is well satisfied for the electric field present between the HV plates.

In section 4.3.2 it was stated that in the fringe fields near the edges of the HV plates the electric field rotates towards the grounded surrounding structures. As we carefully balance the potentials on each plate, this rotation is small enough to neglect for the molecules that participate in the experiment.




FIGURE 4.11: Electric field rotations due to: (a) bent HV plates and (b) a patch potential on one of the field plates. The purple arrows indicate the direction of the electric field. The large gray arrow shows the path of the molecular beam. These drawings are purely illustrative and are not drawn to scale.

Warped field plates are another potential source of field rotation. A map of the electric field along the plate length (along the y -axis) [84] tells us that the HV plates are slightly bent. From the field map we calculate that the field rotations about the x -axis are less than $\pm 0.5 \text{ mrad}$. We expect the rotations about the y -axis to be a similar size. These rotations could generate a significant geometric phase. However, as illustrated in figure 4.11(a), the field rotations in both axes are unaffected by an electric field reversal, so the geometric phase remains unchanged and produces no systematic error.

We can also consider the field rotations that would result from a small patch of charge, δQ , on the surface of one of the field plates. By modelling the electric field from a disc of charge, we find that the field components in the x and y axes are largest along the edges of the patch. If the patch is positioned to one side of the plates, as illustrated in figure 4.11(b), the x and y field components would cause the electric field between the HV plates to rotate anticlockwise about the z -axis (as viewed in the xy -plane). Nonetheless, the patch of charge does not cause a systematic error, because reversing the applied electric field flips the angles that the electric field vector make with the x and y axes, so the sense of rotation remains unaffected and the geometric phase remains the same.

However, if the field rotations due to the bent field plates are combined with the rotations due to a patch of charge, then the geometric phase can produce a systematic error. This is because the bent plates generates a tilt that does not reverse with the E -switch, whereas the patch of charge produces a tilt that does reverse with the E -switch. So together they generate a field rotation that changes when the electric field is reversed.

To place a limit on the potential systematic shift we consider a worst case scenario, where a 1 V patch potential is offset from the plate centre, and fills roughly half the plate length, as shown in figure 4.11(b). By modelling the field from this patch we predict it generates rotations of approximately $1\ \mu\text{rad}$, when averaged over the molecular beam. In the worst case we consider a situation where the bent HV plates generate a 1 mrad rotation about x that does not reverse with the E -switch, and the patch generates a $1\ \mu\text{rad}$ rotation about y that does reverse with the E -switch. We consider a situation where, as the molecules enter the second half of the plate length, the bent plates cause the electric field to rotate about x . The field then rotates about y as the molecules pass the plate region occupied by the patch. The rotation about x then reverses as the bend in the plate reverses, and finally the rotation about y reverses as the molecules leave the patch region. In one electric field state the rotation is clockwise about z , in the other it is anticlockwise. The geometric phase, therefore, reverses with the E -switch and hence imitates an EDM. A numerical simulation which accurately models the electric field inside the machine, and averages over the volume occupied by the molecules, limits the systematic error for this scenario to be:

$$\sigma_{\text{geo}} < 1 \times 10^{-30} \text{ e cm.} \quad (4.26)$$

4.4 Non-Zero Channels

As explained in section 3.9 we evaluate all 511 analysis channels¹¹ to check there is nothing untoward in the dataset. In an ideal experiment only the unsigned averages of $\{SIG\}$ and $\{\delta B\}$ would be non-zero, but in reality this is never the case. Table 4.1 lists the channels which are non-zero by more than 3σ . Note that if $\{X\}$ is non-zero, $\{X \cdot \delta B\}$ will also tend to be non-zero. With the exception of $\{LF1 \cdot \delta B\}$, these dotted channels are not listed as they contain no extra information. The E , B , and rf manual reversals provide eight possible ways to sign and average the channel values for each block. We check each signing for a non-zero value. If a channel is non-zero for more than one signing, the table gives the value for the most relevant signing. In particular, for the channels containing E or B , we are most interested in the respective signings with $\widehat{M}_E$ or $\widehat{M}_B$, as these account for the fact that the electric and magnetic field manual reversals change the assignment

¹¹The $\{E \cdot B\}$ channel is blinded!

between switch state and field direction. The table lists the channel ratios $\{X\}/\{\delta B\}$ and $\{X \cdot \delta B\}/\{\delta B\}$ so that they may be converted into physical units using the equations derived in appendix A. The conversions of $\{RFiF \cdot B\}/\{\delta B\}$ and $\{RFiF \cdot E\}/\{\delta B\}$ into physical units are as defined in equations 4.4 and 4.6 respectively.

TABLE 4.1: Non-zero channels for the 2011 dataset.

	Channel	Value (10^{-3})	Significance
1	$\{\delta B\}/\{SIG\}$	117 ± 1	δB step changes the signal by 12%
2	$\{LF1\}/\{\delta B\}$	60.7 ± 0.5	probe laser beam frequency detuned by -6.5 MHz from the nearest peak in the optical spectrum ($F = 0$ and other underlying spectral lines): 33% of the ~ 20 MHz optical spectrum linewidth
3	$\{LF1 \cdot \delta B\}/\{\delta B\}$	2.0 ± 0.1	probe laser beam frequency detuned by -1.2 MHz from $F = 0$: 6% of the ~ 20 MHz optical spectrum linewidth
4	$\{RF1F\}/\{\delta B\}$	-31.8 ± 0.2	rf1 frequency detuned by 1.2 kHz: 2.4% of the ~ 50 kHz linewidth
5	$\{RF2F\}/\{\delta B\}$	8.7 ± 0.2	rf2 frequency detuned by -0.3 kHz: 0.6% of the ~ 50 kHz linewidth
6	$\{RF1A\}/\{\delta B\}$	-6.8 ± 0.3	rf1 amplitude detuned by 0.9% from π -pulse requirement
7	$\{RF2A\}/\{\delta B\}$	-2.5 ± 0.2	rf2 amplitude detuned by 0.3% from π -pulse requirement
8	$\{RF1F \cdot RF2F\}/\{\delta B\}$	8.1 ± 0.1	a correlation between rf1 and rf2 frequency detuning most likely due to a slow drift in the electric field amplitude, which would shift the resonant frequency of both rf pulses
9	$\{RF2A \cdot \pi\}/\{\delta B\}$	-3.9 ± 0.2	π -flipper changes the amplitude of rf2 by 0.5%
10	$\widehat{M}_B \{B\}/\{\delta B\}$	-80 ± 2	average uncanceled stray magnetic field of 136 pT aligned against the z -axis
Continued on the next page...			

Table 4.1 – continued from previous page

	Channel	Value (10^{-3})	Significance
11	$\widehat{M}_E \{E\}/\{SIG\}$	0.52 ± 0.04	electric field reversal changes signal by 0.05% possibly due to an asymmetric electric field reversal, which would change the efficiency of the rf transitions
12	$\widehat{M}_B \{RF1F \cdot B\}/\{\delta B\}$	4.69 ± 0.09	phase due to rf1 frequency detuning: 288 nrad/Hz
13	$\widehat{M}_B \{RF2F \cdot B\}/\{\delta B\}$	-1.47 ± 0.09	phase due to rf2 frequency detuning: -91 nrad/Hz
14	$\widehat{M}_E \{RF1F \cdot E\}/\{\delta B\}$	-0.84 ± 0.09	$ E_0 $ changes by 0.21 V/cm upon reversal, in rf1 pulse region
15	$\widehat{M}_E \{RF2F \cdot E\}/\{\delta B\}$	0.87 ± 0.09	$ E_0 $ changes by -0.21 V/cm upon reversal, in rf2 pulse region
16	$\widehat{M}_B \{RF1A \cdot B\}/\{\delta B\}$	-0.7 ± 0.1	phase correlated with rf1 amplitude
17	$\widehat{M}_B \{RF2A \cdot B\}/\{\delta B\}$	0.7 ± 0.1	phase correlated with rf2 amplitude
18	$\widehat{M}_B \{LF1 \cdot B\}/\{\delta B\}$	-1.3 ± 0.1	phase correlated with laser frequency

Rows 3–7 show that the laser and rf parameters were detuned by a small amount. There are two reasons for this: 1) The analysis gate used here is not the same as that which was used to lock the experimental parameters during data acquisition. 2) We retain all acquired data blocks, including the ones at the beginning of each cluster where parameters are still locking in (see figure 3.4 for example). The same technique was used at the beginning of each day to manually set each parameter before acquisition was started, so it is likely that the initial detunings of each parameter were always biased in the same direction. Fortunately the parameter detunings are small when averaged over the whole dataset.

Note that for the probe laser frequency detuning we consider $\{LF1 \cdot \delta B\}$ not $\{LF1\}$ (which is slightly larger). This is because $\{LF1 \cdot \delta B\}$ measures the frequency detuning of the probe laser from the $F = 0$ state. $\{LF1\}$ measures only the detuning from the nearest peak in the optical spectrum. This is evidently skewed from the $F = 0$ peak

as a result of underlying spectral lines. These channels, evaluated using the normalised probe PMT signal, only measure the detuning of the probe laser beam. This is because normalising the probe PMT signal with the pump PMT signal almost perfectly removes all dependence on the pump laser detuning.

In section 4.3.2 it was shown that when $\bar{V} \neq 0$ V, $d_e \propto \{RF1F\}$. We find no other such correlations (except those described in the previous sections). The systematic error contribution due to a non-zero $\{RF1F\}$ is small in comparison to the statistical uncertainty of the dataset.

Row 8 shows a positive correlation between the frequency detuning of the first and second rf pulses (rf1 and rf2). The non-zero value is most likely the result of a slow drift in the supply voltages, as this would shift the resonant frequency equally for both rf pulses. This is not a major concern, as the polarisation factor (equation 1.13) is insensitive to small changes in electric field around 10 kV/cm.

Row 9 shows that the π -flipper modulated the amplitude of the second rf pulse in addition to modulating the relative phase difference between the two pulses. This is not a large effect and should not have introduced any Ramsey type effects.

Row 10 gives the average uncanceled stray magnetic field along $\hat{z}$. In section 4.3.3 it was shown that this would not cause a significant systematic shift.

Row 11 shows that the electric field reversal caused the molecular signal to change slightly. This could result from an imperfect electric field reversal, which would change the efficiency of the rf transitions. When combined with an uncanceled stray magnetic field (row 10), the channel combination $\widehat{M}_E \widehat{M}_B \{B\} \{E \cdot \delta B\} / \{\delta B\}^2$ contributes a non-zero ‘lineshape correction’ to the crude EDM channel $\widehat{M}_E \widehat{M}_B \{E \cdot B\}$ as shown in equations 3.13 and 3.14. Overall the lineshape correction is four times smaller than the statistical uncertainty in d_e .

Rows 12–15 are discussed thoroughly in section 4.3.1.

The final three rows (16–18) measure interferometer phases correlated to the rf pulse amplitude modulations and the laser frequency modulation. We have no physical model for these channels and, like the rf frequency phase correlation, any one of these channels could generate a fake EDM signal if the corresponding E -switch correlated channel $\left(\widehat{M}_E \{X \cdot E\}\right)$ is non-zero. Fortunately none of the E -switch correlated channels (other than rows 14 and 15) are non-zero so we have not investigated these phases any further.

Note that the channels $\{\pi\}$ and $\{\pi \cdot E\}$ do not feature in the table. A non-zero $\{\pi\}$ channel would indicate that the Ramsey component is non-zero and the π -flipper is actively cancelling this out over the course of each block. The rf parameters are sufficiently

locked in to ensure that the Ramsey component is very small. Furthermore, the phase scrambler suppresses any residual effect by $1/\sqrt{N}$, where N is the number of blocks in the dataset. It is not surprising, therefore, that when averaged over the whole dataset, these channels are consistent with zero. In his thesis [74], Kara conservatively estimated that, when averaged over the whole dataset, the size of the fake EDM due to the Ramsey component¹² would be no larger than 9×10^{-31} e cm, which is small enough to neglect.

In addition to looking for non-zero channels in the normalised probe PMT, we also evaluate all 511 channels, with each of the eight possible manual state signings, using the signals measured by: the pump PMT, the three mini-fluxgate magnetometers, the *Bartington* magnetometer, the two leakage current monitors and the two dummy inputs. In general all these detectors give the expected results, but there are a few exceptions where certain channel values are non-zero.

The pump PMT does return a significantly non-zero value in $\{LF1\}$. This is because we lock the laser frequency to ensure the probe laser beam is on resonance with the $Q(0)(F = 0)$ transition, and there is inevitably a small Doppler shift of a few MHz between the pump and probe laser beams due to slight beam misalignments.

The mini-fluxgate magnetometer positioned right next to HV supplies measures a non-zero $\{E\}$, which indicates that the HV relays generate a magnetic field that changes when they switch state. However, the 5m separation between the supplies and the beam machine sufficiently attenuate this field so that no other magnetometer measures a significant value for this channel. We are therefore confident that this field has no significant effect on the molecules.

The *Bartington* magnetometer does detect both applied magnetic field steps, which is no surprise. However, this detector does also detect the non-zero channels: $\{\pi\}$ and $\{RFiA\}$, which shows that the rf system generates different magnetic fields for each of the rf amplitude and π -flipper states. These fields are very small and they do not correlate with the E -switch so are of no concern.

The leakage monitors register E -switch correlated currents of less than 1 nA, which is tolerable (see section 4.3.3). Weirdly these detectors do return a non-zero $\widehat{M}_B\{B\}$, which cannot easily be explained. Nonetheless these fields are too small to be of any concern.

¹²I have scaled his quoted value according the number of measurements in the 2011 dataset.

4.5 Other Correlations

Further to evaluating each of the non-zero analysis channels, we investigate whether d_e depends upon the pump and probe laser polarisation angles, the machine manual state, or the analysis gate position.

4.5.1 Pump/Probe Laser Polarisation Dependence

Recall that at the beginning of each block the laser polarisations are independently rotated to randomly chosen angles about the x axis (see section 2.3.3). We can bin the EDM data by polariser angle and analyse each bin individually, albeit at lower experimental sensitivity.




FIGURE 4.12: The dependence of the measured EDM value upon the angle of the pump/probe laser linear polarisation. 0° was arbitrarily chosen and does not lie along any particular lab axis.

As shown in figure 4.12, we find no clear dependence on the polarisation angle of either laser beam, and all fluctuations appear to be purely statistical. A χ^2 test that assumes a constant expected value returns a p-value¹³ of 73% for the pump polariser binned data and 13% for the probe polariser data. The relatively small p-value for the probe data is entirely due to the two outliers at 200° and 320° . However, the distribution of d_e values (figure 3.6) has long wings, so occasional outliers are to be expected.

The most likely alternative model for either set of binned data would be a sinusoidal variation with a period of either 180° or 360° . If we fit the amplitude and phase of such models, we measure no significant variation for either set of binned data.

¹³This is defined as the probability of obtaining a χ^2 value at least as small as the one measured for the assumed model, with the given number of degrees of freedom.

As the binned data is consistent with a constant EDM value, and because we find no clear evidence for a sinusoidal variation, we assert that the EDM value does not depend upon the polarisation angle of either the pump or probe laser beam.

