

Measuring the electric dipole moment of the electron with YbF molecules

Jonathan James Hudson

Submitted for the degree of D. Phil.

University of Sussex

September, 2001

Declaration

I hereby declare that this thesis has not been submitted, either in the same or different form, to this or any other university for a degree.

Acknowledgements

The “I”s in this thesis aren’t just me. A number of other people have contributed to the experiment, all in different ways; here their contributions are acknowledged.

Ed Hinds, my supervisor and principal investigator, has always found the time to answer my questions — both the stupid ones and the not so stupid ones. His patience and good advice are always appreciated.

Ben Sauer has worked in the lab with me on almost everything described within. His friendship and good humour have made the lab a pleasant place to be. Besides, if it wasn’t for the fact he knows almost everything about everything technical we probably wouldn’t have got anything done !

My group is lucky to have excellent technical and administrative support. Much of this experimental work only got done because of their efforts. I specially thank Alan Butler, Sarah Neil and Heather Burton.

I thank the members of SCOAP past and present - they’ve made Sussex a good place to work, and an excellent place to learn.

Finally, trite as it may sound, I thank my friends and family. Their genuine interest in and enthusiasm for what I do is a constant motivation.

Measuring the electric dipole moment of the electron with YbF molecules

Jonathan James Hudson

Submitted for the degree of D. Phil.

University of Sussex

September, 2001

Summary

A molecular beam interferometer has been built. The interferometer is capable of measuring spectacularly small shifts in energy levels of the YbF molecule, such as those that might be induced if the electron has an electric dipole moment (EDM). This device has been used to make a precise measurement of the electron EDM. The result, $(0.3 \pm 4.0) \times 10^{-26} \text{ e} \cdot \text{cm}$, is the most sensitive measurement of the electron EDM made in a molecular environment. Furthermore, the results indicate that further work could substantially increase this sensitivity.

Contents

1	Introduction and overview	1
1.1	Time reversal symmetry and the electron's EDM	1
1.1.1	Time reversal symmetry	1
1.1.2	The electron's EDM	4
1.2	Measuring the EDM in a molecule	6
1.3	EDM measurements - past and present	9
1.4	The YbF molecule	10
1.5	Principle of the experiment	12
1.5.1	Energy differences	12
1.5.2	An interferometric technique	12
1.6	Overview of the experiment	15
2	Making an interferometer	17
2.1	Infrastructure	17
2.1.1	Vacuum system	17
2.1.2	Laser systems	20
2.1.3	Computer control system	21
2.2	Making a beam of YbF molecules	21
2.2.1	Characterising a molecular beam	21
2.2.2	Implementation	23
2.2.3	Operation	24
2.3	Detecting the molecules — laser induced fluorescence	25
2.3.1	Sub-Doppler resolution + many photons = good !	25
2.3.2	Implementation	26
2.3.3	Operation	29
2.4	Pump and probe beams — locking the first laser	30
2.4.1	The I ₂ spectrometer	31
2.4.2	Pump beam	34
2.5	Radio frequency transitions	35
2.5.1	Theory	35

2.5.2	Implementation	38
2.5.3	Operation	41
2.6	Repump — locking the second laser	44
2.6.1	Repump transition	44
2.6.2	Cavity lock	44
2.6.3	Operation	47
2.7	Electric field	49
2.7.1	The Stark shift	50
2.7.2	Implementation	51
2.7.3	Operation	52
2.8	Interference	55
2.8.1	Theory	56
2.8.2	Magnetic field hardware	60
2.8.3	Interference measurements	61
2.9	Comments	62
3	Measuring the EDM	64
3.1	Noise performance	64
3.2	Switching and noise suppression	69
3.3	The control software	73
3.4	Data analysis	76
3.5	Some real clusters	79
3.6	Systematic effects	81
3.6.1	Magnetic effects	82
3.6.2	Measured systematic effects	86
3.7	The final analysis	92
4	Conclusions	96
4.1	Result ?	96
4.2	Future prospects	97
	Bibliography	100

List of Figures

1.1	Fundamental particle modelled as an egg.	5
1.2	Polarization of a ground state rigid rotor.	8
1.3	The important energy levels and transitions in YbF.	11
1.4	A schematic of the interferometer.	13
1.5	Measuring an EDM with the interferometer.	14
1.6	The interferometer, redrawn to suggest the experimental set-up.	15
2.1	A block diagram of the interferometer.	18
2.2	The vacuum chamber.	19
2.3	A picture of the vacuum chamber.	20
2.4	Distribution of rotational population in the $\nu = 0$ manifold.	23
2.5	Quantifying the degree of collimation of a molecular beam.	26
2.6	A plan view of the probe region.	27
2.7	An LIF scan around the Q(0) transitions.	30
2.8	Layout of the I ₂ spectrometer's optical components.	31
2.9	A picture of the I ₂ spectrometer.	32
2.10	The I ₂ spectrometer's control electronics.	33
2.11	A scan over some I ₂ spectrometer lines.	34
2.12	Population transfer vs. Rabi frequency.	38
2.13	The position of the rf antennae.	39
2.14	The position of the magnetic shields.	40
2.15	rf transition intensity vs. rf drive power.	41
2.16	The rf transition lineshapes.	42
2.17	rf transition intensity vs. chamber pressure.	43
2.18	The cavity lock optical layout.	45
2.19	The cavity lock control electronics.	46
2.20	Scan of the second laser offset frequency over the repump transition.	48
2.21	The repump really works !	49
2.22	Stark shifts of the groundstate hyperfine levels.	50
2.23	A close-up view of the electric field plate assembly.	52
2.24	The location of the electric field plates.	53

2.25	rf lineshapes with the field plates charged to their operating voltages. . . .	54
2.26	Watching the electric field decay using rf transitions.	55
2.27	Power dependence of the interference lineshape.	58
2.28	Lineshape slope vs. rf interaction strength.	59
2.29	The optimal interference lineshape.	59
2.30	A measured interference lineshape.	61
3.1	A simple EDM experiment.	65
3.2	The noise power spectrum.	67
3.3	The measured delayed Allan variance.	68
3.4	A more sophisticated EDM experiment.	70
3.5	Points, blocks and waveforms.	71
3.6	Labels for the eight interferometer states.	77
3.7	The laser unlocking.	80
3.8	The magnetic field tracking in action.	81
3.9	Non-zero <i>E</i> Cal and <i>B</i> shift give an <i>EDM</i>	89
3.10	A histogram of the block EDMs.	94
3.11	EDMs of the 51 clusters.	95

List of Tables

1.1	Effective fields for some interesting atoms.	7
1.2	Effective fields for some interesting molecules.	8
3.1	The analysis channels.	72
3.2	A typical day of taking EDM data (taken from the 3rd May 2001).	80
3.3	Analysis channel values for cluster 03May0107.	82
3.4	Some clusters from November 2000.	88
3.5	The corrected <i>EDMs</i>	90
3.6	Some clusters from February 2001.	91
3.7	Analysis channel values for the entire dataset.	93

Chapter 1

Introduction and overview

In this thesis I will describe a very sensitive measurement of the electron’s electric dipole moment (EDM¹).

This chapter begins with a brief introduction to time reversal symmetry and its relevance to physics. This will be followed by a discussion of the relationship between the electron EDM and time reversal symmetry. After explaining the reasons for measuring the EDM using YbF molecules, a short summary of the most sensitive EDM experiments in atoms and molecules will be given. Finally, the experimental technique will be outlined. In chapter 2, the device used to measure the EDM, a spin interferometer based around a beam of YbF molecules, will be described in detail. In chapter 3 I will describe how the interferometer has been used to measure the EDM and present the results of the measurement. Chapter 4 contains some concluding remarks.

1.1 Time reversal symmetry and the electron’s EDM

1.1.1 Time reversal symmetry

I guess it all starts with a question, “Why does time flow ?” Time seems to be marching on inexorably² and there is a strong “intuitive asymmetry” [2] between past and future. Answering the question why this is so has been a longstanding endeavour. The question is intimately bound with our experience of being; as with most fundamental metaphysical questions it seems unlikely to ever be resolved, but the pursuit of its answer is, in itself, a worthwhile activity. The study of such questions crosses many disciplines and it is natural to ask whether physics can contribute anything to the debate. I think an honest but

¹Throughout this thesis, EDM will be used as an abbreviation for both ‘electric dipole moment’ and ‘electron’s electric dipole moment’ — it should be clear which is meant from the context.

²This isn’t the only viewpoint — see [1].

interested physicist would have to answer, “I don’t know, but I’ll try anyway.”

Physics, as the science of the structure and dynamics of the physical world, is ill-equipped to tackle questions of human experience, which necessarily limits its contribution. However, physics *is* well suited to making careful, precise observations, and making these observations in as close to an objective sense as is possible. It is my hope that the careful observations of physics may be used as ‘fuel’ for the ongoing debate over the nature of time.³

To make progress in this thesis though, I will have to leave behind these fascinating questions of time perception and adopt the pragmatic physical view of time. That is, time considered as a numerical parameter that indicates the ordering of events relative to those of a ‘reference’ dynamical system⁴ that are assumed to be periodic. In this framework, the question of the direction of time is approached using the time-reversal transformation (T-transformation) which changes the sign of the time parameter, $t \rightarrow -t$. Our task is to investigate the symmetry of the physical world under this T-transformation.

It doesn’t take a very sophisticated experiment to discover that the physical world is manifestly asymmetric under the T-transformation — dropping a teapot will do. This kind of T-symmetry violation, an ‘entropic’ asymmetry, is summed up physically by the Second Law of Thermodynamics. Through the work of Maxwell, Boltzmann, Gibbs and Loeschmidt in the late 19th century, we now know entropic asymmetry to be a probabilistic feature of systems with a large number of components and highly ordered initial states. The observable universe is in a rather ‘unlikely’ state very far from equilibrium, and continually moves towards the more probable equilibrium state.⁵

There is another, perhaps more fundamental, question concerning the universe’s T-symmetry. Can we perform an experiment on a system with just one component, perhaps a supposedly indivisible component like an electron, that reveals a T-asymmetry of the universe ? In terms of an appealing picture: it’s clear we can tell whether a movie of a teapot being smashed is being played forward or backward, can we do the same for a ‘movie’ of a single electron ? It is performing just this sort of experiment that I will focus on in this thesis.

It is often said that observing T-asymmetry in such a fundamental system is evidence of the ‘laws of physics’ being T-asymmetric. I think caution should be exercised in taking this inferential step. The problem is the usual one, that observation of an asymmetry in the behaviour of a system can be attributed to either an asymmetry in the laws governing

³A well known example of physical observations being used to motivate a theory of time perception is the thesis of Boltzmann, that “what we mean by the future, as opposed to the past, time direction, just is that direction of time in which the entropy change is an increase”. See chapter 10 of [2] for a discussion of Boltzmann’s thesis.