4.5.2 Manual State Dependence

As shown in figure 4.13, if the EDM data is binned and analysed by manual state, we see no statistically significant dependence on the manual state of the machine. This is shown more quantitatively in table 4.2, where each row quotes the difference in the measured EDM, Δd_e , between states of opposite manual configuration; there are seven pairs in total. For example the row labelled M_E compares the weighted average of the second four points in figure 4.13, where $M_E = \text{true}$, to the weighted average of the first four points, where $M_E = \text{false}$. As each difference in table 4.2 is consistent with zero we can be confident that the E , B and rf field producing hardware were working properly during acquisition; the measured EDM value did not change when any electrical connections were reversed.




FIGURE 4.13: Corrected (blinded) EDM value filtered by manual state (M_E, M_B, M_{RF}). Each point contains only ~ 750 data blocks, so the error bars are roughly three times larger than the statistical uncertainty on the final result. The solid and dashed horizontal lines represent the final (blinded) EDM measurement and its statistical uncertainty respectively.

4.5.3 Molecular Arrival Time Dependence

During acquisition of the main dataset $\bar{V}$ was carefully zeroed to minimise the voltage offset systematic effect (section 4.3.2). As was shown in figure 4.7, when $\bar{V} = 0.0 V$ we observe no temporal variation of d_e across the normed PMT time of flight profile. However, the analysis gates used to generate that figure are quite narrow, so the molecular signal in each gate is quite low. This makes the uncertainty for each point relatively high

TABLE 4.2: Differences in measured EDM values between states of opposite manual configuration.

Manual State Signing	Δd_e (10^{-28} e cm)
M_E	-6 ± 12
M_B	-11 ± 12
M_{RF}	3 ± 12
$M_E M_B$	1 ± 12
$M_E M_{RF}$	8 ± 12
$M_B M_{RF}$	-3 ± 12
$M_E M_B M_{RF}$	17 ± 12

compared to the uncertainty that can be achieved by integrating the whole TOF profile (the statistical uncertainty on d_e will be discussed in more detail in the next section). In order to verify the absence of a temporal dependence with higher precision we can split the main analysis gate into two halves and compare measurements made with the relatively fast molecules to those made with the slow.

A full analysis of the main dataset using these two gates returns nothing unexpected — all channel values are consistent with those quoted in the previous sections and we find no manual state dependence. For the fast gate, defined over the period $(2132 - 2198) \mu\text{s}$ (see figure 2.11), the blinded EDM value is $(86.6 \pm 7.0) \times 10^{-28}$ e cm. This agrees very well with the EDM value of $(91.4 \pm 6.7) \times 10^{-28}$ e cm, which was measured using the slow gate, defined over the period $(2198 - 2264) \mu\text{s}$. It must be pointed out that the two error bars are correlated by almost all sources of noise except shot noise, so saying they agree within error is somewhat dubious. But it is at least clear that the choice of analysis gate does not affect the final result significantly. The difference is within the statistical uncertainty of the final result, discussed in the next section.

4.6 EDM Sensitivity Variation with PMT Gates

In section 4.1 it was stated that the PMT analysis gates were chosen to minimise the statistical uncertainty in d_e . The joined blue points in figure 4.14 show the variation in EDM sensitivity with probe PMT analysis gate centre time and width¹⁴. The sensitivities

¹⁴Recall that the pump PMT gates are defined by dividing the probe PMT gates by the ratio of the pump and probe laser distances from the molecular source ($=3.842$). This ensures both PMTs sample the same molecules.

shown were calculated using the 6194 blocks that were taken under normal running conditions. As can be seen, EDM sensitivity is optimised with a $130\ \mu\text{s}$ wide probe PMT gate centred $2198\ \mu\text{s}$ after the initial YAG laser Q-switch trigger signal. With this gate, the uncertainty on d_e , at the 68.3% confidence level, is $5.7 \times 10^{-28}\ \text{e cm}$.




FIGURE 4.14: Demonstration of how the EDM sensitivity varies with analysis gates. The joined blue points represent the statistical uncertainty at the 68.3% confidence level, as measured by bootstrapping the 5% trimmed mean (see section 3.8). The dashed red line indicates the shot noise limit as predicted by equations 4.27 and 4.28. In (b) the analysis gate is centred $2198\ \mu\text{s}$ after the initial YAG laser Q-switch. We choose to analyse the dataset with a $130\ \mu\text{s}$ wide probe PMT gate centred $2198\ \mu\text{s}$ after the initial YAG laser Q-switch. The pump gate is defined by dividing the probe PMT gate times by 3.842. These gates minimise the statistical uncertainty in d_e .

It is interesting to compare the minimum statistical uncertainty with the shot noise limit. We can estimate this from the PMT signals that we observe under normal running conditions. It can be shown (see appendix B) that

$$\sigma_{\text{shot}} = \frac{\mu_B B_{\delta B}}{E_{\text{eff}}} \sqrt{\frac{\kappa_{\text{probe}} \{SIG\}_{\text{probe}}}{4096 N \{\delta B\}_{\text{probe}}^2}} = \frac{\phi_{\delta B} \hbar}{T E_{\text{eff}}} \sqrt{\frac{\kappa_{\text{probe}} \{SIG\}_{\text{probe}}}{4096 N \{\delta B\}_{\text{probe}}^2}} = \frac{3.9 \times 10^{-26}}{\sqrt{N}} \text{ e cm}, \quad (4.27)$$

where $\kappa_{\text{probe}} = 0.1\ \text{V/MHz}$ is the probe PMT calibration factor, $B_{\delta B} = 1.7\ \text{nT}$, $E_{\text{eff}} = 14.5\ \text{GV/cm}$, $\{SIG\}_{\text{probe}} = 49.6\ \text{V}\mu\text{s}$ and $\{\delta B\}_{\text{probe}} = 6.09\ \text{V}\mu\text{s}$. N is the number of blocks measured. Setting $N = 6194$, in equation 4.27 gives an uncertainty of $4.9 \times 10^{-28}\ \text{e cm}$. However, normalisation of the probe PMT signal with the pump PMT signal increases this value by the ratio

$$\sqrt{\frac{\kappa_{\text{probe}}}{\{SIG\}_{\text{probe}}} + \frac{\kappa_{\text{pump}}}{\{SIG\}_{\text{pump}}}} / \sqrt{\frac{\kappa_{\text{probe}}}{\{SIG\}_{\text{probe}}}} = 1.12, \quad (4.28)$$

where $\kappa_{\text{pump}} = 0.02\ \text{V/MHz}$ is the pump PMT calibration factor and $\{SIG\}_{\text{pump}} = 39.7\ \text{V}\mu\text{s}$. Multiplication of equations 4.27 and 4.28 gives a shot noise limited uncertainty of $5.5 \times 10^{-28}\ \text{e cm}$, which is 4% smaller than the optimum measured uncertainty. The difference is likely due to ambient magnetic field noise (section 4.2.1) and slightly detuned experimental parameters (see below). The long tails

in the distribution of d_e (section 3.8) is also likely to add to the discrepancy. With this in mind a 4% difference does not cause us major concern.

Consider figure 4.14(a) again. The dashed red line shows the shot noise limit as calculated using the average values of $\{SIG\}$ and $\{\delta B\}$ for a $130\mu s$ wide gate centred at various times over the probe PMT TOF profile. The minimum shot noise is achieved with a gate centred at $2213\mu s$. This gate maximises the molecular signal. Interestingly the minimum in the shot noise limit does not coincide with the centre time of $2198\mu s$, which we find minimises the overall statistical uncertainty in d_e . We have no conclusive explanation for this, but it is very likely that the bias towards early times is linked to the fact that, during data acquisition, the active parameter locks used an analysis gate centred at $2156\mu s$. This was a mistake. The locks ensure that parameter detuning averages to zero over the lock analysis gate. Since the lock gate was biased towards the faster molecules (that arrived at earlier times), we find that some parameters, such as the rf amplitude and frequency, are detuned for the slower molecules (that arrived at later times). It is possible that parameter detuning could add noise to the EDM signal through some non-linear mechanism, although we have not investigated this at all in any detail.

If we had acquired EDM data with the lock analysis gate centred at $2213\mu s$ rather than $2156\mu s$, it is likely that the measured uncertainty would have coincided with the minimum in the shot noise limit. However, the associated reduction in uncertainty would be relatively small ($\sim 0.1 \times 10^{-28}$ e cm). Regardless, the statistical uncertainty that we do measure, 5.7×10^{-28} e cm, is more than sufficient to set a now world leading limit on the electron EDM.

4.7 EDM Value

With the blind offset removed the final result, at the 68.3% confidence level, is

$$d_e = (-2.4 \pm 5.7_{\text{stat}} \pm 1.7_{\text{syst}}) \times 10^{-28} \text{ e cm}. \quad (4.29)$$

The two error bars correspond to the statistical and systematic uncertainties respectively. The systematic error is calculated by adding in quadrature the uncertainties summarised in table 4.3.

We can set an upper limit on our measurement by bootstrapping the measured values of $|d_e|$ instead of d_e . The statistical confidence bounds, $\sigma_{\text{stat}, p\%}$, are then calculated by integrating the one-sided probability distribution from zero, and extracting the values of $|d_e|$ where the integral contains $p\%$ of the distribution. Systematic uncertainty bounds $\sigma_{\text{syst}, p\%}$ are calculated by assuming a Gaussian distribution with a mean of zero and a

TABLE 4.3: Summary of the uncertainties which limit our measurement of d_e .

Source	Uncertainty (10^{-28} e cm)
δE effects	1.1
$\bar{V}$ effects	$\sqrt{1.3^2 + 0.1^2}$
leakage currents	0.2
shield magnetisation	0.25
geometric phase	0.01
total systematic uncertainty	1.7
total statistical uncertainty (68.3% CL)	5.7

standard deviation of 1.7×10^{-28} e cm. The total bound on $|d_e|$ is then calculated by combining these two uncertainties according to $\sqrt{\sigma_{\text{stat, p}\%}^2 + \sigma_{\text{syst, p}\%}^2}$. Table 4.4 lists the upper limit on $|d_e|$ at various confidence levels.

TABLE 4.4: Bootstrap calculated bounds on $|d_e|$.

Confidence Interval (%)	Bound on $ d_e $ (10^{-28} e cm)
68.3	6.5
90.0	10.6
95.8	13.1
99.5	18.4

Comparing our 90% confidence limit of 10.6×10^{-28} e cm to the comparable previous best limit of 16×10^{-28} e cm, made using beams of thallium atoms in 2002 [58], it can be seen that our value reduces the upper limit by a factor of 1.5. Admittedly this is only a modest improvement which does not significantly affect the EDM limitations on new physics beyond the standard model, but our experiment is subject to different sources of systematic error to the thallium beam electron EDM experiment, so the two results can be seen as complementary. It is pleasing, therefore, that the two measurements agree within error.

It is worth emphasising that our new measurement is significant, as it is the first time that a new improved limit on the electron EDM has been set using molecules rather than atoms. Comparing the relative size of the statistical uncertainty to each systematic uncertainty value quoted in table 4.3, it is clear that there is still potential for a more sensitive EDM measurement. The large enhancement factor, and insensitivity to perpendicular magnetic fields make YbF molecules great candidates for marked improvement on the electron EDM limit. This measurement is the first step down that research avenue.

Note that the physical mechanism of the enhancement factor is rather complex and is subject to theoretical calculation. There are a number of calculations of $E_{\text{eff}}^{\text{max}}$ for YbF [41–46], and most values agree to within 10%. However, our result can be interpreted as a measurement of the P,T-violating EDM of the YbF molecule at 10 kV/cm. In order to calculate appropriate values for this, multiply the values in table 4.4 by $E_{\text{eff}}/E = -1.45 \times 10^6$.

Chapter 5

Upgrading the YbF Interferometer

After publishing our EDM measurement we decided that we should upgrade the experimental setup to eliminate certain systematic effects and improve the statistical uncertainty.

The most significant upgrade to the apparatus was to improve the uniformity of the rf polarisation over the length of the HV plates. We did this to enable their full length to be used for EDM measurement, which should improve EDM sensitivity by approximately 75%. We also took the opportunity to upgrade the rf amplifier so that shorter rf pulses could be applied to reduce the dependence of the molecular signal on rf detuning, and hopefully reduce the associated systematic effects. Other upgrades included adding a third layer of magnetic shielding to reduce the magnetic field noise, as well as increasing the applied electric field to increase the EDM sensitivity. A full account of all the upgrades is given in section [5.1](#).

After making changes to the experimental apparatus we ran a series tests to measure the gain in sensitivity and check that none of the changes introduced any new sources of systematic error. Section [5.2](#) discusses each test in turn. A summary of the systematic uncertainties and a measurement of the EDM sensitivity is given in section [5.3](#).

5.1 Upgraded Apparatus Setup

5.1.1 RF Transmission Line

The model presented in section 2.2.2 assumes that the magnetic component of the rf field is linearly polarised along the x -axis. As the electric field is also assumed to be along z , the model predicts that equal populations are driven from $|0,0\rangle$ into the $|1,\pm 1\rangle$ states. However, if the rf is not perfectly linearly polarised, one of the states $|1,1\rangle$ or $|1,-1\rangle$ will couple more strongly to the rf. This is because the transitions between $|0,0\rangle$ and $|1,\pm 1\rangle$ are driven by circularly polarised rf with opposite handedness. When the rf is linearly polarised along x the amplitudes of the two circular components equal, but if the rf polarisation is not perfectly linear the amplitudes become unequal.

An unequal coupling to the rf can create a population difference between the $|1,\pm 1\rangle$ states. In addition it can affect the phase difference that is accumulated during each rf pulse, if the $|1,\pm 1\rangle$ states are split by a magnetic field and the rf frequency is detuned from the mean separation between $|0,0\rangle$ and $|1,\pm 1\rangle$. A phase difference between the $|1,\pm 1\rangle$ states can also be generated if the rf polarisation ellipse rotates about the axis of the applied electric field during each rf pulse [55, 96].

A difference in population between the $|1,\pm 1\rangle$ states should only reduce fringe contrast and hence limit the experimental sensitivity. But changes in fringe contrast that correlate with the E -switch inflate the ‘lineshape correction’ $\left(\widehat{M}_E \widehat{M}_B \{B\} \{E \cdot \delta B\} / \{\delta B\}^2\right)$, which is undesirable.

The phase difference between the $|1,\pm 1\rangle$ states that is acquired during each rf pulse is a major source of concern, as it might somehow depend upon the E -switch and hence mimic an EDM. One possible mechanism is through a non-zero voltage offset. Recall from section 4.3.2 that this effect causes the electric field distribution to change when the direction of the electric field is reversed. As the molecules are only sensitive to rf in the plane perpendicular to the local electric field, a change in field distribution could change the phase acquired during each pulse.

We attempt to minimise the effects of polarisation rotations and elliptically polarised rf by driving transitions between $|0,0\rangle$ and $|1,\pm 1\rangle$ in plate regions where the rf polarisation is known to be highly uniform. As Kara demonstrated [74] (see also below), with the original transmission line setup, the rf becomes slightly elliptically polarised towards either end of the parallel plates. So for the 2011 measurement only the central 40 cm region of the 70 cm long plates was used for free evolution of the molecular state in static electric and magnetic fields.

In the shot noise limit, the EDM uncertainty is inversely proportional to the phase evolution period. This is easily understood, as the EDM phase is linearly proportional to the free evolution period. A larger phase produces a bigger change in signal between electric field switch states, and hence a larger signal to noise ratio. It was decided that we should attempt to increase the polarisation uniformity so that the whole length of the HV plates could be used, potentially increasing experimental sensitivity by up to 75%.




FIGURE 5.1: Original scheme for coupling rf onto the parallel plate structure. Only one end of the plates is shown — rf is coupled out at the far end with an equivalent setup. The trimming capacitors help with impedance matching between the coaxial and parallel plate transmission lines. The isolating capacitors protect the rf electronics from the DC high voltage required for the EDM interaction. Note the asymmetric setup and the exposed capacitor legs which may act as rf antennae.

In the original apparatus setup the rf is coupled onto the parallel plate structure through DC isolating capacitors, as illustrated in figures 2.12 and 5.1. The coupling is off centre and there are many exposed capacitor legs which might act as antennae — radiating rf into the parallel plate region. Both artefacts could distort the rf polarisation, so we decided to improve the rf coupling to minimise polarisation rotations and any ellipticity generated at each end of the parallel plate transmission line.

Before describing the upgrades to the rf system, it is worth describing the technique used to map the rf polarisation along the beam line. In general a field map is generated from many single rf pulse parameter scans, where the rf pulse time is incremented in between each scan so that molecules interact with the rf field at different positions within the machine. The parameter scans are then converted to field measurements by fitting each to an appropriate physical model, as described in section 2.3.5 and [84].

In order to map the rf polarisation, the rf pulse frequency is continuously scanned over the ~ 170 MHz resonance. Relatively large electric and magnetic fields of 2 kV/cm and $2\mu\text{T}$ are applied along z to polarise the molecules and split the $F = 1$ manifold so that the $|0,0\rangle \rightarrow |1,\pm 1\rangle$ transitions are resolved. As each transition is driven by a different rf polarisation¹, a comparison of the rf spectral peak amplitudes provides a measure of

¹ $\Delta m_F = 0$ from linearly polarised rf along z , $\Delta m_F = \pm 1$ from left and right handed circularly polarised rf about z .

the relative amplitude of each polarisation. We set the rf pulse power to the $\pi/2$ -pulse requirement so that the difference in peak height is linear in transition amplitude. The rf pulse length is set to be $50\ \mu\text{s}$, which is small enough to ensure the molecules do not move far during each pulse, yet large enough to easily resolve each spectral peak.

Figure 5.2 shows an example of a frequency scan over the $|0,0\rangle \rightarrow |1,\pm 1\rangle$ resonances, which clearly demonstrates differing amounts of left and right handed circular polarisations. The amplitudes are determined by fitting each peak to equation 2.16a, with the addition of an overall amplitude scaling and background offset. The solid red line in figure 5.2 demonstrates the best fit in this instance.




FIGURE 5.2: The rf lineshape observed by scanning the rf frequency over the 170 MHz resonance. The red line represents the best fit to a double lineshape function of the form given in equation 2.16a. The central frequency provides a measure of E_0 , the splitting between the peaks measures B_0 , and the difference in heights measures the ellipticity of the rf pulse.

Unfortunately the field mapping technique described above can only measure changes in ellipticity — it cannot measure general rotations of the polarisation ellipse about the z -axis. This is because $|0,0\rangle$ couples equally to both x and y polarised rf. However, a polarisation rotation and change in ellipticity will most likely result from the same apparatus imperfection, since the parallel plate transmission line should generate a perfect TEM mode, with the oscillating magnetic field polarisation along x . We can therefore argue that a uniform linear polarisation along the length of the beam line is a good indication that there are no large polarisation rotations between the parallel plates.

To gain a better insight into the nature and homogeneity of the rf polarisation along the plate length the field maps were supplemented with loop antenna probe measurements, which can detect rotations of the polarisation ellipse, albeit rather crudely, being subject to errors such as capacitive pick-up and the electrical properties of the local environment². These probe measurements were performed out of vacuum, on an optical bench.

²Namely where I stood and how the plates were positioned in lab.




FIGURE 5.3: A map of the rf polarisation as measured by the molecules. Circular blue points represent measurements that were taken before the transmission line upgrade. The square red points are for the coupling scheme illustrated in figure 5.4(a), the diamond green points are for the final transmission line design, as illustrated in figure 5.4(b).

Figure 5.3 shows the molecular field mapping measurements made on the original rf transmission line and the subsequent upgraded setups that are detailed below. The ellipticity, e , is a measure of the relative amplitude of the left and right handed circular components of the rf about the z -axis. It is defined according to

$$e = \frac{a_+ - a_-}{a_- + a_+}, \quad (5.1)$$

where $a_{\pm}$ are the measured heights of the two peaks shown in figure 5.2, that correspond respectively to the transitions between $|0,0\rangle$ and $|1,\pm 1\rangle$. $e = -1$ corresponds to full left hand circularly polarised rf about the z axis, $e = +1$ to right hand circular. $e = 0$ represents an equal superposition of left and right handed polarisations, i.e. linear polarisation.

Consider the molecular field mapping measurements made on the original rf transmission line setup. As shown by the circular blue points in figure 5.3, the polarisation is linear along x over the central region, however, some ellipticity is generated towards either end of the plates. This is not surprising considering the highly asymmetric setup illustrated in figure 5.1. As stated, the coupling point is off centre and there are many exposed capacitor legs which might radiate rf between the parallel plates. Probe measurements on a replica set of parallel plates have confirmed that this setup does indeed generate some ellipticity at each end, but the predominant effect is to rotate the polarisation ellipse about z towards the coupling point. Both effects are undesirable and result from a non-ideal rf field component in the y -axis, $B_y^{(\text{rf})}$.

To minimise the $B_y^{(\text{rf})}$ component the setup illustrated in figure 5.4(a) was developed on the spare set of parallel plates. This setup is designed to be highly symmetric, with all





FIGURE 5.4: The two designs for the upgraded rf coupling schemes. Only one end of the plates is shown — rf is coupled out at the far end with an equivalent setup. The trimming capacitors help with impedance matching between the coaxial and parallel plate transmission lines. The isolating capacitors protect the rf electronics from the DC high voltage required for the EDM interaction.

exposed wire kept to a minimum, or hidden above and below the plates to prevent them from radiating between the plates. Currents flowing around the conducting loop and out from the coupling point act to cancel $B_y^{(\text{rf})}$ in the region of the molecular beam (along the centre line of the plates).

Figure 5.5 shows a comparison of the loop antenna measurements along the beam line for the original and new coupling schemes. The ‘relative normalised amplitude’ is the amplitude of the x , y and z components of the rf normalised such that when added in quadrature, their sum is always one. Comparing figures 5.5(a) and 5.5(b), it is apparent that the new coupling scheme notably improves the rf polarisation. $B_y^{(\text{rf})}$ is consistent with zero right up to the coupling point. Unfortunately negation of $B_y^{(\text{rf})}$ is incomplete away from the centre line, however this component is antisymmetric about the coupling point so should cancel when averaged over the full width of the molecular beam. Furthermore, probe measurements suggest that, near the coupling point where the polarisation is worst, for small departures from the centre line (< 1 cm — the maximum width of the molecular beam), the amplitude ratio $B_y^{(\text{rf})}/B_x^{(\text{rf})}$ is ~ 20 times smaller with upgraded setup than the original.

Despite the encouraging predictions of the loop antenna measurements, when the coupling scheme illustrated in figure 5.4(a) was installed within the vacuum chamber, we found that a significant $B_y^{(\text{rf})}$ persisted. As shown by the red square points in figure 5.3, the molecular field map exhibits significant deviations from linear near the coupling points at both ends, but particularly over the first 0.2 m of plate length. The results of this field map are inconsistent with the probe measurements shown in figure 5.5(b). We don’t fully



(a) Original coupling scheme as illustrated in figure 5.1. (b) Upgraded coupling scheme as illustrated in figure 5.4(a).

FIGURE 5.5: Normalised (x, y, z) rf field amplitudes mapped down the centre line of the plates, over the first 20 cm of the transmission line using the loop antenna probe. Note the marked improvement of the upgraded coupling scheme shown in (b): the field is almost entirely polarised along x .

understand this, but it should be noted that the probe and molecular measurements were made on two different setups. The probe measurements were made on the spare plates during development of the new scheme, whereas the molecular field map was made after the new scheme had been replicated on the plates used for our last EDM measurement. It is likely that a difference between the two builds caused the conflict between the two polarisation measurements.

One possible explanation for the significant ellipticity observed near the coupling points is that the conducting loops did not radiate symmetrically. Another possibility is that the coupling points were displaced slightly from the beam line. Regardless of the cause for the non-zero ellipticity, the upgraded setup presents no advantage over the original. We would have to set the rf pulses to be at 0.2 m and 0.6 m, giving a 0.4 m phase evolution length which is exactly the same as it was before.

To allow the whole length of the HV plates to be used for phase evolutions a third setup was developed wherein the rf is propagated along a second set of parallel plates that straddle the HV plates and extend beyond them by 5 cm at either end. This setup is illustrated in figure 5.4(b). The new rf transmission line plates are 5.6 cm wide, 85 cm long sheets of 2 mm aluminium that are separated by 5.2 cm. As shown in the photos and schematics presented in appendix C, these plates are held in place with PEEK feet which attach to new support extensions that screw into the existing HV plate support structure. Teflon clamps in the centre prevent the plates from bowing.

Like the previous setup, this upgrade continues to couple rf to the centre of the plates to avoid large polarisation rotations and cancel $B_y^{(\text{rf})}$ along the molecular beam. However, the conducting loops, used to pass rf onto the lower plate of the transmission line, have now been replaced with semi-rigid coaxial cables which should contain the rf within their

outer shield and not radiate into the HV plates. The coaxial loops are quite long, so we routed them around the grounded HV plate support structure to further separate/isolate them from phase evolution region³. The low impedance conductor ensures an equal voltage reference on the coaxial shield at either end of the loops. This is important for impedance matching, as described below. The 5 cm extensions at either end of the HV plates provide a final measure to help improve rf polarisation within the HV plate region. This extra length should allow transient rf field components to die away before the electric field plate region.




FIGURE 5.6: A schematic diagram of the new rf transmission line.

A complication of increasing the rf plate spacing is that it increases the transmission line characteristic impedance, Z_0 . For a parallel plate transmission line, $Z_0 \propto s/w$, where s is the plate spacing and w is the plate width [97]. For all practically feasible combinations of plate width and spacing, the predicted impedance is roughly four times that of the $50\ \Omega$ coaxial rf feeds. This poses a problem as the power transmitted from one transmission line to another is only maximised when their impedances are matched [98]. Large impedance jumps at both coupling points would cause reflections and hence a large standing wave within the plate region, which we may not be able to compensate for with the rf attenuators (see section 2.3.5).

To maximise the transmitted rf power, we smoothly transition the impedance from the $50\ \Omega$ semi-rigid coaxial up to $200\ \Omega$ with a 4:1 balun, made from a $\lambda/2 = 0.606\text{ m}$ length of coaxial cable, as designed to work with 172 MHz rf [99]. As shown more clearly in figure 5.6, the balun adds a 180° phase delay to the signal going to the second plate. This arrangement couples the unbalanced coaxial transmission line to the balanced parallel plate line⁴.

The 5.2 cm plate spacing as quoted above was empirically chosen to match the line impedance to the required $200\ \Omega$. We varied the plate spacing with a custom jig until the transmitted rf power was maximised, and the standing wave (as sampled with the loop antenna) was minimised. The 5 pF trimming capacitors soldered at either end

³Unlike in the illustration in figure 5.4(b), where the cables pass next to the HV plates.

⁴In a balanced transmission line, each of the conductors has equal impedance to ground.

attempt to make up for any residual impedance mismatch that occurred when the new transmission line was added to the HV plate support structure used within the machine. Unfortunately impedance matching is now frequency dependent, however we have found that rf transmission does not change significantly over a 10 MHz band about the 172 MHz operating frequency.

As shown by the green diamond points in figure 5.3, the third rf setup successfully ensures a linear rf polarisation along x over the whole length of the HV plates. After careful consideration of the plate geometry we decided to set the rf pulse centre times to be $820 \mu\text{s}$ and $1800 \mu\text{s}$ after the initial valve pulse. These timings increase the phase evolution period from $642 \mu\text{s}$ to $976 \mu\text{s}$, which should reduce the uncertainty in d_e by 34%.

The pulse timings were chosen to ensure that all molecules within the gated region of the TOF profile are at least 1 cm from the HV plate edges during the rf transitions. With these settings approximately 60 cm of the 70 cm plate length is utilised for phase evolution. Unfortunately the timing of second rf pulse is constrained by a small amount of rf pick-up on the probe PMT amplifier. If the second rf pulse is left too late, the trailing edge of the pick-up coincides with the molecular signal on the probe PMT. Setting the centre time of the second rf pulse to $1800 \mu\text{s}$ conservatively prevents this from occurring.




FIGURE 5.7: A typical TOF profile observed on the probe PMT. The blue line demonstrates the molecular signal with no rf2 pulse. The red line shows the signal with an rf2 pulse but no molecular beam. The negative (red) dip shows the unwanted rf pick-up by the PMT amplifier, this dies away before the gated interval illustrated by the shaded blue region.

Figure 5.7 demonstrates the pick-up on the probe PMT. Note that the pick-up appears to die away long before the gated region. To be sure that the pick-up does not generate a systematic error we can gate the interval containing the rf pick-up and analyse this for an EDM signal. When we do this we find no evidence that the rf pick-up generates a systematic shift.

By increasing the pulse separation we decrease the pitch of the $\cos^2(\phi)$ interference fringes (equation 2.19b). We therefore need to reduce B_0 and $B_{\delta B}$ to maintain the respective

phase steps of $2\pi/4$ and $2\pi/32$. We scale these fields to be $B_0 = \frac{642}{976} \times 13.6 \text{ nT} = 8.9 \text{ nT}$ and $B_{\delta B} = \frac{642}{976} \times 1.7 \text{ nT} = 1.1 \text{ nT}$ by adjusting the magnetic field stepping currents to $I_{B_0} = 0.530 \text{ mA}$ and $I_{\delta B} = 0.065 \text{ mA}$.

5.1.2 High Voltage Plate Flattening

During the installation of the first rf system upgrade, we took the opportunity to attempt to reduce the electric field variation along the HV plate length. This would be beneficial to reduce the detuning across the molecular pulse and reduce certain systematics such as the geometric phase.




FIGURE 5.8: A map of the variation in HV plate separation. This was calculated by mapping the applied electric field and assuming uniform plate voltages. The blue points are before the apparatus upgrade. The red points are for softer springs on the HV feedthroughs. The green points are for softer springs and a reduced number of ceramic dowels supporting the plates. Evidently neither modification had a significant effect on the bending of the plates.

The field variation was first discovered in 2007 [73, 74, 84], and was thought to be due to a bend in the plates caused by the spring loaded HV feedthroughs pushing on the plates in between the two pairs of ceramic dowels that hold each field plate in position at either end (see figure 2.15(a)). Assuming a uniform plate voltage, a variation in plate spacing of up to $100 \mu\text{s}$ was observed. This is demonstrated by the blue points in figure 5.8. The dowel pairs are located at roughly 0.1 m and 0.6 m, and the feedthroughs make contact at $\sim 0.35 \text{ m}$. This electric field map was obtained by applying a known voltage to the plates and determining the resonant rf frequency for different molecular cloud positions within the machine. As the quadratic Stark interaction has been well calibrated [36], these frequency measurements were easily converted to applied electric fields and hence plate spacing measurements.