⁴such as a planet orbiting a star, an hour glass, or a Cs atom.

⁵This doesn’t explain how the universe got into its very improbable ordered state though . . .

the system's dynamics *or* an asymmetry in the system's initial conditions.^{6,7} Nonetheless, no matter what the cause of the asymmetry turns out to be, I think the observation of T-asymmetry in a fundamental system is both profound and shocking.

It was, then, a profound and shocking event when, in 1964, the neutral Kaon was found to have a CP-violating decay channel [3]. This result, in combination with the well-known CPT theorem, which is widely *assumed* to be true, implies that the Kaon decay also violates T-symmetry. This discovery of a CP-violating system reinforced the need to study T-symmetry carefully.

Hopefully the reader is convinced by now that T-symmetry is a fascinating aspect of physics that may have profound implications for our view of the universe. Next I hope to show that it's also very relevant to modern physics; T-symmetry could provide clues to solving some of its most pressing problems.

The first problem is that of the Standard Model. Few would argue that the Standard Model is complete, despite its success at describing fundamental particles and their interactions. Some of the often levelled criticisms are: the abundance of 'free parameters' that must be adjusted to bring theoretical predictions into coincidence with experimental results, giving the Model an uncomfortably descriptive feel; the rather arbitrary asymmetric particle structure introduced to explain parity violation; the gauge-hierarchy problem; and, of course, the complete omission of gravity. Recent results from the Super-K collaboration [4] also indicate that neutrinos are massive; the first experimental fact at odds with the Standard Model.⁸

There are many proposed extensions/replacements for the Standard Model that attempt to present a more natural account of particle physics. These theories 'beyond the Standard Model' are all necessarily T-asymmetric, to account for the T-asymmetry displayed by neutral Kaons. Measurements of the degree of T-asymmetry in as many other fundamental systems as possible are necessary to guide and constrain the development of these new theories. Electron EDM measurements turn out to be particularly valuable in this capacity, as will be detailed in the next section.

The second problem is that of baryogenesis, one of the biggest current cosmological mysteries. The observable universe has an overwhelming abundance of matter (versus antimatter) and it has been suggested that CP- and T-violating effects could be responsible [5]. Helping to solve such an interesting and important puzzle as baryogenesis is a strong motivation for the study of T-symmetry.

Hopefully the reason to study T-symmetry is now clear. In the next section I will

⁶or both.

⁷More will be said about the origin of T-asymmetry later, when discussing the electron EDM's relationship to T-symmetry.

⁸A relatively minor and widely accepted extension to the Standard Model does incorporate this result.

explain why the electron is an excellent system on which to carry out these studies.

1.1.2 The electron's EDM

CP-violating effects have been observed in decays of neutral Kaons. To date, this is the only system which is known to violate CP-symmetry. These early results suggested that T-symmetry was also violated, but the inference hinged on the validity of the CPT theorem. In recent years, further experiments on Kaons have confirmed that T-symmetry is violated without recourse to the CPT theorem, but there is still some debate as to whether 'direct evidence' of T-asymmetry has been found [6],[7].

It would seem worthwhile to look for more direct evidence of T-asymmetry, and to investigate systems other than the Kaon. In these respects measurements of the EDMs of fundamental particles are ideal. Consider a light-hearted model of a particle with spin and an EDM: an egg spinning about its symmetry axis — the egg has to spin around its symmetry axis for the model to be consistent with the Wigner-Eckhart theorem. There are two types of spinning egg: those spinning clockwise when viewed from the pointy end, and those spinning counter-clockwise. These two types of egg are related by the T-transformation (figure 1.1(a)). It is important to point out that, so far, T-symmetry remains intact, after all there's nothing strange about spinning eggs. By themselves, non-zero particle EDMs *do not* constitute model-independent evidence of T-asymmetry. However, if only one type of egg were to be found in nature, perhaps the counter-clockwise variety, that would be strange — that would constitute a gross violation of T-symmetry (figure 1.1(b)). There is very strong evidence — inorganic chemistry — that the number of internal states of the electron is consistent with it having spin as its only internal degree of freedom. Therefore, if the electron does have an EDM, only one relative orientation of the EDM to the spin occurs in nature, and T-symmetry is grossly violated. Similar evidence from nuclear structure indicates that the proton and neutron only have one internal degree of freedom.

The usual model of fundamental particles with EDMs does not involve spinning eggs. Conventionally, an EDM is viewed as a phenomenological description of the particle's response to an electric field: this T-asymmetric response is due to T-asymmetric interactions between the 'components' of the particle.⁹ I must stress that this step of associating the T-asymmetry with the interactions — with the laws of nature — however likely to be justifiable, is dependent on a particular model of the particle and should be acknowledged as such. It is possible to conceive of models where the T-asymmetry is a result of initially asymmetric conditions in the universe. Similarly, one can devise models where

⁹These components include the polarised vacuum field and all the associated exotic particles — in this sense even the electron has components.

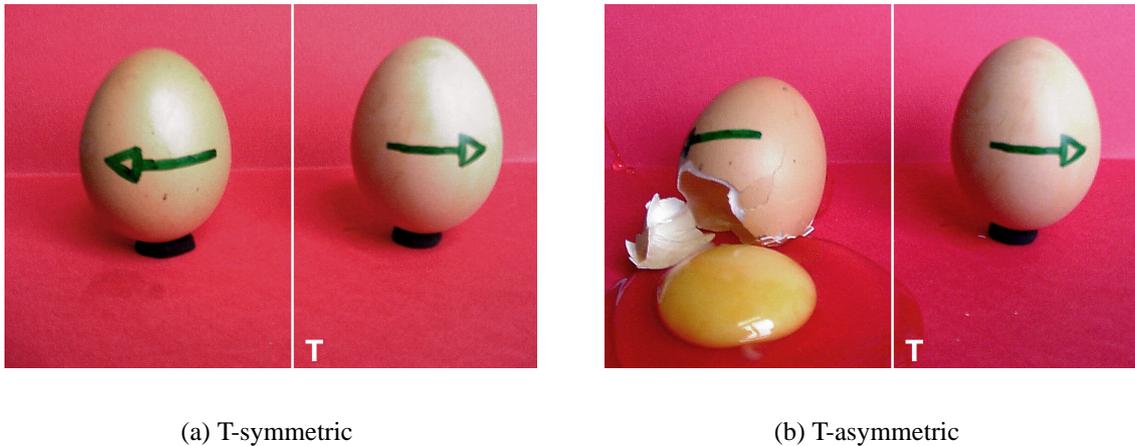


Figure 1.1: Fundamental particle modelled as an egg. The white line is a ‘time reversal mirror’.

the underlying laws are T-symmetric, and the observed T-asymmetry has the status of a spontaneously broken symmetry.

Particle EDM measurements are especially applicable to the first of the problems mentioned in the previous section: extending the Standard Model. The Standard Model predicts exceedingly tiny particle EDMs ($d_{electron} < 10^{-40} \text{ e} \cdot \text{cm}$, $d_{neutron} \simeq 10^{-32} \text{ e} \cdot \text{cm}$ [8]). Extensions to the Standard Model usually introduce extra possible sources of T-asymmetry and predict correspondingly larger EDMs. Typically the EDMs predicted are close to the current experimental limits ($d_{electron} = (6.8 \pm 8.1) \times 10^{-28} \text{ e} \cdot \text{cm}$, $d_{neutron} = (1.9 \pm 5.4) \times 10^{-26} \text{ e} \cdot \text{cm}$ [9],[10]). Trying to measure particle EDMs therefore provides both a sensitive search for physics beyond the Standard Model and a valuable discriminant between potential Standard Model refinements.

Given that we are going to measure a fundamental particle’s EDM, why choose the electron? There are two reasons:

- Leptons are theoretically well understood. QED has proven itself to be the most spectacularly accurate of all physical theories. Moreover, QED calculations of the leptons’ properties are tractable, if not easy. This is in contrast to QCD calculations, which can not be carried out to the same high precision.
- The electron is stable and abundant. Techniques for manipulating electrons and their usual containers, atoms and molecules, have developed to a high degree of sophistication.

Finally, whilst not really a reason, there is a certain attraction to answering high energy particle physics questions using the standard small-scale, low-budget techniques of atomic

physics.

It is appropriate here to give brief mention to the next generation of particle physics experiments concerned with CP- and T-violation. Intense sources of B-mesons have recently come online, and experiments at these facilities are planned that will be very sensitive to CP- and T-violating effects. The only comment I wish to make on these experiments is to highlight that their focus is slightly different from particle EDM experiments. Current particle EDM experiments are *only* sensitive to physics beyond the Standard Model. In contrast, the Standard Model predicts that B-mesons will show significant CP-violation, more so than the Kaon; B-meson experiments are looking for small discrepancies between the measured CP-violation and the Standard Model's predictions.

1.2 Measuring the EDM in a molecule

Now that we've decided to make a measurement of the electron EDM, we need to decide where to get our electrons from. Free electrons are readily available, and indeed the first EDM experiments were carried out on free electrons, as part of an experiment to measure the electron's anomalous magnetic moment [11]. However, free electrons have the significant disadvantage that they are accelerated in electric fields, adding unnecessary complication to and limiting the sensitivity of the experimental technique. Perhaps the next most obvious sources are atoms and molecules. We should investigate what effect an electron EDM would have on the atom/molecule's response to an electric field. Schiff took up this investigation¹⁰ some decades ago and it was quickly realised that the atom/molecule as a whole had the potential to shield an electron EDM from an externally applied field. Schiff's famous theorem [12] showed that for an electrostatically bound system of point particles the shielding is complete — the EDMs of the constituent parts have no observable effect on the whole. Of course, real atoms/molecules aren't electrostatically bound systems of point particles; relativistic effects are important and the nuclei are of finite size. Schiff showed that through these effects the constituents' EDMs *can* have an observable effect on the atom/molecule. Some years later, Sandars discovered the surprising result that for high-Z atoms not only can these effects result in incomplete shielding, they can result in an 'enhancement' of the EDM [13]. His further discovery, that the enhancement could be enormous in highly polar molecules containing a high-Z atom, inspired the search for T-violation in heavy, polar molecules [14]. The first explicit suggestion that the electron's EDM could be measured in heavy, polar molecules was given by Sushkov and Flambaum [15].

The effect of an electron EDM on an atom can be described by adding an effective

¹⁰Strictly, Schiff only investigated the effect of a nuclear EDM on atoms.

Species	α
Cs	114
Tl	-585
Fr	1150

Table 1.1: Effective fields for some interesting atoms. See [8].

interaction term to the Hamiltonian,

$$\hat{H}_{\text{EDM}} = -\alpha g_e d_e \frac{\hat{F} \cdot \vec{E}_{\text{ext}}}{|\hat{F}|},$$

where d_e is the EDM, \hat{F} the total angular momentum and \vec{E}_{ext} the applied electric field. The ‘enhancement factor’ α reflects the effectiveness of the atom/molecule at screening ($\alpha < 1$) or enhancing ($\alpha > 1$) the EDM interaction. The ‘ g -factor’ g_e contains information about how much of the total angular momentum is electron spin. Calculating the enhancement factor for high- Z atoms, or molecules containing high- Z atoms, is far from trivial, nonetheless such calculations have been performed. Table 1.1 lists the enhancement factors for some of the more interesting high- Z atoms; notice that the enhancement factors are very large.

The form of the interaction Hamiltonian for molecules is slightly different as there is a second symmetry axis, the internuclear axis. Molecules suitable for EDM experiments are usually polar and diatomic. In these polar molecules the ‘applied’ electric field is provided by the strong internal electric field. The high- Z ion core is strongly polarised, saturating the EDM enhancement. As the enhancement factor and the ‘applied field’ now depend only on the internal structure of the molecule it is usual to combine them into an effective electric field \vec{E}_{eff} . This effective electric field is directed along the internuclear axis. The interaction Hamiltonian is now written

$$\hat{H}_{\text{EDM}} = -g_e d_e E_{\text{eff}} \frac{\hat{F} \cdot \hat{\lambda}}{|\hat{F}|}, \quad (1.1)$$

where $\hat{\lambda}$ is a unit-vector operator directed along the internuclear axis. Table 1.2 lists the effective electric field for several diatomic fluorides containing a high- Z atom — these molecules are amongst the most promising candidates for an EDM experiment.

The rôle of the external electric field in a molecular EDM experiment is not to polarise the high- Z ion, but instead to align the molecule. The effectiveness of the field at producing this alignment can be calculated by finding the expectation value of (1.1) in a given electric field. This gives the energy shift due to the interaction, which can be written

$$\Delta U = -\eta g_e d_e E_{\text{eff}}, \quad (1.2)$$

Species	\vec{E}_{eff} (GV/cm)	Ref.
YbF	25	[16]
BaF	7.4	[17]
HgF	99	[18]
PbF	-29	[18]

Table 1.2: Effective fields for some interesting molecules.

where η is the ‘polarization factor’, essentially the projection of the internuclear axis onto the total angular momentum, that expresses the alignment of the molecule. The dependence of the polarization factor on the electric field for the ground state of a rigid rotor is shown in figure 1.2.¹¹ The field-axis (x) is in dimensionless units normalized to the molecule’s permanent electric dipole moment μ and rotational constant B .

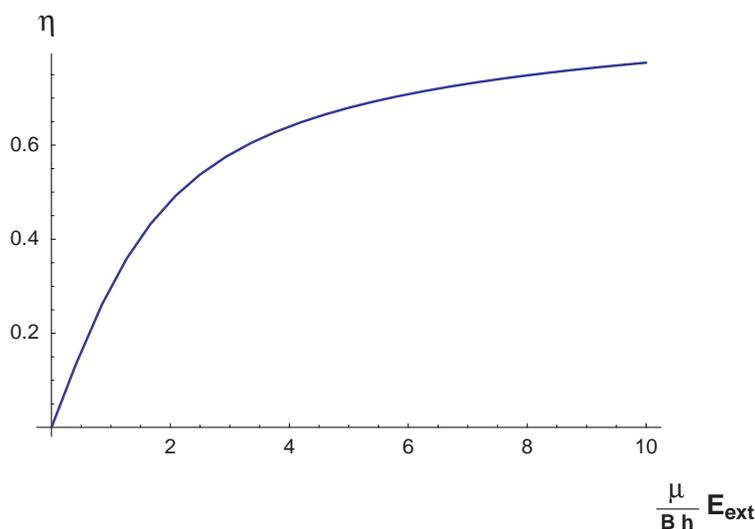


Figure 1.2: Polarization of a ground state rigid rotor.

YbF was chosen for this experiment as it has a large effective electric field (25 GV/cm), much larger than available in experiments that don’t use molecules, and is experimentally feasible.¹² Including all of the other factors in (1.2) the interaction is still very large compared to that with a bare or atomic electron. In the hyperfine state used in the experiment (§2.2) $g_e = 1$ and at 8 kV/cm, the electric field used in the experiment, $\eta = 0.5$, which

¹¹The polarisation factor is most easily calculated by evaluating the gradient of the Stark shift wrt. electric field. The Stark Hamiltonian is $\hat{H}_{Stark} = -\vec{\mu} \cdot \vec{E} = \mu E \cos \theta$; the derivative of this operator wrt. E is proportional to $\cos \theta$. The polarisation factor η is the expectation value of $\cos \theta$ and is therefore proportional to the gradient of $\langle \hat{H}_{Stark} \rangle$ wrt. electric field.

¹²RaF probably has a much higher effective electric field, but the additional experimental and safety concerns do not make it a good choice.

gives an interaction energy,

$$\Delta E[\text{eV}] = 13 \times 10^9 d_e [\text{e} \cdot \text{cm}] .$$

1.3 EDM measurements - past and present

In this section I do not aim to give a comprehensive history of electron EDM measurements — [8] and [19] give good overviews. Instead I shall concentrate on outlining the status of the most sensitive EDM experiments in atoms and molecules. I will also describe the state of the YbF experiment before I started work on it.

The most sensitive EDM measurement to date has been made by Commins and co-workers [9]. This measurement was made on a beam of atomic Tl. The technique is similar to the experiment described in this thesis: they look for a shift of the resonance frequency of an atomic transition that depends on the relative orientation of applied magnetic and electric fields. They measure the resonance frequency using Ramsey's separated oscillatory fields technique. Compared to YbF, Tl has a relatively small enhancement factor, $\alpha = -585$, but they more than make up for this by having a much stronger signal (Tl beams are much easier to make than YbF beams). Their latest, and final, result is $d_e = (6.8 \pm 8.1) \times 10^{-28} \text{ e} \cdot \text{cm}$. The Tl experiment is severely limited by systematic effects, most notably the motional magnetic field effect. The last seven years of heroic effort have controlled these systematic effects allowing the Commins team to improve the limit on d_e by a factor of 1.75 over their 1994 result [20], but it does not seem that they can improve their method any further.

The most sensitive EDM measurement made in a molecule previously was the work of Hinds and co-workers [21]. Their measurement, made on the TlF molecule, yielded a result of $d_e = (-1.4 \pm 2.4) \times 10^{-25} \text{ e} \cdot \text{cm}$. It should be noted that the TlF experiment was not designed to be especially sensitive to the electron EDM, and was primarily aimed at measuring the EDM of the proton. The technique was again similar to the work reported in this thesis: a separated oscillatory fields method was used to measure a frequency shift of the Tl nuclear magnetic resonance. The TlF experiment was stopped because of systematic problems. A renewed effort could probably overcome these problems but conducting an experiment on YbF seems a more attractive possibility.

At present there is only one other molecular EDM experiment being conducted: the experiment using PbO of DeMille and co-workers. Groundstate PbO is spinless so the experiment must be carried out on a metastable excited state. Furthermore, PbO's effective field is not as large as YbF's. PbO does have a significant advantage though: the excited state has very closely spaced rotational levels, meaning that it can be fully polarised ($\eta = 1$) in fields of order a few V/cm. This allows the PbO experiment to be carried out

in a cell, potentially giving long coherence times and large signals. The PbO experiment, being a cell experiment, is likely to encounter different systematic effects than the YbF experiment. It can be viewed as a competitive but complementary technique.

The experiment described in this thesis was started three years before my arrival. The work carried out during that period is described in [22]. In summary, the interferometer (chapter 2) had been built and interference had been observed, but systems weren't in place to allow computer-automated data acquisition. Moreover, the signal:noise was not adequate to allow reliable automated data acquisition. A crude measurement of the EDM had been made 'by hand', giving the result $d_e = (-4.8 \pm 4.2) \times 10^{-23} \text{ e} \cdot \text{cm}$. This thesis reports on work that has improved this result by a factor of 10^3 — many of the interferometer's systems have been upgraded or replaced, an additional 'repump' laser has been added to increase the signal and the experiment has been brought under computer control, to allow automated data acquisition.

1.4 The YbF molecule

To be able to understand the rest of this thesis, a working knowledge of the YbF molecule is needed; providing that knowledge is the aim of this section. I make no attempt to explain the thorny subject of molecular spectra — the interested reader is referred to [23] and [24]. For the purposes of this thesis, it will be sufficient to simply list the molecular states that play a part in the experiment and the transitions between them.

Before that though, it may help to give an idea of the gross structure of the molecule. The bond in YbF is very ionic. Think of Yb ($[\text{Xe}] 4f^{14}6s^2$) giving one of its $6s$ electrons to F ($[\text{He}] 2s^22p^5$) to make a F^- ion — the resulting Yb^+ ion's other $6s$ electron orbits the Yb^{2+} core, and the F^- ion just sticks around. This model misses some important features of the molecule's structure — notably the significant admixture of $4f$ into the valence electron [25]— but is sufficient to understand the important features of the relevant molecular states.

The ground electronic, vibrational, rotational state of the molecule — known as the $\text{X}^2\Sigma^+(\nu = 0, N = 0)$ state — is similar to an alkali groundstate. The most important states in the experiment are the hyperfine states of the groundstate. The molecular ion core has a net spin of $1/2$, due to the F nucleus, and the hyperfine interaction with the "6s" electron is large, giving the groundstate an $F = 1$ triplet and an $F = 0$ singlet.¹³ The hyperfine splitting is approximately 170 MHz. These groundstate hyperfine levels will be referred to as $|1, 1\rangle$, $|1, 0\rangle$, $|1, -1\rangle$ and $|0, 0\rangle$ — if no projection axis is explicitly

¹³The 174 isotope of Yb is chosen because its spinless nucleus considerably simplifies the molecule's level structure. It is also the most abundant isotope.

stated as a subscript, the projection will be assumed to be along the z-axis. If direction is unimportant the states will be simply referred to as $|F = 1\rangle$ and $|F = 0\rangle$.

The molecules are manipulated by driving transitions with lasers. These lasers excite the molecules into the first excited electronic state, the $A^2\Pi_{1/2}(\nu = 0, N = 0)$ state. This state has a similar hyperfine structure to the groundstate, but the hyperfine splitting is small enough (~ 3 MHz) that it can usually be treated as a single state. The transition from the $|F = 0\rangle$ state to the $A^2\Pi_{1/2}(\nu = 0, N = 0)$ state is labelled Q(0) F=0. Similarly, the transition from the $|F = 1\rangle$ state is labelled Q(0) F=1. These transitions require laser light at around 553 nm.