In our first attempt to fix the bend, we used much softer springs on the feedthroughs to reduce their pressure on the plates. When the plates were reinstalled within the chamber, a new field map revealed that this had no effect on the observed plate spacing. This is

shown by the red points in figure 5.8. It appears that the springs were not responsible for the plate bowing.

As explained previously, the plates were removed a second time to install the new rf transmission line. In a second attempt to fix the bending, the number of dowels suspending the HV plates was reduced to one at either end. We found that certain dowels would not freely slide through the plates and the surrounding support structure. It was believed that these dowels might be bending the plates rather than simply holding them in position. Once the plates were replaced, a repeat electric field map revealed that removing these dowels also had no major effect on the observed field. This is shown by the green points in figure 5.8.

The origin of the observed field variation still remains unresolved. However the slight variation is not a major concern to us, as the geometric phase systematic is negligible, and we should be able to reduce the systematic effects that depend upon rf detuning by broadening the transition linewidth with shorter rf pulses.

5.1.3 Shorter RF Pulses and a More Powerful RF Amplifier

Both the rf detuning phase correction and the electric field voltage offset systematic depend upon rf frequency detuning. While we have no physical model for these effects, it is certainly true (in the RWA) that the dominant molecular interaction with the rf field does depend upon the product of the detuning, δ , with the rf pulse length, τ . We might therefore hope that any spurious effects induced by the rf will also depend upon the $\delta\tau$ product. If this is the case we should be able to make these systematic effects smaller by shortening the rf pulse length. This is indeed the case, as will be shown in section 5.2.

From equation 2.19b it can be shown that interference contrast also decreases as a function of $\delta\tau$. Reducing the pulse length would make the experiment run closer to its optimum sensitivity more often. It would also reduce the attenuation in EDM sensitivity that results from a variation in detuning across the molecular pulse.

To allow the application of shorter rf pulses, we purchased a *CPC MRI plus 7T1000M* rf amplifier to replace the *ar worldwide KAW1050* amplifier used previously. Tests have shown that the upgraded rf system is capable of applying $2\mu\text{s}$ long π -pulses to the molecules, over the whole length of the HV field plates. This is roughly a factor of 10 reduction from the $18\mu\text{s}$ pulse length used previously, which is excellent considering the fact that increasing the plate spacing reduced the rf field strength. Furthermore, the

new rf setup increased the reflectance⁵ at either end of the transmission line from 20% to approximately 45%. A large standing wave does not change the physics presented in section 2.2.2 and above, but it does mean certain pulse times/positions require more injected rf power than others to apply π -pulses to the molecules.

With shorter rf pulses we should require correspondingly larger rf amplitude and frequency steps to ensure the locks continue to work efficiently. As the amplitude steps are multiplicative, these scale naturally with rf power, so the attenuator control voltages need not be changed. For the 2011 dataset, the signal to noise ratio on the rf frequency channels far exceeded requirement. We therefore do not need to scale the frequency steps in proportion to the decrease in pulse length. We have found that 2 kHz frequency steps are more than sufficient for our needs even with short $2\ \mu\text{s}$ pulses. This is beneficial, as the experiment will run closer to zero rf detuning.

5.1.4 New Magnetic Shields

Replacement Inner Magnetic Shield




FIGURE 5.9: Lower half of science chamber following rf transmission line upgrade. Note the pump PMT optics now reside inside the inner magnetic shield. Not drawn to scale.

In order to accommodate the longer rf transmission line, a new ‘inner’ magnetic shield was installed in the upper vacuum chamber. This shield has the same diameter as the previous (170 mm), but is slightly longer at 985 mm in length. In this setup, the pump PMT optics now reside inside the shielded region rather than below it, as shown in figure 5.9 (see also figure 2.7). To protect the optical pumping region from fringe electric fields, which might cause some systematic effect, we placed an earthed aluminium table over

⁵We calculate the transmittance and reflectance from the voltage standing wave ratio (VSWR) of an rf amplitude field map. These are generated from many Rabi flopping curves which are scanned out at various rf pulse times, see [84] and previous theses for full details.

the pump collection optics; a small 12 mm hole in the middle allows the molecular cloud to pass through.

In the new setup the HV plates sit 5 cm lower than they used to. This is advantageous as it moves the parallel plates away from the inner shield end caps, where stray magnetic fields are known to rapidly increase to 200 nT (see [74, 84]). After installing the new field plate assembly, we mapped the ambient magnetic field (in z) along the beam line. This was done by comparing the $|1, -1\rangle \leftrightarrow |1, +1\rangle$ splitting for opposite orientations of a large applied B_z field, at different rf pulse times. As before, all the details of the mapping technique are given in references [74, 84]. With the new HV plate position, the ambient field remains uniform over the whole plate length. All fluctuations are within the statistical uncertainty of the measurement technique, and we observe no deviations towards either end of the HV plates. This is excellent as it means we may use the full length of the plates for EDM measurement⁶.

A Second Outer Magnetic Shield

To reduce the size of ambient magnetic field noise, we purchased a third layer of magnetic shielding to surround the old ‘outer’ magnetic shield. This new shield is made from 1.0 mm μ -metal, and has the dimensions $\varnothing 0.8\text{ m} \times 1.6\text{ m}$.

We measured the impact that this additional shield has on the magnetic field noise by collecting a small dataset containing 2579 blocks, and comparing the *Bartington* magnetometer readings to those from the 2011 dataset. Recall that this magnetometer is located on the vacuum chamber, within all ‘outer’ layers of magnetic shielding.

From the test dataset, we find that the third layer of magnetic shielding has little effect on the size of the ambient magnetic field between the hours of 7pm and 7am. During the day, however, the shields have a much larger effect. This can be seen in figure 5.10 and table 5.1, which compare the width of the $\widehat{M}_E \{E\}_{\text{Bart}}$ distribution for the test dataset to that obtained from the 2011 dataset.

It is not surprising that the distributions are similar during the night, as the experiment previously ran very close to the shot noise limit during this period. However, during the day, there is much more magnetic field noise in the physics department, so the reduction in ambient field is much larger. Consider the values presented in table 5.1. Interestingly the ratio of the distribution widths is less than one, in both columns, for the 90.0 and 95.8 percentiles. This likely indicates that the new outer shield removes the outliers in the distribution both day and night.

⁶Ignoring the technicality introduced by the rf pick-up on the PMT amplifiers.




FIGURE 5.10: Distributions of the E -switch correlated ambient magnetic field as measured by the *Bartington* magnetometer. The blue bars in the histograms, and dashed lines in the cumulative distribution plots are measurements taken with only one layer of outer magnetic shielding. The red bars and solid lines represent measurements with two layers of magnetic shielding. Overnight the distributions appear very similar as there is very little activity in the physics department, so in both cases the experiment runs very close to the shot noise limit during this period. During the day, however, there is much more magnetic field noise in the physics department, so the second outer shield has a much larger influence.

TABLE 5.1: Reduction in width of $\widehat{M}_E \{E\}_{\text{Bart}}$ distribution due to a third layer of magnetic shielding.

Percentile	Ratio of distribution widths	
	7am–7pm	7pm–7am
68.3	0.55	1.1
90.0	0.61	0.91
95.8	0.74	0.76

Averaged over the whole test dataset the *Bartington* magnetometer measures a magnetic field noise of $0.17 \text{ nT} \sqrt{\text{block}}$. We can use the measured correlation between d_e and $\widehat{M}_E \{E\}_{\text{Bart}}$ (as in section 4.2.1) to estimate the magnetic field noise as seen by the molecules in EDM units. We estimate a value of $9.2 \times 10^{-27} \text{ e cm} \sqrt{\text{block}}$, which is 30% smaller than the value averaged over the 2011 dataset ($1.3 \times 10^{-26} \text{ e cm} \sqrt{\text{block}}$). As expected, this improvement is roughly half way between the day and night predicted reductions quoted in table 5.1 for the 68.3% percentile. Of course this 30% reduction in

magnetic field noise does not directly translate to a 30% reduction in the uncertainty in d_e , as the measurement of d_e is predominantly limited by shot noise. However, we expect the shot noise limit to decrease significantly as a result of the upgrades detailed in the following sections, so it is important that we reduce the magnetic field noise so that its fractional contribution to the d_e uncertainty does not increase in the next measurement.

5.1.5 Increased Electric Field

In the old apparatus setup, with the rf system coupling directly onto the HV plates, the maximum voltage that we could be reliably charge the plates to was only ± 6 kV. Above this voltage, the plates would often discharge⁷. This is undesirable as we have to recondition the plates after each discharge, which can take many days. This process involves very slowly increasing the plate voltage to the operating value, whilst switching the HV relays once every ~ 20 s.

When building the new rf setup, we carefully smoothed all sharp edges and corners on both the support structure and the transmission line plates to reduce the chances of a discharge. We also ensured that all low voltage components, such as capacitors and semi-rigid coaxial, were positioned as far away from the HV plates as possible — preferably behind the rf transmission line plates.

Following the apparatus upgrade, we have found that the experiment is now capable of reliably charging the HV plates to ± 7.5 kV, which has increased the applied electric field from 10 kV/cm to 12.5 kV/cm. The increased field translates to an effective electric field of 15.8 GV/cm, giving a 8% reduction in the statistical uncertainty in d_e .

5.1.6 Independent Pump and Probe Laser Frequency Modulations

During acquisition of the 2011 dataset, the LF1 waveform, W_{LF1} , was set to modulate the overall dye laser frequency — stepping the frequency of both the pump and probe laser beams together. We would like the frequency of these beams to be stepped independently so that each can be better locked to its ideal frequency, and to enable correlation analysis to be performed on each.

As described in section 2.3.3 the pump laser beam is split from the main beam after the I_2 frequency lock system, which is used for the stabilisation and modulation of the laser frequency. Recall that we pass the pump beam (twice) through a second AOM, to blue shift its frequency by ~ 172 MHz to become resonant with the $Q(0)(F = 1)$ transition. We

⁷HV discharges are easily diagnosed by off the scale readings from leakage monitors.

added an external signal input to the VCO controlling this AOM, so that its frequency could be computer controlled.

The independent control of the pump and probe beam frequencies is somewhat complex: We modulate the probe laser frequency as we did before, by stepping the overall laser frequency with the I_2 AOM VCO drive. This modulation remains assigned to W_{LF1} . We assign the pump beam frequency modulation to a new waveform, W_{LF2} . On each shot we adjust the voltage to the pump AOM VCO to cancel the frequency steps of the laser (due to W_{LF1}) as well as step the frequency of the pump beam according to W_{LF2} . The pump AOM must also cancel the block-to-block frequency shifts of the laser which result from the $\{LF1 \cdot \delta B\}/\{\delta B\}$ parameter lock. The four voltages sent to the pump AOM during a block are automatically calculated at the beginning of each block, when the I_2 and pump AOM frequencies are measured.

As both the pump and probe lasers drive $Q(0)$ transitions, their linewidths should be roughly the same. We therefore modulate the pump laser by the same amount as the probe laser, which we set to 0.33 MHz. Like the probe beam frequency lock, we lock the pump laser frequency to $\{LF2 \cdot \delta B\}/\{\delta B\}$. However, unlike the probe laser lock, this channel ratio is measured by the un-normalised probe PMT. This is because the dependence of the probe PMT signal on the pump laser frequency cancels out when the probe PMT signal is normalised by the pump PMT signal — the pump PMT signal is also dependent on the pump laser frequency so $\{LF2 \cdot \delta B\}/\{\delta B\}$ just returns noise when measured with the normalised probe PMT.

5.1.7 Third Turbo Pump on the Source Chamber

To improve the vacuum within the source chamber, we purchased a new turbomolecular pump. This pump replaces the flange that the stepper motor drive and Penning gauge used to be mounted on (see figure 2.6). The motor drive and gauge were moved onto the side of a new 8" 'full nipple', which was inserted between the source chamber and the rear turbo pump. The stepper motor drive was re-routed to the Yb target assembly with a long flexible drive shaft.

We have found that increasing the number of turbos pumping on the source chamber from two to three increases the YbF signal by $(21 \pm 3)\%$. The uncertainty is fairly high as target spot degradation is quite hard to account for if the pattern is continually started and stopped. Using equation 4.27, we predict that the signal increase should give a 9% reduction in the statistical uncertainty in d_e .

5.1.8 Updated Lock Channel Analysis Gates

During the acquisition of the 2011 dataset, the analysis gate used to lock experimental parameters in real time was not the same as that which minimised the statistical uncertainty in d_e . As a result, in the final analysis, certain channels reported non-zero laser and rf parameter detunings. In general this was not a problem, however the largest contribution to the $\bar{V}$ systematic uncertainty (section 4.3.2) came from the product of the average frequency detuning of the first rf pulse, $\bar{\delta}_{\text{rf}}^{(1)}$, with the uncertainty in the correlation between d_e and $\delta_{\text{rf}}^{(1)}$. It is unlikely that we can better constrain this gradient, short of acquiring significantly more data blocks or intentionally detuning the rf. However, by updating the lock analysis gate to that which maximises sensitivity we should be able to reduce the average rf detuning and hence the $\bar{V}$ systematic uncertainty.

To get an idea of the level to which we should be able to lock $\delta_{\text{rf}}^{(1)}$, we can re-analyse the 2011 dataset using the same analysis gate as that which the locks used. From this analysis we find that $\{RF1F\}/\{\delta B\} = (2.13 \pm 0.21) \times 10^{-3}$, giving an average detuning: $\bar{\delta}_{\text{rf}}^{(1)} = (80.9 \pm 8.0)$ Hz (ignoring the uncertainty in the conversion factor). Combining this with the gradient $(0.3 \pm 1.1) \times 10^{-28}$ e cm/kHz, as was measured with $\bar{V} = 0.0$ V, this average detuning would limit the $\bar{V}$ uncertainty to be less than 0.15×10^{-28} e cm. This is almost an order of magnitude improvement over the published value: 1.3×10^{-28} e cm.

We have now updated the lock gate timings to those which minimise the experimental uncertainty. We have also adjusted the code so that real time lock channel analysis uses the channel ratios $\langle\{X \cdot \delta B\}/\{\delta B\}\rangle$, as calculated using the non-linear analysis.

5.2 Systematic Limitations of the Upgraded Interferometer

Following the systems upgrade as described in the previous section, we proceeded to test the experiment for systematic error using the same methods as before (section 4.3). After running the experiment for a short while with $2 \mu\text{s}$ rf pulses, we noticed that the high power rf pulses would cause the leakage monitors to register large ~ 100 nA currents when no voltage was applied to the HV plates. Once the plates were charged, these currents would completely disappear. Furthermore, the leakage would only trigger above a certain rf power threshold. As this process was not well understood, we decided to adjust the rf pulse parameters such that the rf induced leakage would never occur. We reduced the required rf power by increasing rf pulse length. The systematic tests presented here are those that were taken with $4 \mu\text{s}$ long rf pulses.

5.2.1 Detuning Phase Correlation

In section 4.2.2 I discussed an empirical correction that we apply for an interferometer phase correlated to the frequency of both rf pulses. This dependency causes a systematic shift if the magnitude of the electric field changes by a small amount, δE , when its direction is reversed.