There is one more set of molecular states that are involved in the experiment. These states are the second rotationally excited states of the electronic, vibrational groundstate, the $X^2\Sigma^+(\nu = 0, N = 2)$ states. Of these states the most important is the hyperfine state with $F = 2$ and $J = 3/2$. This state is connected to the $A^2\Pi_{1/2}(\nu = 0, N = 0)$ state by the bizzarely monikered ${}^{\circ}P_{12}(2)$ transition — this transition is driven by laser light 40 GHz red detuned from the Q(0) light.

The levels and transitions are summarised in figure 1.3.

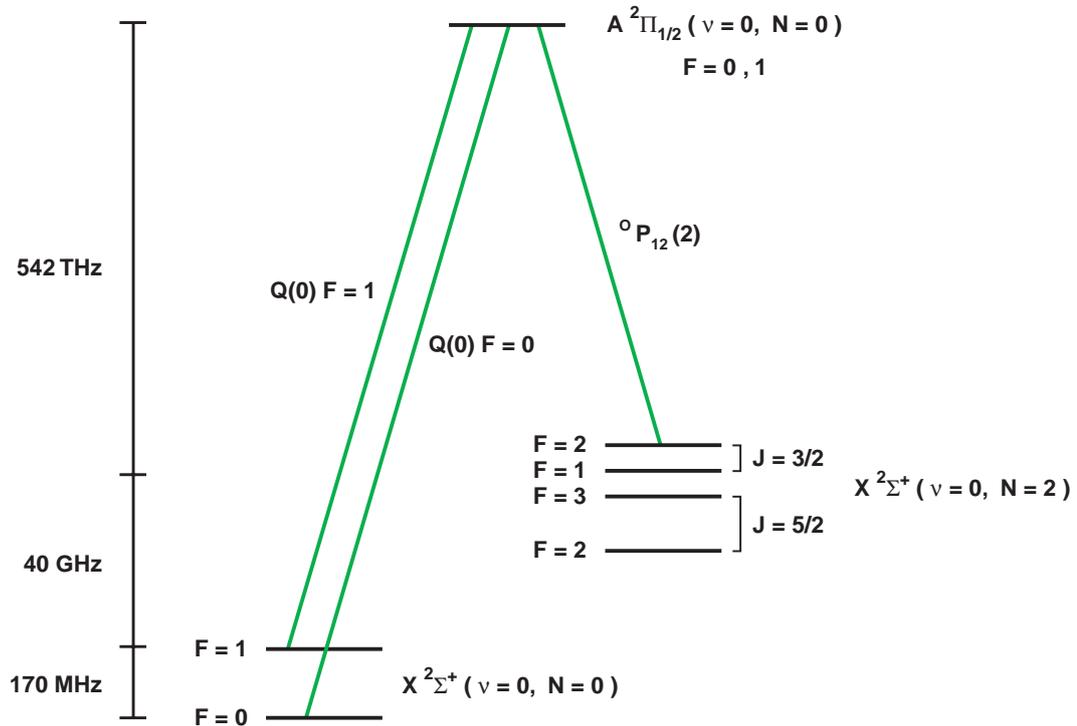


Figure 1.3: The important energy levels and transitions in YbF.

More details on the structure and spectra of YbF can be found in [25], [26], [27], [28] and [29].

1.5 Principle of the experiment

Below I will describe the principle of the experiment without reference to concrete experimental details — these will be the subject of much discussion later on.

1.5.1 Energy differences

Most of the previous EDM experiments, described in section 1.3, are spin precession experiments — changes in the precession rate when electric and magnetic fields are applied are used to infer the EDM. This experiment can also be viewed as a spin precession experiment, but it is much more natural to think in terms of measuring the shift of energy levels under the applied electric and magnetic fields. From this energy shift viewpoint, the experiment can be thought of as an interferometer — this is explained in the next section.

The particular energy shift that I am interested in is the relative shift of the $|1, 1\rangle$ and $|1, -1\rangle$ levels. These levels have oppositely oriented electron spins. As such the energy shifts (given by equation (1.2)) are of opposite sign. In an electric field of magnitude E_{ext} aligned parallel to the spin quantization axis one would expect an energy level difference of

$$\Delta U = 2 \eta g_e d_e E_{eff} .$$

For an applied electric field of 50 kV/cm, enough to significantly polarise the molecule, and an assumed dipole moment of 5×10^{-28} e-cm, just below the current experimental limit, this energy difference is 4 mHz. This is not easily visible as a spectral line splitting as it is a very small fraction of the natural linewidth.

To measure this very small splitting, a technique directly sensitive to the energy *difference* has been adopted. The idea is to prepare a coherent superposition of the two states and follow the time evolution of this superposition, which depends directly on the energy difference. This is the approach, explained further below, that I have adopted.

1.5.2 An interferometric technique

The clearest way to view the experiment is as an interferometer.¹⁴ A schematic outline of the interferometer is given in figure 1.4, time running left to right. In the **Prepare** phase, the molecules are transferred into the $|0, 0\rangle$ state. The **Split** phase leaves the molecules in a coherent superposition of $|1, 1\rangle$ and $|1, -1\rangle$ states — the two arms of the interferometer. The molecules then evolve in an electromagnetic field (the colours in the figure are meant to indicate a changing phase relationship — this will be made rigorous later). The molecules are **Recombined** by driving them back into the $|0, 0\rangle$ state — it is the efficiency

¹⁴It should be noted there is no spatial separation between the ‘arms’ of this interferometer. The interference is between different states of the molecules’ *internal* degrees of freedom.

of *this* transition that depends on the detailed evolution of the superposition in the field. Finally, the number of molecules in the $|0, 0\rangle$ state is measured in the **Probe** phase.

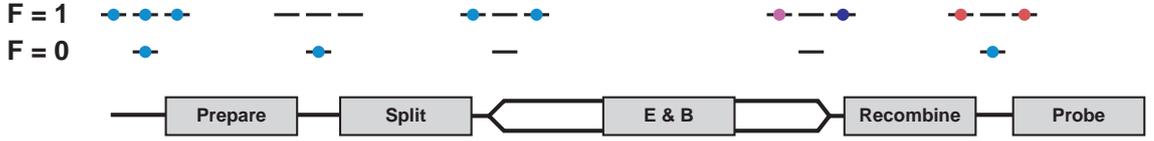


Figure 1.4: A schematic of the interferometer.

A presentation of a simplified model of the interferometer should make the principle of its operation clear.¹⁵ The action of the **Prepare** phase is to project the system onto the state $|0, 0\rangle$. The actions of the **Split** and **Recombine** phases are identical. They can be summarised as the transformations,

$$\begin{aligned}
 |0, 0\rangle &\rightarrow \frac{1}{\sqrt{2}} [|1, 1\rangle + |1, -1\rangle] \\
 |1, 1\rangle &\rightarrow \frac{1}{\sqrt{2}} |0, 0\rangle + \frac{1}{2} [|1, 1\rangle - |1, -1\rangle] \\
 |1, -1\rangle &\rightarrow \frac{1}{\sqrt{2}} |0, 0\rangle - \frac{1}{2} [|1, 1\rangle - |1, -1\rangle] .
 \end{aligned} \tag{1.3}$$

During the **E & B** phase, the interaction with external fields introduces a phase difference $2\Delta\phi$ between the $|1, 1\rangle$ and $|1, -1\rangle$ states,

$$\begin{aligned}
 |1, 1\rangle &\rightarrow e^{i\Delta\phi} |1, 1\rangle \\
 |1, -1\rangle &\rightarrow e^{-i\Delta\phi} |1, -1\rangle .
 \end{aligned} \tag{1.4}$$

These phase shifts can be calculated using first order time-independent perturbation theory. Considering only the magnetic and electric dipole interactions (equation (1.1)) the energy shifts of the $|1, \pm 1\rangle$ levels are,

$$\Delta E^\pm = \langle 1, \pm 1 | H_{int} | 1, \pm 1 \rangle = \pm (\mu_B B_{ext} - \eta d_e E_{eff}) ,$$

for magnetic and electric fields aligned along the z-axis. This leads to a phase shift

$$\Delta\phi = \frac{1}{\hbar} (\mu_B B_{ext} - \eta d_e E_{eff}) \tau = \Delta\phi_B - \Delta\phi_E , \tag{1.5}$$

where τ is the interaction time. The **Probe** phase will simply measure the probability of being in state $|0, 0\rangle$.

¹⁵A more complete presentation will be given in section 2.8.1.

Following the state of the molecule through the interferometer, using equations (1.3) and (1.4),

$$\begin{aligned}
 \text{Prepare :} & \quad |\Psi_1\rangle = |0, 0\rangle \\
 \text{Split :} & \quad |\Psi_2\rangle = \frac{1}{\sqrt{2}} [|1, 1\rangle + |1, -1\rangle] \\
 \text{E \& B :} & \quad |\Psi_3\rangle = \frac{1}{\sqrt{2}} [e^{i\Delta\phi}|1, 1\rangle + e^{-i\Delta\phi}|1, -1\rangle] \\
 \text{Recombine :} & \quad |\Psi_4\rangle = \frac{1}{2} (e^{i\Delta\phi} + e^{-i\Delta\phi}) |0, 0\rangle + \dots
 \end{aligned}$$

The output of the interferometer, the probability of being measured in state $|0, 0\rangle$, is then given by,

$$I = |\langle 0, 0 | \Psi_4 \rangle|^2 = \frac{1}{4} (e^{i\Delta\phi} + e^{-i\Delta\phi})^2 = \cos^2(\Delta\phi) . \quad (1.6)$$

Combining equations (1.6) and (1.5) we reach the result,

$$I = \cos^2 \left[\frac{1}{\hbar} (\mu_B B_{ext} - \eta d_e E_{eff}) \tau \right] . \quad (1.7)$$

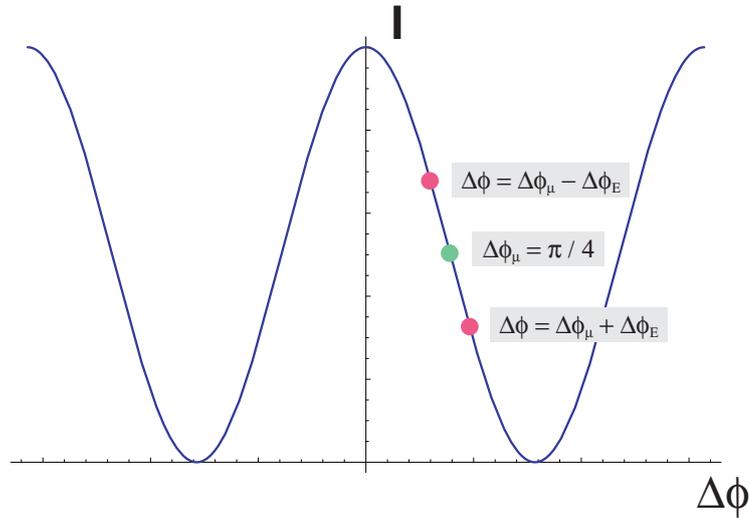


Figure 1.5: Measuring an EDM with the interferometer — the EDM induced phase $\Delta\phi_E$ has been greatly exaggerated !

Armed with such an interferometer, we are in a position to make a measurement of the EDM (figure 1.5). For maximum sensitivity the magnetic field B_{ext} should be adjusted to induce a magnetic phase difference $\Delta\phi = \pi/4$. Looking for a change in signal synchronous with the reversal of the direction of a large electric field constitutes a measurement of the EDM phase $\Delta\phi_E$ and hence the EDM.

1.6 Overview of the experiment

Before we head into the detailed description of the implementation of the interferometer described above, it is appropriate to present a brief overview — chapter 2 will fill in the details. A complete description of how this interferometer is used to measure the EDM is given in chapter 3.

Figure 1.4 can be redrawn to correspond to the experimental setup — figure 1.6. The YbF molecules are produced as an effusive beam, travelling left to right in the figure. The beam issues from a crucible containing the YbF precursors, heated to about 1500 K (§2.2). The beam is housed inside a vacuum chamber — the operating pressure is typically 4×10^{-7} Torr (§2.1.1).

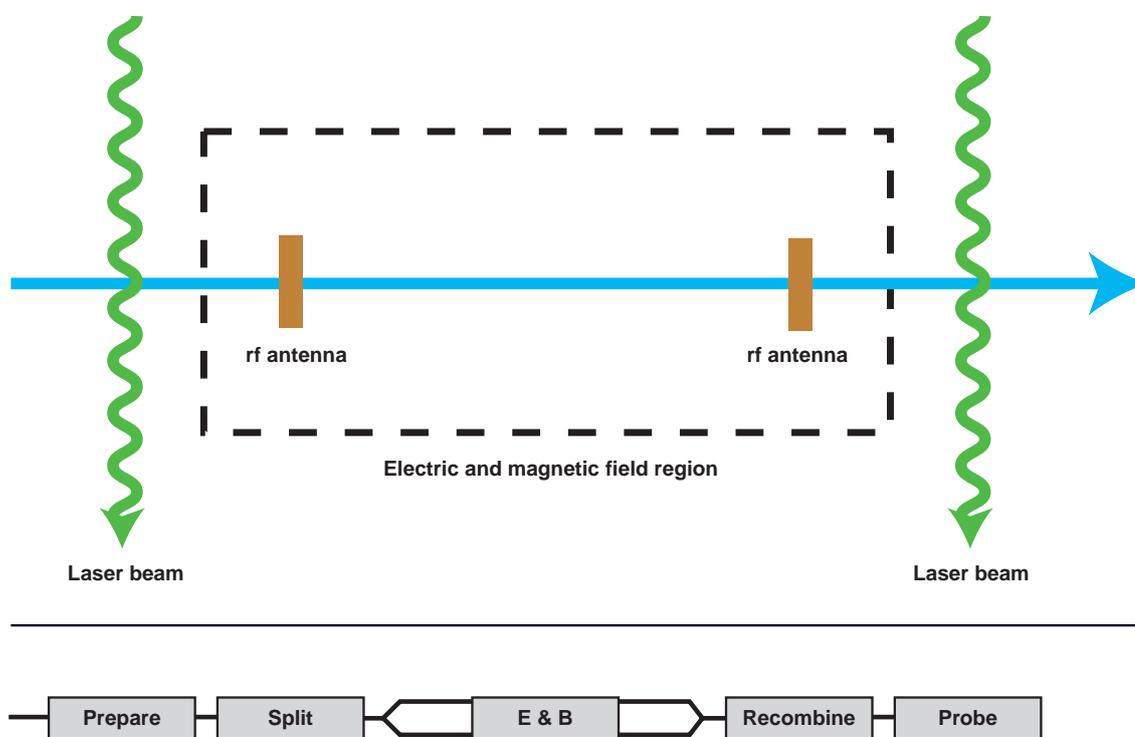


Figure 1.6: The interferometer, redrawn to suggest the experimental set-up.

During their flight through the beam machine the molecules interact with three laser fields, derived from two tunable single-mode dye lasers producing light near 553 nm (§2.1.2). The first of these lasers is stabilised by reference to an I_2 saturated absorption spectrometer (§2.4). This laser produces light resonant with the $Q(0) F=0$ transition and an acousto-optic modulator shifts some of the laser's output to produce light to drive the $Q(0) F=1$ transition. The second laser is locked to the first laser via a stabilised Fabry-Perot cavity (§2.6). This laser produces light resonant with the ${}^oP_{12}(2)$ transition.

During the **Prepare** phase the molecules interact simultaneously with two laser fields as they fly through two overlapped laser beams. The first field, known as the pump beam, drives the Q(0) F=1 transition. This field pumps molecules out of $|1, \pm 1\rangle$ and $|1, 0\rangle$ states. The second field is resonant with the ${}^{\circ}\text{P}_{12}(2)$ transition and acts as a repump, driving molecules lost to the $\text{X}^2\Sigma^+(\nu = 0, N = 2)$ states back into the $\text{X}^2\Sigma^+(\nu = 0, N = 0)$ manifold. The combined action of these two fields results in a population imbalance within the $\text{X}^2\Sigma^+(\nu = 0, N = 0)$ manifold favouring the $|0, 0\rangle$ state.

The **Split** and **Recombine** phases are both carried out by driving radio frequency (rf) transitions between the $|1, \pm 1\rangle$ and $|0, 0\rangle$ states. Two independent rf synthesis chains generate radiation at a frequency of around 170 MHz. The rf radiation is fed to two circular loop antennae that the beam flies through — one for splitting, one for recombining (§2.5). The radiation is resonant with the $|1, \pm 1\rangle \rightarrow |0, 0\rangle$ transition. The orientation of the loops is such that the oscillating magnetic field drives transitions between the states $|0, 0\rangle \leftrightarrow \frac{1}{\sqrt{2}} [|1, 1\rangle + |1, -1\rangle]$.

The splitter and recombiner antennae are at opposite ends of the field region. It is in this region that the **E & B** phase takes place. In the field region the molecular beam flies between electric field plates that can be charged to create fields up to 15 kV/cm (§2.7). The direction of this electric field can be easily reversed by a relay system connecting the field plates to the voltage supplies. The field region is also equipped with magnetic field coils to allow a small magnetic field (~ 10 nT) to be applied in any direction (§2.8). The molecules take around 1 ms to travel the length of the field region.

After the molecules leave the field region, the **Probe** phase takes place. A laser induced fluorescence method is used (§2.3). The molecules fly through a final laser beam, the probe beam, which drives the Q(0) F=0 transition. A photomultiplier, its axis perpendicular to both the molecular and laser beams, counts photons spontaneously emitted from excited molecules.¹⁶ The fluorescence rate measured is the ‘read-out’ of the interferometer.

¹⁶... and, sadly, quite a lot of other photons from other transitions, black body light from the oven, laser scatter etc.

Chapter 2

Making an interferometer

In the context of measuring the EDM, the molecular interferometer described briefly in chapter 1 can be viewed as a tool for measuring small energy differences. In this chapter I will describe the construction of this tool. A block diagram of the interferometer's control systems is presented in figure 2.1.

Much of the interferometer existed before the work reported in this thesis was started. The state of the interferometer at that time, and the limit placed on the EDM, have been described previously (§1.3). A complete description of the interferometer will be given, but emphasis will be placed on parts that are different from those described in [22].

The structure of the presentation follows the structure of a day in the lab. Most of the tasks described below are carried out, at a somewhat increased pace, every time I want to make the interferometer work.¹

2.1 Infrastructure

Some components of the interferometer, the vacuum system, the lasers and the computer control system, are integral parts of many of the interferometer's systems. They are described here.

2.1.1 Vacuum system

It is necessary to conduct the experiment in high vacuum, in order that molecules are not scattered out of the beam by collisions with background gas molecules (§2.2). To this end, the beam is created in a vacuum chamber, shown in figure 2.2. The chamber is made from stainless steel sections connected by Conflat style connectors. These connectors are

¹The reader is spared descriptions of “rebooting the computer” and “mopping the floor” etc.