By applying equation 4.4 to a small test dataset containing 2579 blocks, we measure the average rate of change of phase with rf detuning to be $(19.4 \pm 5.4) \text{ nrad/Hz}$ and $(-10.1 \pm 5.2) \text{ nrad/Hz}$ for the first and second rf pulses respectively.

Before upgrading the apparatus we observed that these gradients depended upon the direction of rf propagation. This difference is no longer so significant. When the data is subdivided by rf manual state, measurements from both states agree within error: For rf1, from $M_{RF} = \text{true}$ data we measure the gradient $(24 \pm 16) \text{ nrad/Hz}$, and $M_{RF} = \text{false}$ data gives $(18.1 \pm 5.9) \text{ nrad/Hz}$; For rf2, $M_{RF} = \text{true}$ data gives $(2 \pm 16) \text{ nrad/Hz}$, and $M_{RF} = \text{false}$ data gives $(-13.8 \pm 5.9) \text{ nrad/Hz}$.

Whilst the overall average gradients do have reasonably significant non-zero values, they are roughly an order of magnitude smaller than those observed before the upgrade. This reduction is more than adequate to reduce the EDM shift (equation 4.9) well below $1 \times 10^{-28} \text{ e cm}$. We therefore no longer need to apply this empirical correction.

Interestingly the reductions in the rf phase gradients are more than the factor of 4.5 reduction in rf pulse length. This may indicate that the mechanism does not have a linear dependence upon $\delta \tau$. The reduction may also be a result of the improved rf polarisation. One other option, somewhat related to the rf polarisation, is that the rf detuning phase correlation might depend upon the conditions local to the rf pulse region; recall that the pulses are now at different positions within the machine. However, since the EDM shift is now well below our statistical uncertainty, the mechanism isn't worth investigating.

5.2.2 Asymmetric Electric Field Reversal

As before, we tested for residual δE effects by applying a large external asymmetry with a modification to the hardware controller software to boost the supply voltages in one electric field relay state. The asymmetries $\pm 4.17 \text{ V/cm}$ were each applied in turn⁸, and data was collected in the manual states: $(M_E, M_B, M_{RF}) = (\text{T}, \text{T}, \text{T}), (\text{F}, \text{T}, \text{T}), (\text{T}, \text{F}, \text{T})$ and $(\text{T}, \text{T}, \text{F})$.

⁸The error bars on the measured asymmetries are completely negligible.




FIGURE 5.11: Gradient in d_e with respect to δE as divided by machine manual state. The measurements are consistent with no dependence on electric field asymmetry. Averaged over all sampled manual states we measure the gradient: $(2.1 \pm 2.9) \times 10^{-28} \text{ e cm / (V/cm)}$.

As shown in figure 5.11 we find no evidence for a dependence on δE in any machine manual state. Averaging all manual state data together, we measure a gradient of $(2.1 \pm 2.9) \times 10^{-28} \text{ e cm / (V/cm)}$, which is consistent with no dependence. If we estimate, as we did in section 4.3.1, that the typical size of internally generated δE is $(0.1 \pm 0.1) \text{ V/cm}$, this limits the EDM shift due to all δE induced effects to be $< 0.6 \times 10^{-28} \text{ e cm}$.

5.2.3 Electric Field Voltage Offset

To test for a voltage offset systematic effect, we acquired EDM data with the large applied offsets⁹: $\bar{V} = -500 \text{ V}$ and $\bar{V} = +500 \text{ V}$, each in the manual states: (T,T,T), (F,T,T), (T,F,T) and (T,T,F). As in section 4.3.2, we observe two distinct effects: 1) a variation in d_e across the time of flight in direct proportion to the detuning of the first rf pulse (which also varies across the time of flight), $d_e(t) \propto \delta_{\text{rf}}^{(1)}(t) \propto \{RF1F\}(t)$; 2) an overall shift in d_e , which is a function of only $\bar{V}$. As before, we model these two components with the equation $d_e^{\bar{V}} = f_1(\bar{V}) \delta_{\text{rf}}^{(1)} + f_0(\bar{V})$.

We find that the correlation between d_e and $\delta_{\text{rf}}^{(1)}$, $f_1(\bar{V})$, is machine state independent. As shown in figure 5.12, we find the correlation reverses with $\bar{V}$, and becomes consistent with zero when $\bar{V}$ is carefully nulled. Note, however, that the errors in each point are correlated by almost all sources of noise (except shot noise), so the error bars are not representative. The uncertainties on the measured values of $f_1(\bar{V})$ are therefore somewhat

⁹With the plates charged to apply a 12.5 kV/cm electric field, we were unable to achieve the $\pm 1 \text{ kV}$ offsets, as we did before. This is because the plates would discharge too often to allow a substantial amount of data to be taken, when charged above $\pm 8 \text{ kV}$. The 500 V offsets were more than sufficient to limit the potential size of this systematic error.




FIGURE 5.12: The direct proportionality between d_e and the detuning of the first rf pulse, as calculated via subdivision of the TOF profile. A linear fit to each dataset gives the gradients: (a) (13.2 ± 2.9) , (b) (-17.9 ± 3.1) , (c) (2.4 ± 2.8) . Each are quoted in units of 10^{-28} e cm/kHz.

inflated because of this correlation. This is not a major concern as the larger uncertainties will just result in a more conservative estimate of the $f_1(0.0) \bar{\delta}_{\text{rf}}^{(1)}$ shift.

Interestingly the $\bar{V} = \pm 500$ V gradients, as quoted in the figure caption, are a similar size to the 102.5 V offset data gradient $(-14.3 \pm 2.0) \times 10^{-28}$ e cm/kHz, which was obtained with the ‘old’ experimental setup using $18 \mu\text{s}$ long rf pulses. This attenuation in $\bar{V}$ dependence is possibly a result of the reduced pulse length and broadened rf linewidth. However, it is hard to give a quantitative comparison here, as we don’t know how $f_1(\bar{V})$ depends on $\bar{V}$.

The analysis of 1073 blocks taken recently with $\bar{V} = 0.0$ V is shown in figure 5.12(c). From this data we measure a gradient of $(2.3 \pm 2.8) \times 10^{-28}$ e cm/kHz, which is consistent with zero. We also measure the average rf1 detuning: $\bar{\delta}_{\text{rf}}^{(1)} = (-0.107 \pm 0.094)$ kHz. This limits the EDM shift due to the $f_1(0.0) \bar{\delta}_{\text{rf}}^{(1)}$ product to be $< 0.6 \times 10^{-28}$, which is encouraging.

After correcting the $\bar{V} = \pm 500$ V data for the detuning correlated shifts, we observe the overall $f_0(\bar{V})$ EDM shift. As demonstrated in figure 5.13, if we subdivide the data by machine manual state, we find the shift is independent of manual state. This is to be




FIGURE 5.13: Gradient in d_e with respect to $\bar{V}$ as filtered by machine manual state. This overall shift is in addition to the variation across the TOF profile. All gradients are consistent within error. Averaged over all sampled manual states, we measure the gradient: $(6.3 \pm 2.0) \times 10^{-30} \text{ e cm/V}$.

expected of an internally induced EDM masquerading interferometer phase. Averaging all the data together we measure the gradient $(6.3 \pm 2.0) \times 10^{-30} \text{ e cm/V}$. Curiously this has the opposite sign to that observed with the original setup (figure 4.7). The rf pulses timings have been changed, however, so this is perhaps not entirely surprising. Nonetheless, the measured gradient is not dangerously large.

Using a HV probe we are able to zero the offset to within 1 V. With this maximum offset we find that the $f_0(\bar{V})$ shift is $< 0.063 \times 10^{-28} \text{ e cm}$. Combining this with the limit on the $f_1(0.0) \bar{\delta}_{\text{rf}}^{(1)}$ product, we can limit the total shift due to both $\bar{V}$ effects to be $< \sqrt{0.6^2 + 0.08^2} \times 10^{-28} \text{ e cm} < 0.6 \times 10^{-28} \text{ e cm}$. This is half the previous uncertainty, what's more its size should decrease with the acquisition of more data blocks, as the uncertainties in $f_1(0.0)$ and $\bar{\delta}_{\text{rf}}^{(1)}$ will decrease.

5.2.4 Perpendicular Magnetic Field Test

As was discussed in section 4.3.3, the motional magnetic field can generate a systematic shift when coupled with a stray magnetic field, $B_s \hat{x}$. This shift is easily calculated by substituting the static perpendicular magnetic field, $B_x = B_s + v \times E_0 / c^2$, into the Hamiltonian quoted in equation 4.17, and comparing the $|1, +1\rangle \leftrightarrow |1, -1\rangle$ splitting in each of the four applied electric and magnetic field switch states. The dashed green line in figure 5.14 represents the calculated EDM shift for the applied electric and magnetic fields: $E_0 = \pm 15 \text{ kV/cm}$ and $B_z = B_0 = \pm 8.9 \text{ nT}$, and a molecular velocity, v , of 590 m/s . As our measured values are blinded, an offset has been added to the theoretical calculation to minimise the mean square deviation from the experimentally measured values, which are discussed below.




FIGURE 5.14: Variation in d_e with applied perpendicular magnetic field, $B_s \hat{x}$. The points are measured values. The dashed green line is a theoretical prediction which assumes a perpendicular field $B_x = B_s + v \times E_0/c^2$. The solid red line is a straight line fit to the data points. This line has a gradient $(7.9 \pm 1.2) \times 10^{-30}$ e cm/nT.

The calculated shift is completely negligible ($< 5 \times 10^{-30}$ e cm) even for a 100 nT stray magnetic field. But to be sure that perpendicular magnetic fields do not affect our EDM measurements, we acquired a small amount of data with a large static magnetic field applied along x . The points in figure 5.14 represent the measured EDM values.

For fields < 100 nT we see no significant shift in d_e . This agrees well with the theoretical prediction. But for large perpendicular magnetic fields, it is clear that the measured values do not fit the prediction. In fact the p-value for the χ^2 value that compares the theoretical model to all the measured data points, is only 0.1%. However, 1000 nT is a very large magnetic field, it is roughly 100 times the $v \times E/c^2 = 8$ nT motional magnetic field. Applying this field would cause the machine to operate very far from its ideal conditions.

Regardless of the unexpected deviations from theory, a linear model is a good fit to the measured points (we calculate a p-value of 33%), and it certainly provides a more conservative estimate of the shift due to stray perpendicular magnetic fields. Using the linear model fit shown by the solid red line in the figure, we measure the gradient $(7.9 \pm 1.2) \times 10^{-30}$ e cm/nT. Stray fields within the machine are typically < 2 nT in magnitude, so we can limit the possible shift to be $< 0.2 \times 10^{-28}$ e cm.

5.2.5 Laser Tests

With the aim of making a new improved measurement of d_e , we asked ourselves if there were any further systematic tests that we could perform on the experiment. Whilst the correlations $\{X \cdot LF1\}$ do provide a wealth of information about how the experiment depends upon laser frequency, we decided there was a number of additional tests that we could perform on the laser system to further support our next measurement.

Probe Laser Tuned to $Q(0)(F=1)$

In principle we could measure d_e with the probe laser frequency tuned to the $Q(0)(F=1)$ transition. This should just invert the interference fringes, giving a signal proportional to $\sin^2 \phi$ instead of $\cos^2 \phi$. In fact probing on $F=1$ would be beneficial as $Q(0)(F=1)$ scatters more photons than $Q(0)(F=0)$, so the LIF signal would be larger (see [80]). We don't probe on $F=1$ as it was previously observed that, when measured this way, d_e becomes dependent on molecular velocity and the probe laser polarisation angle [74]. We believe this dependence to be related to magnetic substructure of the $F=1$ state, that can interact anisotropically with all fields present after the second rf pulse.

When we probe on $F=0$, we inevitably detect a small fraction of the $F=1$ population, as the $Q(0)(F=1)$ transition has a fairly significant (30 MHz) power broadened linewidth. To measure this fraction we blocked the pump laser beam and scanned the probe laser frequency over the $Q(0)$ transition. Figure 5.15 shows the spectrum that was obtained with the laser power set to its usual running value (~ 0.55 mW).




FIGURE 5.15: Optical spectrum showing the two hyperfine components of the $Q(0)$ transition in ^{174}YbF . The red line shows a double Lorentzian fit to the fluorescence signal (blue points). The green points show the transmission through a 150 MHz cavity. The green line is an Airy function fit to the cavity data.

Scans were taken for a range of laser powers about the typical running value. This enabled us to calibrate the dependence in both laser power and frequency, which is necessary as the transition linewidth depends upon laser intensity. During each scan, we simultaneously recorded the transmission signal through a 150 MHz reference cavity. This signal was then used to linearise the spectral data and calibrate the frequency axis. To each scan we fitted a double Lorentzian lineshape. From this we were able to calculate the probability of detecting a molecule in $F = 1$, as a function of probe laser detuning, δ_L , using the equation:

$$\epsilon(\delta_L) = \frac{\mathcal{L}_1(\delta_L)/3}{\mathcal{L}_1(\delta_L)/3 + \mathcal{L}_0(\delta_L)}, \quad (5.2)$$

where $\mathcal{L}_1(\delta_L)$ and $\mathcal{L}_0(\delta_L)$ are the Lorentzian functions fit respectively to the $Q(0)(F = 1)$ and $Q(0)(F = 0)$ spectral peaks. Note that $\mathcal{L}_1(\delta_L)$ is divided by three. This is because, when taking EDM data under normal running conditions, we operate the interferometer at $\phi = \pi/4$; following the second rf pulse, the population in $F = 1$ almost exactly equals the population in $F = 0$.

With the probe tuned to $F = 0$, we calculate that $\epsilon(0 \text{ MHz}) = 0.021 \pm 0.003$. Conversely, when the probe laser is tuned to $F = 1$, $\epsilon(-172 \text{ MHz}) = 0.985 \pm 0.002$. To estimate the uncertainties it was assumed that the laser power drifts by $\approx 20\%$. We also accounted for the possibility that the $F = 0$ and $F = 1$ features may be detuned from the spectral peaks by $\approx 5 \text{ MHz}$. The spectral fit parameter errors were also included in this calculation.

Following suit with all other systematic tests, we took EDM data with the probe laser beam tuned to $Q(0)(F = 1)$, then with the laser tuned to $Q(0)(F = 0)$. By comparing the two values of d_e , and using the two calculated values of ϵ , we were able to measure the gradient $\partial d_e / \partial \epsilon = (-1.3 \pm 2.6) \times 10^{-27} \text{ e cm}$. Combining this gradient with the typical running value $\epsilon(0 \text{ MHz}) = 0.021 \pm 0.003$, we limit the EDM shift to be $< 0.8 \times 10^{-28} \text{ e cm}$. Admittedly this value is relatively large (approaching $1 \times 10^{-28} \text{ e cm}$), but it is entirely limited by the uncertainty in the difference between the two EDM values. With the acquisition of more data we should be able to reduce this limit below the $0.5 \times 10^{-28} \text{ e cm}$ level should we need to.

Pump Laser Frequency Detuning

The pump laser is tuned to the $Q(0)(F = 1)$ transition. Detuning the pump laser frequency should, in theory, only reduce the optical pumping efficiency. This would decrease the pump PMT signal and reduce interference contrast. As the optical pumping process is incoherent, even if there is some population still in $F = 1$ after the molecules leave the pump laser beam this should not cause a problem as there should be no fixed coherence between the $|1, \pm 1\rangle$ states (assuming the laser is linearly polarised). We should,

however, place a limit on the possible EDM shift, assuming this assertion is indeed untrue, like all other subtle effects that we don't fully understand. EDM data was taken with the pump laser frequency lock offset, in turn, by -20 MHz, 0 MHz and +20 MHz. From this data we measured the gradient $(-0.2 \pm 1.9) \times 10^{-28}$ e cm/MHz. With the addition of the $\{LF2 \cdot \delta B\}/\{\delta B\}$ and the new lock gates we are confident that the pump laser can be locked, on average, to within 0.5 MHz. In fact the analysis of some recent data taken under normal running conditions gives the channel value $\{LF2 \cdot \delta B\}/\{\delta B\} = (0.91 \pm 0.15) \times 10^{-3}$. This gives a detuning of ≈ 0.3 MHz. If we assume this value is typical, we limit the EDM shift to be $< 0.6 \times 10^{-28}$ e cm.

Laser Polarisation Tests

A simple way in which a coherence between the $|1, \pm 1\rangle$ states could be introduced by optical pumping is through a predominantly circularly polarised laser¹⁰. Consider, for example, a σ_+ polarised pump laser which is aligned along the x axis. For simplicity, also assume that there are no static electric and magnetic fields in the optical pumping region. As shown in figure 5.16, after passing through the laser beam, all population in $F = 1$ would be transferred into $|0, 0\rangle_x$ and $|1, +1\rangle_x$, as written in the x -basis (ignoring any population lost to dark states). Written in the z basis $|1, +1\rangle_x = \frac{1}{2} \left(|1, -1\rangle + |1, +1\rangle - \frac{2}{\sqrt{2}} |1, 0\rangle \right)$, which has a well defined coherence between the $|1, \pm 1\rangle$ states. These states would acquire a phase difference as the molecules move into the electric field region. If the fringe fields at the bottom of the HV plates change when the electric field is reversed, then it is possible that we would observe a systematic shift in d_e .




FIGURE 5.16: Optical pumping of $F = 1$ by a right hand circularly polarised pump laser. As the laser is oriented along $\hat{x}$, the levels are drawn in the x -basis.

It is hard to think of a mechanism in which a circularly polarised probe laser could generate an EDM masquerading signal, without allowing the detection of a small fraction of $F = 1$ population. Nevertheless, we can empirically measure the possible shift from

¹⁰Note that if the pump laser is more linear than circular all $F = 1$ sublevels would be pumped, so no coherence would be established.

either beam by intentionally making the polarisations highly elliptical. For both beams this simply requires removing the polarising beam cube placed immediately before the vacuum chamber laser port. Since the laser beams are largely linearly polarised before passing through the quarter wave plates, we can access a wide range of ellipticities by rotating the quarter wave plate through 90° .

Measurement of the laser ellipticity is relatively straight forward. As shown in figure 2.6, after intersecting the molecular beam, both the pump and probe laser beams pass out the back of the vacuum chamber through a second laser port. We place a polarising beam cube in a rotating mount after this window, and measure the transmitted intensity on a photodiode as the beam cube is rotated through 360° . Through consideration of the polarisation ellipse shown in figure 5.17, we can relate the maximum and minimum photodiode signals ($S_{\max}$ and $S_{\min}$) to the ratio, ρ , of the circular component, σ , to the linear component, π , of the laser polarisation with the equation:

$$\rho = \frac{\sigma}{\pi} = \left(\sqrt{\frac{S_{\max}}{S_{\min}}} - 1 \right)^{-1}. \quad (5.3)$$




FIGURE 5.17: Decomposition of the laser polarisation ellipse into a linear component (π) and a circular component (σ).

EDM data was taken for at least three values of σ/π in the range -1 to +1, for both the pump and probe laser beams. From this data we measured the gradient $\partial d_e / \partial \rho_{\text{pump}} = (-0.2 \pm 2.2) \times 10^{-27}$ e cm for the pump beam, and $\partial d_e / \partial \rho_{\text{probe}} = (-0.6 \pm 2.0) \times 10^{-27}$ e cm for the probe beam.

With the optics setup for normal running conditions, we measure $\rho_{\text{pump}} = (19 \pm 3) \times 10^{-3}$ and $\rho_{\text{probe}} = (13 \pm 3) \times 10^{-3}$. These values are not consistent with zero but they are very small. Combining these offsets with the measured gradients we limit the possible shift due to imperfect laser polarisation to be $< 0.5 \times 10^{-28}$ e cm for the pump laser and $< 0.3 \times 10^{-28}$ e cm for the probe laser beam.

5.3 EDM sensitivity

Table 5.2 presents a summary of the systematic tests that were discussed in the previous sections. All uncertainties are below the 10^{-28} e cm level and are constrained to $\sim 6 \times 10^{-29}$ e cm. This is excellent news for our next EDM measurement, which will require only the ambient magnetic field correction. The table also provides reassuring verification that none of the upgrades have resulted in any new systematic dependencies.

TABLE 5.2: Summary of systematic uncertainties as measured on the upgraded experimental apparatus.

Source	Uncertainty (10^{-28} e cm)
δE effects (inc. the $\delta_{\text{rf}}^{(i)}$ phase correlation)	0.6
$\bar{V}$ effects	0.6
$B_{\perp}$ effects	0.2
$F = 1$ population in probe LIF	0.8
pump laser detuning	0.6
non-linear probe laser polarisation	0.3
non-linear pump laser polarisation	0.5

In section 5.1.7 it was stated that the new turbo pump increased molecular signal by 20%. However, with the reduced dependence on rf detuning and following a certain amount of source parameter optimisation, we now reliably obtain probe PMT signals of $\{SIG\}_{\text{probe}} = 85 \text{ V}\mu\text{s}$. Incredibly, this is a 70% increase on the average value for the 2011 dataset. In accordance with this we also measure $\{\delta B\}_{\text{probe}} = 11 \text{ V}\mu\text{s}$ and $\{SIG\}_{\text{pump}} = 45 \text{ V}\mu\text{s}$. The signal increase alone should reduce the shot noise limit on d_e by 28%, but when combined with the increased machine length and the larger applied electric field, we predict the shot noise limit to be 57% smaller than the value predicted for the 2011 dataset measurement. To be explicit, we measure a shot noise limit of $1.9 \times 10^{-26} \text{ e cm} \sqrt{\text{block}}$. If we combine this value with the magnetic field noise quoted in section 5.1.4 ($9.2 \times 10^{-27} \text{ e cm} \sqrt{\text{block}}$), we predict a noise limit of $2.1 \times 10^{-26} \text{ e cm} \sqrt{\text{block}}$.

From a small test dataset containing 1074 blocks, that were taken under normal running conditions, we measure the statistical uncertainty on d_e to be $7.4 \times 10^{-28} \text{ e cm}$, at the 68.3% confidence level. Interestingly this value is 14% larger than the uncertainty predicted by the shot noise and magnetic field noise measurements combined. Unfortunately this disparity cannot be easily explained.

One possible cause of the difference might be that, with the new source experimental parameters, source noise has grown in comparison to the shot noise limit. Another

possibility might be that the departure from a Gaussian distribution in the measured values of d_e has also become larger in comparison to the shot noise limit. The current experimental sensitivity is still an active subject of investigation.

Regardless of the small discrepancy between the predicted noise limit and the measured uncertainty, the experiment is now capable of measuring d_e with a sensitivity of $2.4 \times 10^{-26} \text{ e cm } \sqrt{\text{block}}$, which is excellent. With the acquisition of 10,000 blocks (roughly three months running) we should be able to set a new limit on the electron EDM with a statistical uncertainty of $\simeq 2.4 \times 10^{-28} \text{ e cm}$, at the 68.3% confidence level. This is more than a factor of two smaller than the current published limit.

Chapter 6

Conclusions and Outlook

In this thesis I have presented a detailed account of our latest measurement of the electron electric dipole moment using YbF molecules. From a dataset of 6194 blocks, and a series of supplementary measurements that were taken under exaggerated non-ideal conditions, we measured $d_e = (-2.4 \pm 5.7_{\text{stat}} \pm 1.5_{\text{syst}}) \times 10^{-28} \text{ e cm}$. This result is consistent with zero, so we set an upper bound of $|d_e| < 10.6 \times 10^{-28} \text{ e cm}$ at the 90% confidence level. At present this is the tightest limit on the d_e . Furthermore, it is the first time that a molecular measurement of the electron EDM has surpassed the best limit set using atoms. Their large enhancement factors and insensitivity to many systematic effects make polar molecules, such as YbF, great candidates for a new improved measurement.

We have made a number of upgrades to the experimental apparatus that increase EDM sensitivity and reduce systematic effects. By shortening the rf pulse length the systematic dependencies on rf detuning have been reduced to a negligible level. By separating the rf transmission line from the electric field plates, we are able to use the full length of the electric field plates for EDM measurement. A new layer of magnetic shielding has significantly reduced the magnetic field noise inside the machine during the day, when there is much activity in the physics department. Combined with improvements to the molecular source parameters, an increased applied electric field, and better gating of PMT TOF profiles, these improvements have reduced the uncertainty in d_e to $2.4 \times 10^{-26} \text{ e cm} \sqrt{\text{block}}$. This is almost a factor of two increase in the sensitivity of the experiment compared to the sensitivity during the acquisition of the published dataset $(4.5 \times 10^{-26} \text{ e cm} \sqrt{\text{block}})$. With the acquisition of 10,000 blocks (roughly three months continuous running), we should be able to set a new limit on d_e with a statistical uncertainty of approximately $2.4 \times 10^{-28} \text{ e cm}$ — almost a factor of three improvement on the current published upper limit.

In the longer term, after making the next measurement at the 2×10^{-28} e cm level, further improvement in the statistical sensitivity could be gained by installing a cryogenic buffer gas molecular source, which would reduce the molecular velocity by approximately a factor of three, and increase the source flux by an order of magnitude [60, 82, 100, 101]. A significant improvement would also be gained by increasing the efficiency that we detect molecules stimulated by the probe laser beam. Currently we detect less than 0.5% of the fluorescence emitted by the molecules [83]. We could potentially improve this fraction by removing elements in the light collection optics and placing the probe PMT as close to the molecular beam as possible. Adding more probe PMTs, or even an array of highly sensitive photo-diodes are also viable options. With the buffer gas source and improved detection efficiency we should be able to reach a sensitivity of the order of 10^{-29} e cm.

Just recently, we have begun building a new experiment to measure the electron EDM using a fountain of YbF molecules [102]. This is an ambitious project that will be developed in parallel to the current YbF electron EDM experiment. While the sensitivity of the current EDM experiment should be competitive with other electron EDM experiments around the world for many years to come, we are inspired to develop a new highly sensitive experiment that will lead the way in the next generation of EDM experiments.

The development of this experiment will be technically very challenging as many aspects of the experiment are still in their infancy. In the proposed new EDM experiment, we plan to laser cool molecules produced in a cryogenic buffer gas cell. Once cooled, the molecules will be launched up between two parallel plates and an EDM measurement will be made in much the same way as described in this thesis. Development of a cryogenic source of YbF molecules is ongoing, but recent results from the buffer gas team within the research group suggest that we can reliably produce a stable effusive beam of YbF molecules with a speed that is slow enough to be laser cooled [103]. The laser cooling and slowing of molecules has only recently been demonstrated on a very limited number of species, namely: SrF, CaF and YO [104–106]. For YbF, we have measured the Franck-Condon factors, the ratios that characterise the change of vibrational state during the transitions between ground and excited electronic states, and find them suitable for laser cooling [79, 107]. We are currently in the process of building the laser systems that will be required for laser cooling. A molecular fountain of diatomic molecules has yet to be demonstrated. No doubt this aspect of the experiment will be equally complex, and a number of unforeseeable hurdles will need to be overcome before we can even begin making an EDM measurement with it.

Assuming all goes to plan, we expect that the new fountain EDM experiment should be able to run with a repetition rate of 2 Hz, and we expect that we should be able to detect

4.4×10^6 molecules per shot. As the interaction period within the fountain will be ~ 0.5 s — three orders of magnitude larger than the current YbF electron EDM experiment — the new EDM experiment will have the potential to achieve a statistical sensitivity of 6×10^{-31} e cm in eight hours [102]. This level is deep within the region where a non-zero EDM is predicted by extensions to the standard model of particle physics.

Bibliography

- [1] L. Boltzmann, On certain questions of the theory of gases, *Nature* **51**, 413, (1895).
- [2] G. Lüders, Proof of the TCP theorem, *Ann. Phys. (NY)* **2**, 1, (1957).
- [3] W. Pauli, (1955), *Exclusion principle, Lorentz group and reflection of space-time and charge*, In W. Pauli and L. Rosenfeld, *Niels Bohr and the development of physics*, McGraw-Hill, p30.
- [4] V.A. Kostelecký and N. Russell, Data tables for Lorentz and CPT violation, *Rev. Mod. Phys.* **83**, 11, (2011).
- [5] G. Steigman, Observational tests of antimatter cosmologies, *Annu. Rev. Astro. Astrophys.* **14**, 339, (1976).
- [6] C.L. Bennett *et al.* Nine-year Wilkinson microwave anisotropy probe (WMAP) observations: final maps and results, submitted to *Astrophysical Journal*, arXiv:1212.5225 [astro-ph.CO], (2013).
- [7] A.H. Guth, Inflationary universe: A possible solution to the horizon and flatness problems, *Phys. Rev. D* **23**, 347, (1981).
- [8] A.D. Sakharov, Violation of CP invariance, C asymmetry, and baryon asymmetry of the universe, *JETP Lett.* **5**, 24, (1967).
- [9] E.W. Kolb and M.S. Turner, Grand unified theories and the origin of the baryon asymmetry, *Ann. Rev. Nucl. Part. Sci.* **33**, 645, (1983).
- [10] C.S. Wu *et al.* Experimental test of parity conversion in beta decay, *Phys. Rev.* **105**, 1413, (1957).
- [11] M. Goldhaber, L. Grodzins, A.W. Sunyar, Helicity of neutrinos, *Phys. Rev.* **109**, 1015, (1957).
- [12] M. Bardon, P. Franzini, J. Lee, Helicity of μ^- mesons; mott scattering of polarised muons, *Phys. Rev. Lett.* **7**, 23, (1961).