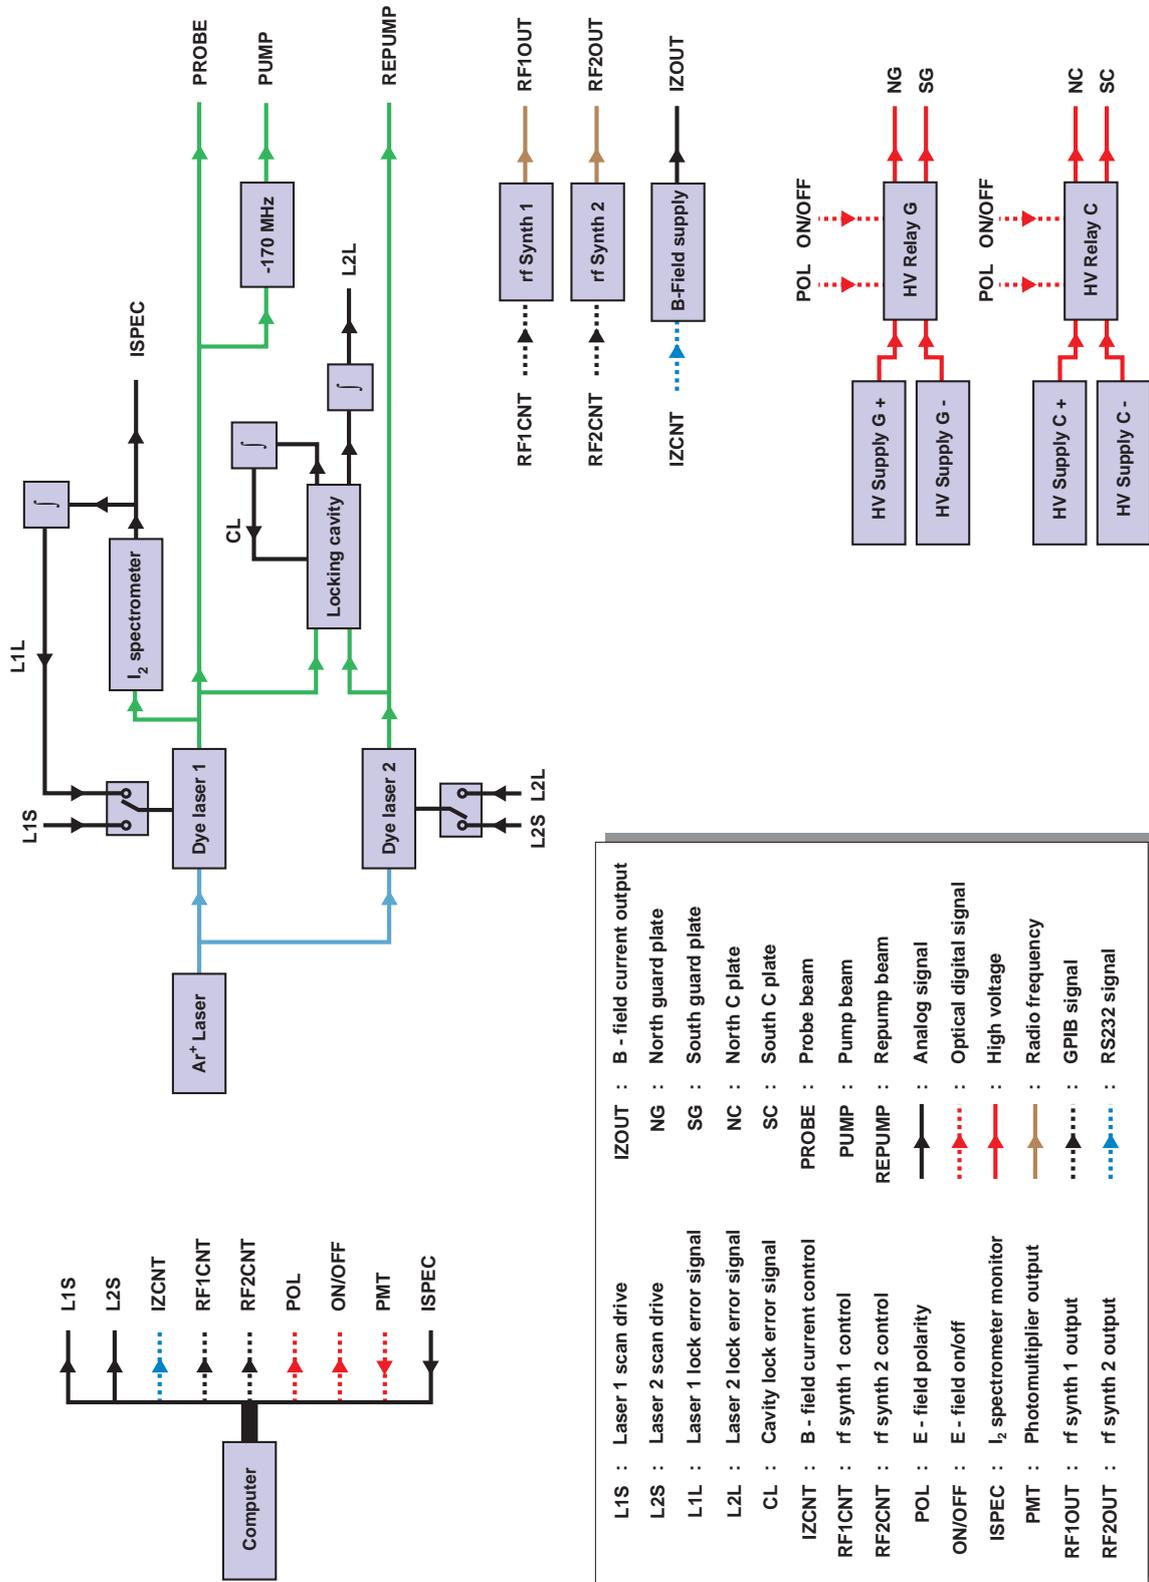


Figure 2.1: A block diagram of the interferometer.

sealed with copper gaskets or Viton rubber o-rings, the choice depending upon how often they are opened.

The chamber is pumped by two turbomolecular pumps, one near the bottom, and one on the top. The turbo pump exhausts are backed by a rotary pump. Together, these pumps reduce the pressure in the chamber to $< 4 \times 10^{-7}$ Torr. As will be discussed in section 2.5.3, this pressure is low enough to reduce the effect of collisions with background gas molecules to a negligible level.

Figure 2.3 shows a picture of the vacuum chamber. The view in the picture would correspond to that in figure 2.2 if the machine were rotated such that the labelled laser ports pointed directly away from the viewer.

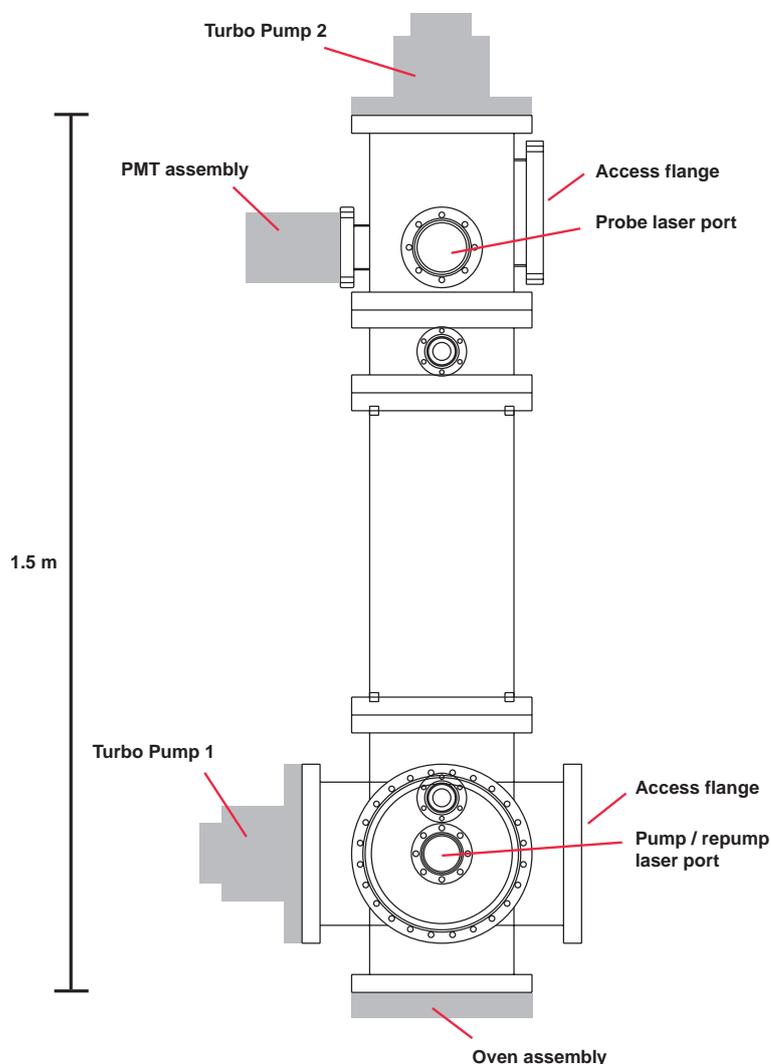


Figure 2.2: The vacuum chamber.

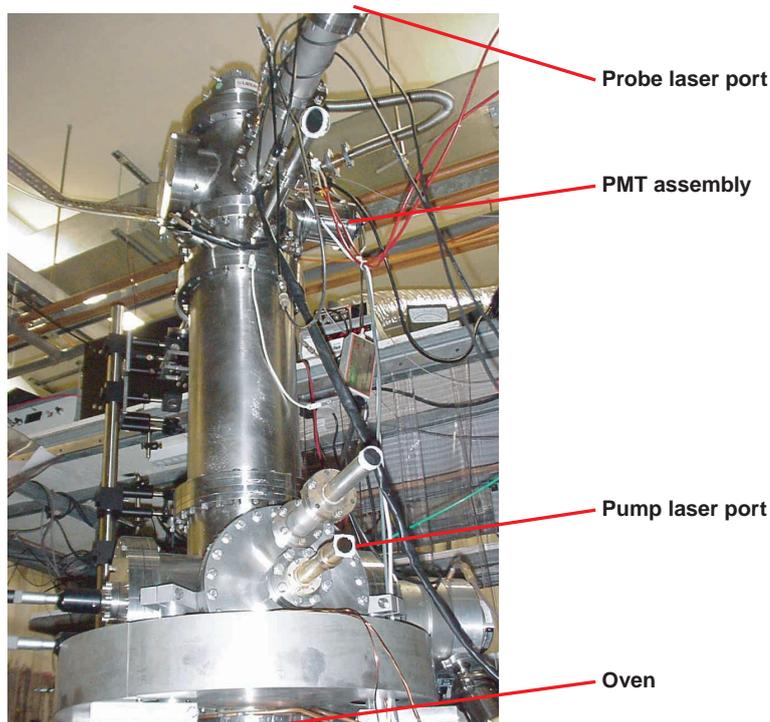


Figure 2.3: A picture of the vacuum chamber.

2.1.2 Laser systems

Two independent single-mode tunable dye lasers are used to generate the three frequencies of laser light needed.² Both of these lasers share a common pump laser, a Spectra Physics 2580 Ar⁺ laser, lasing on all visible lines. The output from this laser is divided in variable proportion between the two dye lasers using a half waveplate and a polarising beam splitter. In this way the pump power for each dye laser can be controlled independently. This considerably simplifies the operation of the dye lasers, whose dyes age at different rates.

The first dye laser, a Spectra Physics 380D, is used to generate the Q(0) pump and probe beams. Rhodamine 110 dye is pumped by 4–6 W of Ar⁺ light, producing 200–500 mW of light at 553 nm. Under these conditions the dye must be changed after approximately 7 days of operation. The laser is equipped with a control unit and reference station. The reference station contains two temperature stabilised Fabry-Perot cavities. With reference to these cavities the control unit stabilises the laser frequency to approximately 500 kHz, and scans the laser output frequency smoothly over a few GHz. The absolute frequency stability provided by the control unit and reference station, around 5 MHz·hr⁻¹, is not adequate for proper operation of the interferometer — the I₂ saturated

²Q(0) F=0, Q(0) F=1 and ^oP₁₂(2).

absorption lock remedies this (§2.4).

The second dye laser, a Coherent 699, produces light for the $^{\circ}\text{P}_{12}(2)$ repump beam. 5–8 W of Ar^+ light again pump Rhodamine 110. Typically 200–400 mW of light near 553 nm is output. The lifetime of the dye in this laser is much shorter, typically 2 days of operation — it is believed that this is due to an element of the circulator poisoning the dye. The laser is stabilised by its control unit to a single, temperature stabilised, Fabry-Perot reference cavity. The required absolute frequency stability is achieved by locking it to the first laser with the cavity lock (§2.6).

2.1.3 Computer control system

Coordination and automation of the experiment is provided by a computer control system. It is the responsibility of the computer to scan the lasers, control magnetic and electric field supplies, and read and record data.