- [13] J.H. Christenson *et al.* Evidence for the 2π decay of the K_2^0 meson, *Phys. Rev. Lett.* **13**, 138, (1964).
- [14] T. Ibrahim and P. Nath, CP violation from the standard model to strings, *Rev. Mod. Phys.* **80**, 577, (2008).
- [15] E.M. Purcell and N.F. Ramsey, On the possibility of electric dipole moments for elementary particles and nuclei, *Phys. Rev.* **78**, 809, (1950).
- [16] N.F. Ramsey, Time reversal, charge conjugation, magnetic pole conjugation and parity, *Phys. Rev.* **109**, 222, (1958).
- [17] A. Angelopoulos *et al.* First direct observation of time-reversal non-invariance in the neutral-kaon system, *Phys. Let. B* **444**, 43, (1998). A. Angelopoulos *et al.*, T-violation and CPT-invariance measurements in the CPLEAR experiment: a detailed description of the analysis of neutral-kaon decays to $e\pi\nu$, *Eur. Phys. J. C* **22**, 55, (2001).
- [18] J.P. Lees *et al.* (the BaBar Collaboration), Observation of time-reversal violation in the B^0 meson system, *Phys. Rev. Lett.* **109**, 211801, (2012).
- [19] L.D. Landau, Conservation laws in weak interactions, *Sov. Phys. JETP* **5**, 336, (1957).
- [20] N.F. Ramsey, Electric-dipole moments of elementary particles, *Rep. Prog. Phys.* **45**, 95, (1982).
- [21] J.J. Sakurai, (1994), *Modern Quantum Mechanics*, Addison-Wesley, 241.
- [22] D. Javorsek *et al.* New experimental test of the Pauli exclusion principle using accelerator mass spectrometry, *Phys. Rev. Lett.* **85**, 2701, (2000).
- [23] M.S. Sozzi, (2008), *Discrete Symmetries and CP Violation*, Oxford University Press.
- [24] G.L. Kane, Supersymmetry: What? why? when?, *Contemporary Physics* **41**, 359, (2000).
- [25] P.G.H. Sandars, Electric dipole moments of charged particles, *Contemporary Physics* **42**, 92, (2001).
- [26] I.B. Khriplovich, Electron dipole moments, present and future, *Nuc. Phys. A* **663-664**, 147, (2000).
- [27] M. Pospelov and A. Ritz, Electric dipole moments as probes of new physics, *Ann. Phys.* **318**, 119, (2005).
- [28] J.M. Pendlebury and E.A. Hinds, Particle electric dipole moments, *Nucl. Instrum. Meth. Phys. Res. A* **440**, 471, (2000).

- [29] L.I. Schiff, Measurability of nuclear electric dipole moments, *Phys. Rev.* **132**, 2194, (1963).
- [30] J.S.M. Ginges and V.V. Flambaum, Violations of fundamental symmetries in atoms and tests of unification theories of elementary particles, *Phys. Rep.* **397**, 63, (2004).
- [31] P.G.H. Sandars, The electric dipole moment of an atom, *Phys. Lett* **14**, 194, (1965).
- [32] E.D. Commins, Electric dipole moments of leptons, *Adv. At. Mol. Opt. Phys.* **40**, 1, (1999).
- [33] P.G.H. Sandars, The electric dipole moments of an atom: I and II, *J. Phys. B: At. Mol. Phys.* **1**, 499 & 511, (1968).
- [34] E.A. Hinds, Testing time reversal symmetry using molecules, *Physica Scripta*. **T70**, 34, (1997).
- [35] O.P. Sushkov and V.V. Flambaum, Parity breaking effects in diatomic molecules, *Sov. Phys. JETP* **48**, 608, (1978).
- [36] B.E. Sauer, J. Wang and E.A. Hinds, Laser rf double resonance spectroscopy of ^{174}YbF in the $X^2\Sigma^+$ state: Spin-rotation, hyperfine interactions, and the electric dipole moment, *J. Chem. Phys.*, **105**, 7412, (1996).
- [37] V.F. Dmitriev and R.A. Sen'kov, Schiff Moment of the Mercury Nucleus and the Proton Dipole Moment, *Phys. Rev. Lett.* **91**, 212303, (2003).
- [38] V.A. Dzuba and V.V. Flambaum, Calculation of the (T,P)-odd electric dipole moment of thallium and cesium, *Phys. Rev. A* **80**, 062509, (2009).
- [39] H.S. Nataraj, B.K. Sahoo, B.P. Das and D. Mukherjee, Reappraisal of the electric dipole moment enhancement factor of Thallium, *Phys. Rev. Lett.* **106**, 200403, (2011). S.G. Porsev, M.S. Safronova and M.G. Kozlov, Electric dipole moment enhancement factor of Thallium, *Phys. Rev. Lett.* **108**, 173001, (2012).
- [40] V.A. Dzuba *et al.* Energy levels and lifetimes of GdIV and enhancement of the electron electric dipole moment, *Phys. Rev. A* **66**, 032105, (2002).
- [41] M.G. Kozlov and V.F. Ezhov, Enhancement of the electric dipole moment of the electron in the YbF molecule, *Phys. Rev. A* **49**, 4502, (1994).
- [42] M.G. Kozlov, Enhancement of the electric dipole moment of the electron in the YbF molecule, *J. Phys. B: At. Mol. Opt. Phys.* **30**, L607, (1997).
- [43] A.V. Titov, N.S. Mosyagin and V.F. Ezhov, P,T-odd spin-rotational Hamiltonian for YbF molecule, *Phys. Rev. Lett.* **77**, 5346, (1996).

- [44] H.M. Quiney, H. Skaan and I.P. Grant, Hyperfine and PT-odd effects in YbF $^2\Sigma$, *J. Phys. B: At. Mol. Opt. Phys.* **31**, L85, (1998).
- [45] F.A. Parpia, Ab initio calculation of the enhancement of the electric dipole moment of an electron in the YbF molecule, *J. Phys. B: At. Mol. Opt. Phys.* **31**, 1409, (1998).
- [46] N.S. Mosyagin, M.G. Kozlov and A.V. Titov, Electric dipole moment of the electron in the YbF molecule, *J. Phys. B: At. Mol. Opt. Phys.* **31**, L763, (1998).
- [47] A.N. Petrov *et al.* Configuration-interaction calculation of hyperfine and P,T-odd constants on ^{207}PbO excited states for electron electric-dipole-moment experiments, *Phys. Rev. A* **72**, 022505, (2005).
- [48] E.R. Meyer and J.L. Bohn, Prospects for an electron electric-dipole moment search in metastable ThO and ThF $^+$, *Phys. Rev. A* **78**, 010502, (2008).
- [49] J. Lee *et al.* An electron electric dipole moment search in the ground state of tungsten carbide molecules, *J. Mod. Opt.* **56**, 2005, (2009).
- [50] A.N. Petrov *et al.* Theoretical study of HfF $^+$ in search of the electron electric dipole moment, *Phys. Rev. A* **76**, 030501, (2007).
- [51] S.K. Lamoreaux and R. Golub, Experimental searches for the neutron electric dipole moment, *J. Phys. G: Nucl. Part. Phys.* **36**, 104002, (2009).
- [52] B. Lee Roberts and W.J. Marciano, (2010), *Lepton Dipole Moments*, Advanced Series on Directions in High Energy Physics — Vol. 20, World Scientific.
- [53] J.H. Smith *et al.* Experimental limit to the electric dipole moment of the neutron, *Phys. Rev.* **108**, 120, (1957).
- [54] C.A. Baker *et al.* An improved experimental limit on the electric-dipole moment of the neutron, *Phys. Rev. Lett.* **97**, 131801, (2006).
- [55] N.F. Ramsey, (1956), *Molecular Beams*, Oxford University Press, 124.
- [56] P.G. Harris *et al.* New experimental limit on the electric dipole moment of the neutron, *Phys. Rev. Lett.* **82**, 904, (1999).
- [57] W.C. Griffith *et al.* Improved limit on the permanent electric dipole moment of ^{199}Hg , *Phys. Rev. Lett.* **102**, 101601, (2009).
- [58] B.C. Regan *et al.* New limit on the electron electric dipole moment, *Phys. Rev. Lett.* **88**, 071805, (2002).
- [59] A.C. Vutha *et al.* Search for the electric dipole moment of the electron with thorium monoxide, *J. Phys. B: At. Mol. Opt. Phys.* **43**, 074007, (2010).

- [60] N.R. Hutzler *et al.* A cryogenic beam of refractory, chemically reactive molecules with expansion cooling, *Phys. Chem. Chem. Phys.* **13**, 18976, (2011).
- [61] R. Stutz and E. Cornell, Search for the electron EDM using trapped molecular ions *Bull. Am. Phys. Soc.* **J1.047**, (2004).
- [62] E.D. Commins and D. DeMille, The electric dipole moment of the electron, in B. Lee Roberts and W.J. Marciano, (2010), *Lepton Dipole Moments*, Advanced Series on Directions in High Energy Physics — Vol. 20, World Scientific.
- [63] B.J. Heidenreich *et al.* Limit on the electron electric dipole moment in gadolinium-iron garnet, *Phys. Rev. Lett.* **95**, 253004, (2005).
- [64] S.K. Lamoreaux, Solid-state systems for the electron electric dipole moment and other fundamental measurements, *Phys. Rev. A* **66**, 022109, (2002).
- [65] C.-Y. Liu and S.K. Lamoreaux, A new search for a permanent dipole moment of the electron in a solid state system, *Mod. Phys. Lett. A* **19**, 1235, (2004).
- [66] F. Fang and D.S. Weiss, Resonator-enhanced optical guiding and trapping of Cs atoms, *Opt. Lett.* **34**, 169, (2009).
- [67] G. Xu, Manipulation and quantum control of ultracold atoms and molecules for precision measurements, *PhD Thesis, University of Texas at Austin*, (2001).
- [68] C. Chin *et al.* Measurement of an electron's electric dipole moment using Cs atoms trapped in optical lattices, *Phys. Rev. A* **63**, 033401, (2001)/
- [69] J.M. Amini *et al.* Electron electric-dipole-moment experiment using electric-field quantized slow cesium atoms *Phys. Rev. A* **75**, 063416, (2007).
- [70] J.J. Hudson *et al.* Measurement of the electron electric dipole moment using YbF molecules, *Phys. Rev. Lett.* **89**, 023003, (2002).
- [71] J.J. Hudson, Measuring the electric dipole moment of the electron with YbF molecules, *PhD Thesis, University of Sussex*, (2001).
- [72] P.C. Condylis, Measuring the electron electric dipole moment using supersonic YbF, *PhD Thesis, Imperial College London*, (2006).
- [73] H. Ashworth, Towards an improved measurement of the electron electric dipole moment, *PhD Thesis, Imperial College London*, (2008).
- [74] D.M. Kara, Towards an electron electric dipole moment measurement using Ytterbium Fluoride, *PhD Thesis, Imperial College London*, (2010).

- [75] J.J. Hudson, D.M. Kara, I.J. Smallman, B.E. Sauer, M.R. Tarbutt and E.A. Hinds, Improved measurement of the shape of the electron, *Nature*, **473**, 4930496, (2011).
- [76] D.M. Kara, I.J. Smallman, J.J. Hudson, B.E. Sauer, M.R. Tarbutt and E.A. Hinds, Measurement of the electron's electric dipole moment using YbF molecules: methods and data analysis, *New J. Phys.* **14**, 103051, (2012).
- [77] B.E. Sauer, J. Wang and E.A. Hinds, Anomalous spin-rotation coupling in the $X^2\Sigma^+$ state of YbF, *Phys. Rev. Lett.* **74**, 1554, (1995).
- [78] J. Wang, Laser and radiofrequency spectroscopy of Ytterbium Fluoride Ground State, *PhD Thesis, Yale University*, (1996).
- [79] X. Zhuang, A. Le, T.C. Steimle, N.E. Bulleid, I.J. Smallman, R.J. Hendricks, S.M. Skoff, J.J. Hudson, B.E. Sauer, E.A. Hinds, M.R. Tarbutt, Franck-Condon factors and radiative lifetime of the $A^2\Pi_{1/2}-X^2\Sigma^+$ transition of ytterbium monofluoride, YbF, *Phys. Chem. Chem. Phys.* **13**, 19013, (2011).
- [80] T.E. Wall *et al.* Lifetime of the $A(\nu' = 0)$ state and Franck-Condon factor of the $A-X(0-0)$ transition of CaF measured by the saturation of laser-induced fluorescence, *Phys. Rev. A* **78**, 062509, (2008).
- [81] M. Weissbluth, (1978), *Atoms and Molecules*, Academic Press Inc., p54 & p101.
- [82] M.R. Tarbutt, J.J. Hudson, B.E. Sauer, E.A. Hinds, V.A. Ryzhov, V.L. Ryabov and V.F. Ezhov, A jet beam source of cold YbF radicals, *J. Phys B: At. Mol. Opt. Phys.* **35**, 5013, (2002).
- [83] J.A. Devlin, PhD student on YbF eEDM experiment, *private communication*, 2012.
- [84] J.J. Hudson *et al.* Pulsed beams as field probes for precision measurement, *Phys. Rev. A* **76**, 033410, (2007).
- [85] B.E. Sauer, D.M. Kara, J.J. Hudson, M.R. Tarbutt and E.A. Hinds, A robust floating nanoammeter, *Rev. Sci. Instrum.* **79**, 126102, (2008).
- [86] G.E. Harrison, M.A. Player and P.G.H. Sandars, A multichannel phase-sensitive detection method using orthogonal square waveforms, *J. Phys. E: Sci. Instrum.* **4**, 750, (1971).
- [87] J.J. Hudson, M.R. Tarbutt, B.E. Sauer and E.A. Hinds, Stochastic multi-channel lock-in detection, arXiv:1307.4280 [physics.ins-det], (2013).
- [88] J.R. Klein, A. Roodman, Blind analysis in nuclear and particle physics, *Annu. Rev. Nucl. Part. Sci.* **55**, 141, (2005).

- [89] P.F. Harrison, Blind analysis, *J. Phys. G: Nucl. Part. Phys.* **28**, 2679, (2002).
- [90] J.W. Tukey, A survey of sampling from contaminated distributions, in I. Olkin *et al.*, (1960), *Contributions to Probability and Statistics*, Stanford University Press, 448.
- [91] J.F. Kenney and E.S. Keeping, (1954), *Mathematics of Statistics Part One*, D. Van Nostrand, 211.
- [92] B. Efron and R. Tibshirani, Bootstrap methods for standard errors, confidence intervals, and other measures of statistical accuracy, *Stat. Sci.* **1**, 54, (1986).
- [93] R.J. Carroll, D. Ruppert and L.A. Stefanski, (1995), *Measurement Error in Nonlinear Models*, Chapman and Hall, ch 2.
- [94] M.V. Berry, Quantal phase factors accompanying adiabatic changes, *Proc. R. Soc. Lond. A* **392**, 45, (1984).
- [95] M.R. Tarbutt, J.J. Hudson, B.E. Sauer, and E.A. Hinds, Prospects for measuring the electric dipole moment of the electron using electrically trapped polar molecules, *Faraday Discussions* **142**, 37, (2009).
- [96] S. Millman, On the determination of the signs of nuclear magnetic moments by the molecular beam method of magnetic resonance, *Phys. Rev.* **55**, 628, (1939).
- [97] I.S. Grant and W.R. Phillips, 2ed (2004), *Electromagnetism*, Wiley, 323.
- [98] S.Y. Liao, 1990, *Microwave devices and circuits*, Prentice-Hall, 76.
- [99] J.B. Hagen, 1996, *Radio-Frequency Electronics: Circuits and Applications*, Cambridge University Press, 222.
- [100] N.E. Bulleid *et al.* Traveling-wave deceleration of heavy polar molecules in low-field-seeking states, *Phys. Rev. A* **86**, 021404(R), (2012).
- [101] H.I. Lu, J. Rasmussen, M.J. Wright, D. Patterson and J.M. Doyle, A cold and slow molecular beam, *Phys. Chem. Chem. Phys.* **13**, 19013, (2011).
- [102] M.R. Tarbutt, B.E. Sauer, J.J. Hudson, E.A. Hinds, Design for a fountain of YbF molecules to measure the electron's electric dipole moment, arXiv:1302.2870 [physics.atom-ph], (2013).
- [103] N.E. Bullied, Slow, cold beams of polar molecules for precision measurements, *PhD Thesis, Imperial College London*, (2013).
- [104] E.S. Shuman, J.F. Barry, D.R. Glenn and D. DeMille, Radiative force from optical cycling on a diatomic molecule, *Phys. Rev. Lett.* **103**, 223001, (2009). E.S. Shuman, J.F.