A Pentium II 350 MHz powered PC running Windows 2000 forms the heart of the computer system. The computer is fitted with a National Instruments LabPC+ data acquisition board. The board equips the computer with analog and digital input and output capabilities and three versatile programmable counters. The LabPC+ board interfaces to the other parts of the experiment through a ‘breakout box’ containing TTL \leftrightarrow optical transmitters and receivers. Optical transmission of digital signals reduces the effect of electrical noise created by the high voltage relays as well as reducing the chance of digital control signals coupling directly to the interferometer’s output — a possible cause of systematic error. The breakout box also has a transceiver to provide an optical interface to the computer’s RS 232 serial port. A second board allows the computer to communicate with instruments using the GPIB protocol.

Small programs to automate routine tasks, such as recording molecular spectra, have been written using National Instruments’ Labview package. The program responsible for coordinating the acquisition of EDM data is written in C++ using Microsoft’s Visual C++ development environment (§3.3). All data analysis and database software is implemented with Wolfram Research’s Mathematica system.

2.2 Making a beam of YbF molecules

The starting point of any day in the lab is to make a beam of molecules.

2.2.1 Characterising a molecular beam

The YbF beam issues from an effusive source, one for which the mean free path of the molecules inside the source is large compared to the dimension of its exit hole [32]. Un-

der these conditions the molecules simply ‘leak’ out by chance, their thermodynamic properties being largely unchanged by their escape.

The beam can be characterised by two parameters, the effective source temperature, T and a total intensity, Q , defined as the number of molecules reaching the detector³ per second. Of interest is the relationship between these parameters and the states of the molecules’ external and internal degrees of freedom.

The distribution of molecular velocities in the beam is given by the probability distribution

$$P(v) = \frac{2}{\alpha^4} v^3 \exp\left(\frac{-v^2}{\alpha^2}\right), \quad (2.1)$$

with

$$\alpha = \sqrt{\frac{2kT}{m}},$$

where k is Boltzmann’s constant, and m the molecular mass. This differs from the Maxwell-Boltzmann distribution by a factor of v (and a different normalisation, of course) reflecting the increased probability for fast moving molecules to leave the source.

The parameter α can be related to the characteristic velocities of the beam. The most probable velocity is $v_p = 1.22\alpha$ and the mean velocity is $\bar{v} = 1.33\alpha$. For the YbF oven, which has a temperature of ~ 1500 K, these characteristic velocities are calculated to be

$$\begin{aligned} v_p &= 440 \text{ m.s}^{-1}, \\ \bar{v} &= 480 \text{ m.s}^{-1}. \end{aligned}$$

The molecules’ external and internal degrees of freedom are in thermal equilibrium; the rotational and vibrational temperatures are the same as the kinetic temperature. It is straightforward to calculate the distribution over rotational and vibrational states,

$$P(N, \nu) = \frac{(2N + 1) \exp\left(-\frac{h(BN(N+1) + \nu_0(\nu+1/2))}{kT}\right)}{\sum_{n=0}^{\infty} \sum_{m=0}^{\infty} (2n + 1) \exp\left(-\frac{h(Bn(n+1) + \nu_0(m+1/2))}{kT}\right)},$$

where N and ν are the rotational and vibrational quantum numbers of the state, B and ν_0 are the $X^2\Sigma^+$ electronic state rotational and vibrational constants and h is Planck’s constant. The molecular constants for YbF are $B = 7.4$ GHz and $\nu_0 = 1.4 \times 10^4$ GHz. Calculation shows that at $T = 1500$ K over 40% of the molecules are in the $\nu = 0$ rotational state; the situation is very different for the rotational states however. The probability of being in a rotational state N in the $X^2\Sigma^+$ ($\nu = 0$) manifold is plotted in figure 2.4. It can be seen that the peak of the distribution is near $N = 45$. Only $\sim 10^{-4}$ of the molecules leave the oven in the states that contribute to the interferometer signal.⁴

³This is chosen as it is more relevant than the rate of molecules leaving the oven, and takes into account losses to the beam collimation slits etc.

⁴These are the $\nu = 0, N = 0$ state and one of the hyperfine levels of the $\nu = 0, N = 2$ state.

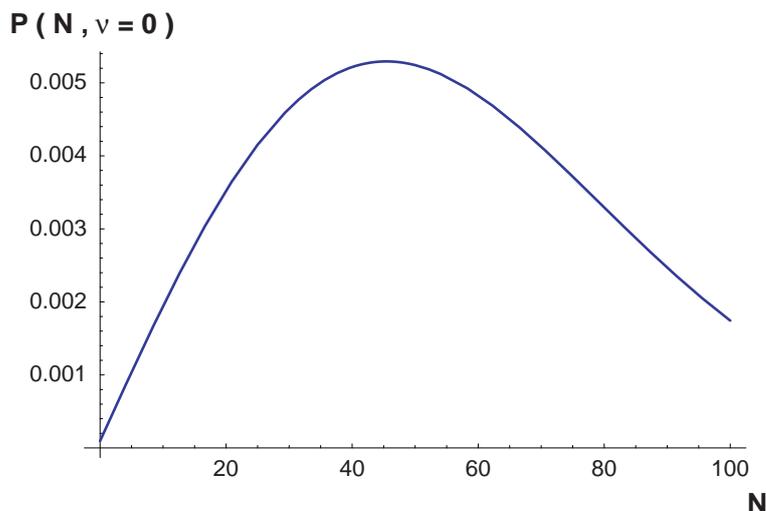


Figure 2.4: Distribution of rotational population in the $\nu = 0$ manifold.

2.2.2 Implementation

The source of the beam is a small cylindrical molybdenum crucible, approximately 10 mm in diameter and 50 mm long. The crucible has a lid with a 0.5×4 mm slot cut into it through which the beam effuses. The crucible sits in a resistive coating-plant heater made of tungsten. The heater / crucible assembly is enveloped in a water cooled heat shield and a mu-metal magnetic shield, to protect the interferometer from deleterious thermal and magnetic effects. The oven is fixed to a large flange on the bottom of the beam machine (figure 2.2). The position of the oven relative to the flange is adjustable from outside the vacuum, facilitating ‘live’ alignment of the molecular beam. Current to the heater is supplied by a variac-fed transformer — typical operating currents are ~ 165 A at ~ 1.7 V.

The crucible cannot be loaded with YbF molecules — YbF is chemically very reactive. Instead the crucible has to be loaded with precursor compounds which react when heated to produce YbF molecules. The chemistry of these reactions is not well understood, indeed alchemy is probably a more appropriate term. Very little is known of the rate constants of the vapour phase reactions that must be taking place, making predictions of precursor efficacy near impossible.

Several chemical mixes have been empirically determined to be effective. Usually a stoichiometric mix of Yb metal, chopped into small hunks (~ 10 mm³), and powdered AlF₃ is used (mass ratio 4:1). The mixture of powdered YbF₃ and Al metal (mass ratio 4.3:1) is also effective. The total oven charge weighs about 3 g. Both of these mixtures produce similar maximum intensities, $Q \simeq 10^{11}$ molecules per second reaching the detector. To produce such a flux, both mixtures have to be heated to approximately 1500 K.

Some limited success has also been achieved with a mixture of YbF_2 and Al metal.

The beam leaving the crucible is highly divergent and must be collimated. Furthermore, the oven's high temperature makes it a bright source of blackbody radiation from which the detector must be shielded (§2.3). To effect this, three baffles with rectangular slots cut into them are arranged along the beam's path. The first baffle the beam encounters is on a movable feedthrough, to provide 'live' alignment. The baffle is a sheet of tantalum foil with a 2×8 mm slit cut into it. The second and third baffles are fixed at opposite ends of the field region, affixed to the inner magnetic shield (described in section 2.8 — see figure 2.14 for the location of the inner magnetic shield). The lower baffle has a 2×8 mm slot and the upper 6×38 mm slot. All three baffles are aligned with their long direction perpendicular to the electric field, as is the slot on the crucible. Roughly speaking, the beam fills the gap between the field plates (details of the field plate geometry can be found in section 2.7).

2.2.3 Operation

The oven places the most significant constraint on when I can run the interferometer. A single oven charge produces a beam of molecules continuously for around 10 hours.⁵ When the oven is exhausted it must first be allowed to cool, this takes around 5 hours. Bringing the chamber up to atmospheric pressure, removing and refilling the crucible, cleaning the lower collimating baffles, and pumping the chamber down to operating pressure takes a further 2–3 hours. The chemicals must then be slowly baked at an intermediate temperature to allow volatile components to evaporate off.⁶ The exit of the oven assembly tends to clog if this is not done for long enough — around 24 hours is good.

On a typical day I heat the oven up over the course of a few hours during the morning, from its baking current of 100 A to the operating current of ~ 165 A.⁷ Then I run the interferometer from early afternoon until late. When the oven charge is nearly exhausted the heater is turned off. The next morning the oven can be recharged, as outlined above. I can then bake the oven, taking the oven current up to 100 A over a few hours. The oven will be ready to use the following day.

Whilst I can reliably produce an intense beam of molecules from the oven, using the above procedure, the beam source is not completely understood. Seemingly identical oven charges and preparation procedures often produce beam intensities that differ by almost

⁵Longer running times are possible at lower oven temperatures, but the intensity is lower.

⁶Consideration of Yb's vapour pressure at the baking temperature shows that it is definitely one of these 'volatile components' — some important chemistry must take place during the baking phase that produces a less volatile YbF precursor.

⁷The oven's temperature cannot be measured directly and there is no straightforward way to relate the heater current to temperature. Nonetheless, heater current provides a reproducible, if uncalibrated, indicator of oven temperature.

a factor of 2. This seems to indicate that the important chemistry that takes place during the baking phase depends sensitively on the parameters of the bake. Some attempt has been made to systematically vary the length and temperature of the baking phase, but no conclusive results have been obtained.

It should also be mentioned in passing that although I speak of an intense beam of molecules, the beam is very dilute in comparison to typical atomic beams. The peak usable intensity (i.e. the rate of molecules in the $X^2\Sigma^+(\nu = 0, N = 0)$ state entering the detector) is $Q \simeq 10^7$ molecules per second (10^{10} mols \cdot (s \cdot Sterad) $^{-1}$). A simple kinetic theory calculation for the number of ground state Tl atoms entering a similar detector, assuming a similar oven slit and collimating geometry, gives a peak intensity of $Q_{Tl} \simeq 10^{13}$ atoms per second (10^{16} atoms \cdot (s \cdot Sterad) $^{-1}$). The difference is largely due to the molecules' rotational excitations.

2.3 Detecting the molecules — laser induced fluorescence

Once the oven is heated up to its operating temperature it's necessary to check that there *really* is a beam of molecules. This is done by laser exciting the molecules and counting their spontaneously emitted photons. This technique, known as laser induced fluorescence (LIF) detection, is the basis of all the measurements that I make with the interferometer — everything finally comes down to looking for changes in the rate of spontaneous emission of photons from the molecules. The LIF detector system, therefore, is a central part of the interferometer. I will describe it in detail below.