- Barry and D. DeMille, Laser cooling of a diatomic molecule, *Nature* **467**, 820, (2010).
J.F. Barry, E.S. Shuman, E.B. Norrgard and D. DeMille, Laser radiation pressure slowing of a molecular beam, *Phys. Rev. Lett.* **108**, 103002, (2012).
- [105] M.T. Hummon, M. Yeo, B.K. Stuhl, A.L. Collopy, Y. Xia and J. Ye, 2D magneto-optical trapping of diatomic molecules, *Phys. Rev. Lett.* **110**, 143001, (2013).
- [106] V. Zhelyazkova, A. Cournol, T.E. Wall, A. Matsushima, J.J. Hudson, E.A. Hinds, M.R. Tarbutt and B.E. Sauer, Laser cooling and slowing of CaF molecules, arXiv:1308.0421 [physics.atom-ph], (2013).
- [107] I.J. Smallman, F. Wang, T.C. Steimle, B.E. Sauer, J.J. Hudson, M.R. Tarbutt and E.A. Hinds, *in preparation*.

Appendix A

Expansion of PMT Signal

In this section I derive expressions for the detuning of experimental parameters in terms of the analysis channel values.

The interferometer signal can be modelled with the equation

$$S = A(\delta_1, \delta_2, \delta_3 \dots) \cos^2(\phi) + \beta, \quad (\text{A.1})$$

where β is a constant background and $A(\delta_1, \delta_2, \delta_3 \dots)$ is an amplitude term that accounts for the detuning, δ_i , of the laser and rf parameters from their ideal settings.

To a good approximation the amplitude term can be separated into a product of many terms, each of which is a function of only one parameter detuning.

$$A(\delta_1, \delta_2, \delta_3 \dots) \approx A_1(\delta_1) A_2(\delta_2) A_3(\delta_3) \dots \quad (\text{A.2})$$

This is an approximation as the rf frequency and amplitude settings are correlated, as can be seen in equation [2.19b](#). But for small detunings this approximation holds well.

Provided that the parameter detunings are small, each term, $A_i(\delta_i)$, is quadratic in δ_i .

$$A_i(\delta_i) \approx a_i \left(1 - \left(\frac{\delta_i}{\gamma_i} \right)^2 \right). \quad (\text{A.3})$$

γ_i specifies the width of the amplitude function and a_i is a scaling coefficient. For the detuning of the laser frequency, δ_L , $\gamma_i = \gamma_L$ and $\alpha_i = \alpha_L$ for Gaussian and Lorentzian lineshapes of the form:

$$A_L(\delta_L) = a_L e^{-(\delta_L/\gamma_L)^2}, \quad (A.4) \quad A_L(\delta_L) = \frac{a_L}{1 + (\delta_L/\gamma_L)^2}, \quad (A.5)$$

where γ_L specifies the width of the optical resonance.

From equation 2.19b it can be shown that for rf frequency detuning $\gamma_i = \pi/\tau$, where τ is the rf pulse length, and for rf amplitude detuning $\gamma_i = 2/\pi$. As stated above, the detuning of the rf frequency and amplitude are correlated. As a result, the widths, γ_i , of the frequency and amplitude curves are also correlated. The quoted width of the rf frequency (amplitude) detuning curve is only as quoted when the amplitude (frequency) is perfectly set to its ideal setting. However, the definitions of γ_i are quite resilient to small rf parameter detunings, so the quoted values are good approximations.

Consider experimental parameter X , which may refer to any one of the laser and rf parameters. Let us assume that the PMT signal, S , depends quadratically upon the detuning of X , δ_x , as in equation A.3. During the acquisition of a data block, parameter X is centred on the detuning δ_0 , and is modulated between shots by an amount $2\Delta_x$. The channel value $\{X\}$ measures the change in signal between both switch states of X . So long as the step size ($2\Delta_x$) is small in comparison to the width, γ_x , then:

$$\{X\} = 2\Delta_x \left. \frac{\partial S}{\partial \delta_x} \right|_{\delta_x=\delta_0} \quad (A.6a)$$

$$= \frac{-2\alpha_x \delta_0 \Delta_x}{\gamma_x^2}. \quad (A.6b)$$

Here α_x is some factor that accounts for the average value of the other amplitude terms $\prod_{i \neq x} A_i$. We have also used the fact that the average value of $\cos^2(\phi)$ is $1/2$ when equally sampled on both sides of the central interference fringe.

Another channel which is often of interest, particularly in the case of the laser frequency detuning, is $\{X \cdot \delta B\}$. This channel compares the signal when X and the calibration magnetic field ($B_{\delta B}$) are in the same switch state, to when they are in opposite switch states. Recall that the calibration magnetic field step $2B_{\delta B}$ modulates the interferometer phase, ϕ , by an amount $2\phi_{\delta B}$. We can calculate $\{X \cdot \delta B\}$ according to:

$$\{X \cdot \delta B\} = \frac{1}{2} \left((2\Delta_x)(2\phi_{\delta B}) \left. \frac{\partial^2 S}{\partial \delta_x \partial \phi} \right|_{\delta_x=\delta_0, \phi=\phi_{bg}-\phi_{B_0}} + (2\Delta_x)(-2\phi_{\delta B}) \left. \frac{\partial^2 S}{\partial \delta_x \partial \phi} \right|_{\delta_x=\delta_0, \phi=\phi_{bg}+\phi_{B_0}} \right) \quad (A.7a)$$

$$= \frac{-8\alpha_x \delta_0 \Delta_x \phi_{\delta B}}{\gamma_x^2}, \quad (A.7b)$$

where we have taken the average value of the signal change on both sides of the central interference fringe, as specified by the phases $\pm\phi_{B_0}$ and ϕ_{bg} . $\pm\phi_{B_0}$ is the phase associated with the large magnetic field step, and ϕ_{bg} is a phase offset due to any uncanceled stray magnetic fields. We assume that this offset is small so that the PMT signal is linear in phase. Note that on one side of the curve the phase calibration step is negative. This is because the calibration step modulates the magnitude of the applied magnetic field. In effect it always steps outward from the centre of the interference fringes. As above α_x is some factor that accounts for the average value of the other amplitude terms $\prod_{i \neq x} A_i$.

To normalise out the scaling factor, α_x , we divide $\{X\}$ and $\{X \cdot \delta B\}$ by $\{\delta B\}$. $\{\delta B\}$ measures the change in signal due to the calibration step $2\phi_{\delta B}$, and averages over other parameter switch states. This channel is calculated according to:

$$\{\delta B\} = \frac{1}{2} \left((2\phi_{\delta B}) \frac{\partial S}{\partial \phi} \Big|_{\phi=\phi_{bg}-\phi_{B_0}} + (-2\phi_{\delta B}) \frac{\partial S}{\partial \phi} \Big|_{\phi=\phi_{bg}+\phi_{B_0}} \right) \quad (\text{A.8a})$$

$$= 2\alpha_x \phi_{\delta B} \left(1 - \left(\frac{\delta_0}{\gamma_x} \right)^2 \right), \quad (\text{A.8b})$$

where, as above, we average the signal change for both sides of the interference curve. We also assume the step Δ_x is small in comparison to γ_x . Taking the ratio of these equations to cancel the coefficient α_x :

$$\frac{\{X\}}{\{\delta B\}} = \frac{-\delta_0 \Delta_x}{\phi_{\delta B} \gamma_x^2 \left(1 - (\delta_0/\gamma_x)^2 \right)} \approx \frac{-\delta_0 \Delta_x}{\phi_{\delta B} \gamma_x^2}, \quad (\text{A.9})$$

and

$$\frac{\{X \cdot \delta B\}}{\{\delta B\}} = \frac{-4\delta_0 \Delta_x}{\gamma_x^2 \left(1 - (\delta_0/\gamma_x)^2 \right)} \approx \frac{-4\delta_0 \Delta_x}{\gamma_x^2}, \quad (\text{A.10})$$

Appendix B

The Shot Noise Limit

In this section I present a derivation of the shot noise limit as expressed in equation 4.27.

For a dataset containing N blocks, we can write the total number of photons detected as

$$n_{\text{tot}} = \frac{\{SIG\}}{\kappa} N_s N, \quad (\text{B.1})$$

where $\{SIG\}$ is the average integrated probe PMT signal and κ is the coefficient of proportionality between the PMT signal and the photon count rate. $\{SIG\}/\kappa$ is the number of photons per shot and N_s is the number of shots per block. Using these photons we measure the correlation between interferometer phase and the various experimental parameters. We can use the calibration phase step to calculate the number of photons per radian of interferometer phase. The calibration phase step, $\delta\phi_{\delta B} = 2g\mu_B B_{\delta B} T/\hbar$, changes the number of detected photons by an amount,

$$\delta n_{\delta B} = \frac{2\{\delta B\}}{\kappa} N_s N. \quad (\text{B.2})$$

Note the extra factor of two. This is because of how the channel is defined in equation 3.1 — the summation is divided by N_s not $N_s/2$, so $\{\delta B\}$ measures half the change in signal due to the phase step, $\delta\phi_{\delta B}$. Combining these two expressions, we can write the number of photons per radian of phase as

$$\eta = \frac{\delta n_{\delta B}}{\delta\phi_{\delta B}} = \frac{2\{\delta B\}}{\kappa} N_s N \bigg/ \frac{2g\mu_B B_{\delta B} T}{\hbar}. \quad (\text{B.3})$$

In the shot noise limit, the uncertainty on the total photon count is $\sqrt{n_{\text{tot}}}$. To convert this to an interferometer phase uncertainty, σ_ϕ , we simply divide $\sqrt{n_{\text{tot}}}$ by the number

of photons per radian of phase:

$$\sigma_\phi = \frac{\sqrt{n_{\text{tot}}}}{\eta} = \frac{g \mu_B B_{\delta B} T}{\hbar \{\delta B\}} \sqrt{\frac{\kappa \{SIG\}}{N_s N}}. \quad (\text{B.4})$$

The EDM phase $\phi_{\text{EDM}} = d_e E_{\text{eff}} T / \hbar$, so we can write the uncertainty in d_e as

$$\sigma_{d_e} = \frac{\sigma_\phi}{E_{\text{eff}} T / \hbar} = \frac{g \mu_B B_{\delta B}}{E_{\text{eff}}} \sqrt{\frac{\kappa \{SIG\}}{N_s N \{\delta B\}^2}}. \quad (\text{B.5})$$

which is quoted in equation 4.27. Written in this form, equation B.5 is very useful for converting measured channel values into a shot noise limit. However, this form is not very useful for gaining a physical insight. We can, of course, rewrite equation B.5 in terms of the photon counts, n_{tot} and $\delta n_{\delta B}$:

$$\sigma_{d_e} = \left(\frac{g \mu_B B_{\delta B}}{E_{\text{eff}}} \right) \left(\frac{n_{\text{tot}}}{\delta n_{\delta B} / 2} \right) \left(\frac{1}{\sqrt{n_{\text{tot}}}} \right). \quad (\text{B.6})$$

The first term is factor that converts phase into units of EDM. The second term, $2 n_{\text{tot}} / \delta n_{\delta B}$, shows that the uncertainty depends upon the sensitivity of our experiment to phase. The third term shows the well known fact that the uncertainty decreases as $1 / \sqrt{n_{\text{tot}}}$.

Appendix C

The Upgraded RF Transmission Line

This section contains technical scale drawings and photos of the new rf transmission line assembly. All marked dimensions are in millimetres.




FIGURE C.1: The electric and rf field plate assembly.




FIGURE C.2: Cross section in the xz -plane of the field plate assembly.




FIGURE C.3: Dimensions of the rf transmission line plates. The plates are cut from 2 mm thick aluminium.




FIGURE C.4: Cross section of the plate support structure in the xy -plane. The marked holes are for the rf plate support extensions shown in figure C.5.








FIGURE C.5: The extensions to the plate support assembly which mount the additional rf transmission line plates. The extensions are cut from 2 mm thick aluminium.







(a)



(b)



(c)



(d)

FIGURE C.8: Images of the rf transmission line and electric field plate assembly before being wired up.




FIGURE C.9: The coupling between the coaxial rf feed and the parallel plate transmission line at the top of the plate assembly. Note that the $\lambda/2$ -balun is fed around the grounded support structure to keep it away from the plate region.




FIGURE C.10: The coupling between the coaxial rf feed and the parallel plate transmission line at the bottom of the plate assembly.

Appendix D

Parameter Timings and Saved Values

TABLE D.1: Pattern Generator timing values.

Parameter	Typical value
clock frequency	1 MHz
pad shots	1
points (shots)	4096
flash-to-Q	$255\ \mu s$
valve-to-Q	$355\ \mu s$
scrambler length	$500\ \mu s$
valve pulse length	$195\ \mu s$
scrambler centre time	$1400\ \mu s$
rf centre times	$(1120, 1780)\ \mu s$
rf pulse lengths	$(18, 18)\ \mu s$
attenuator centre time	$900\ \mu s$
attenuator length	$1200\ \mu s$
fm centre time	$900\ \mu s$
fm length	$1200\ \mu s$

TABLE D.2: Data acquired per block. We also save the PG timings and switch waveforms.

Data source	Symbol	Description	Typical Value
laser system	ν_{LF1}	I_2 centre frequency	120 MHz
	$\delta\nu_{\text{LF1}}/2$	I_2 frequency step	340 kHz
	V_{I_2}	I_2 voltage centre	1.2 V
	$\delta V_{I_2}/2$	I_2 voltage step	0.05 V
		pump AOM frequency	87 MHz
		pump polariser angle	0–360 °
		probe polariser angle	0–360 °
rf system		synth amplitude	2.4 dBm
		DCFM range	20 kHz
	$\nu_{\text{rf}}^{(i)}$	centre frequencies	173 MHz
	$\delta\nu_{\text{rf}}^{(i)}/2$	frequency steps	1.6 kHz
	$a_{\text{rf}}^{(i)}$	amplitude centres	(-0.9, 1.2) dBm
	$\delta a_{\text{rf}}^{(i)}/2$	amplitude steps	0.2 dBm
	V_{RFiA}	attenuator voltage centres	(2.9, 3.7) V
	$\delta V_{\text{RFiA}}/2$	attenuator voltage steps	0.1 V
	V_{RFiF}	DCFM voltage centres	(0.7, 0.6) V
	$\delta V_{\text{RFiF}}/2$	DCFM voltage steps	0.1 V
	V_{ϑ}	phase scrambler voltage	0–1.2 V
	M_{RF}	manual state	true/false
electric field system		bleed time	1 s
		charge time	4 s
		ramp up time	2 s
		ramp down time	2 s
		ramp up delay	2 s
		ramp down delay	1 s
		switch time	1 s
		discharge time	3 s
	$V_{\text{E+}}$	+ve supply voltage	+6 kV
	$V_{\text{E-}}$	-ve supply voltage	-6 kV
	M_E	manual state	true/false
magnetic field system	V_{bias}	bias voltage centre	~2.5 V
	I_{bias}	bias current centre	~0 mA
	I_{B_0}	field reversal current step	0.812 mA
	$I_{\delta B}$	calibration current step	0.097 mA
	M_B	manual state	true/false