2.3.1 Sub-Doppler resolution + many photons = good !

It is clearly essential to the operation of the interferometer that the $|F = 1\rangle$ and $|F = 0\rangle$ states can be individually addressed by the lasers — the **Prepare** and **Probe** phases need this hyperfine resolution.

In a gaseous sample of YbF at 1500K and low pressure, the resolution with which the molecules' transitions could be laser excited⁸ would, of course, be limited by Doppler broadening. The Doppler width of YbF under such conditions is 1 GHz, much larger than the 170 MHz energy difference between the $|F = 1\rangle$ and $|F = 0\rangle$ states.

The situation can be very different in a well collimated molecular beam. In the LIF scheme the molecules in the beam are excited by a laser beam with its propagation direction, \hat{k}_l , perpendicular to the axis of the molecular beam. The broadening due to the first order Doppler shift, which, for the relevant velocities, is linearly dependent on the molecules' velocities along \hat{k}_l , can be greatly reduced.

⁸Considering only straightforward one photon excitation, as employed in the interferometer.

The degree of collimation of the molecular beam can be quantified by the collimation angle (figure 2.5) [30]. The collimation angle, θ_c , is the half angle inside the diverging molecular beam. It can be shown that, for well collimated beams, the first order Doppler width is reduced by a factor of $1/\theta_c$. A rough geometrical estimate of the collimation angle of the YbF beam gives $\theta_c \simeq 10^{-2}$. The corresponding reduction in the Doppler width to ~ 10 MHz permits resolution of the $|F = 1\rangle$ and $|F = 0\rangle$ states.

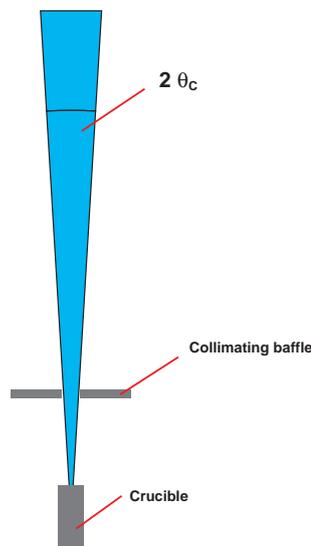


Figure 2.5: Quantifying the degree of collimation of a molecular beam.

It should be noted that there is a significant difference between LIF in molecules and atoms. Typical experiments with atoms drive a transition within a closed system; meaning, after excitation the atom decays back into the lower state of the transition with a rate that is fast compared to the interaction time with the exciting laser. The number of spontaneous photons that can be ‘extracted’ from each atom is limited by the effective decay rate and the interaction time. It is usually impossible to find such transitions in molecules; there are almost always decay modes from the transition’s upper state into metastable states, other ro-vibrational states, that do not interact with the exciting laser. The number of spontaneous photons per molecule is limited by the branching ratio into these ‘lost’ states. The branching ratios for transitions in YbF are not well known, but experiments⁹ suggest that each YbF molecule emits 5–10 photons before being ‘lost’.

2.3.2 Implementation

A plan view of the probe region is shown in figure 2.6 (cf. figure 2.2).

⁹A short pulse of intense laser light was used to drive the Q(0) F=1 transition. The fluorescence signal as a function of pulse length contains information about the branching ratios.

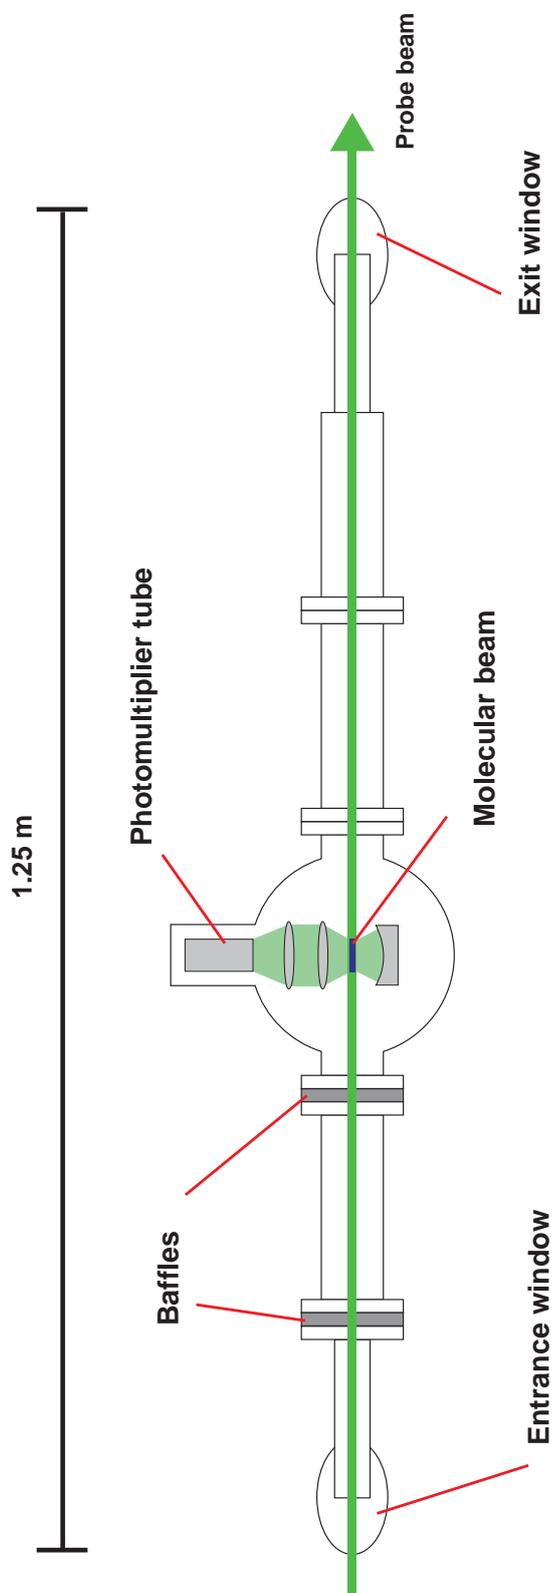


Figure 2.6: A plan view of the probe region. The molecular beam travels vertically, out of the page.

The probe laser beam is transported to the vacuum chamber through a single-mode polarization-preserving optical fibre. After leaving the fibre, the beam passes through beam shaping and cleaning optics, resulting in a spatially well-defined, collimated, elliptically profiled beam of cross-section 2×10 mm. In figure 2.6 the long axis of the ellipse is in the plane of the page.

The probe beam enters the vacuum chamber through a glass window, set at Brewster's angle to minimise reflection. A sidearm, length 40 cm, separates the entrance window (which scatters some light) from the region of intersection with the molecular beam. The sidearm contains two baffles, made from soot-blackened razor blades, through which the beam passes. The holes in the baffles are approximately 2×10 mm. The sidearm and baffles reduce the amount of stray light entering the photomultiplier tube (PMT), both from room lights and from laser light scattered by the window.

The region of intersection between the probe beam and molecular beam sits at the centre of curvature of a spherical mirror. This mirror reflects fluorescence back toward the imaging optics. The imaging optics, 2 lenses, focus the fluorescence from the intersection region onto the face of the photomultiplier tube. The light collection efficiency, due to the solid angle subtended by the imaging optics, is approximately 10^{-2} .

The spontaneously emitted light is of a similar wavelength to that of the exciting laser and at this wavelength the PMT's photocathode efficiency is $\sim 10^{-1}$. In combination with the solid angle factor, this gives an overall efficiency for the detector of $\sim 10^{-3}$.

After the intersection region the probe beam travels down a 60 cm long sidearm before leaving the chamber through another glass window tilted to Brewster's angle. The beam then travels several metres across the lab before being absorbed by a beam dump. Again, the purpose of the sidearm and the distant beam dump are to reduce the amount of light scattered into the PMT.

The PMT is operated in pulse counting mode, to count individual photons. The output from the PMT is shaped into a TTL pulse train and sent down an optical fibre to the computer. The fibre is connected to one of the optical receivers in the 'breakout box'. The computer uses two of the counters on its Lab-PC+ board to measure the pulse rate.

A Labview program, `Laser scan`, can generate a voltage ramp at one of the Lab-PC+'s analog outputs. This analog output can be used to scan either of the lasers. After each small increment in the output voltage `Laser scan` records the PMT count rate.¹⁰

¹⁰Also recorded are the outputs from photodiodes monitoring the transmission through two reference Fabry-Perot cavities. I use the cavities to calibrate and linearise scans when I want to do precise spectroscopy.

2.3.3 Operation

The LIF detector is both simple and reliable in operation — as would be hoped for something that the rest of the interferometer depends upon. Getting it to work is a simple matter of turning it on, and adjusting the alignment of the probe beam relative to the chamber in order to minimise the scattered light. The baffle geometry ensures that the probe beam intersects the molecular beam.

Under typical running conditions the photon count rate from the detector is as follows:

Dark rate The PMT's dark count rate is very low, effectively negligible.

Black-body rate When the oven is at its operating temperature, it is responsible for a count rate of 20–100 kHz — known as the black-body rate. Although it's always the case that, on any given day, the hotter the oven, the larger the black-body rate, the count rate from one day to the next is frequently very different for seemingly identical oven conditions. Often the black-body rate is higher in the morning, just after the oven has been heated to operating temperature, than in the afternoon. I suspect that bright spots form on the face of the crucible, perhaps due to chemicals erupting from the oven, and increase the count rate temporarily.¹¹

Room light rate The interferometer is run under subdued room lighting. Blackened walls and the long probe beam sidearms also help reduce the amount of stray light entering the probe region. The room light count rate is usually <1 kHz.

Laser scatter rate YbF molecules aren't alone in scattering photons from the laser beam into the PMT. The chamber windows also scatter light. The count rate due to this scattering is known as the laser scatter rate. This rate is dependent on the probe beam intensity, indeed the relative variation of this rate and the signal rate with intensity determines the optimum intensity. At optimum intensity the laser scatter rate is ~40 kHz.

Molecular fluorescence rates The spectrum of the molecular fluorescence near the Q(0) transition is quite crowded — fluorescence is observed from many transitions. At the peaks of the strongest of these transitions the count rate can be up to 200 kHz. When the laser is locked to the Q(0) F=0 transition other molecular transitions usually contribute ~ 15 kHz to the count rate.

A plot of the detected fluorescence as a function of the laser excitation frequency is shown in figure 2.7. The laser frequency was scanned over a few GHz around the Q(0)

¹¹This is just a suspicion though. Yet another thing that contributes to the feeling that I really don't understand what's going on in the oven ...

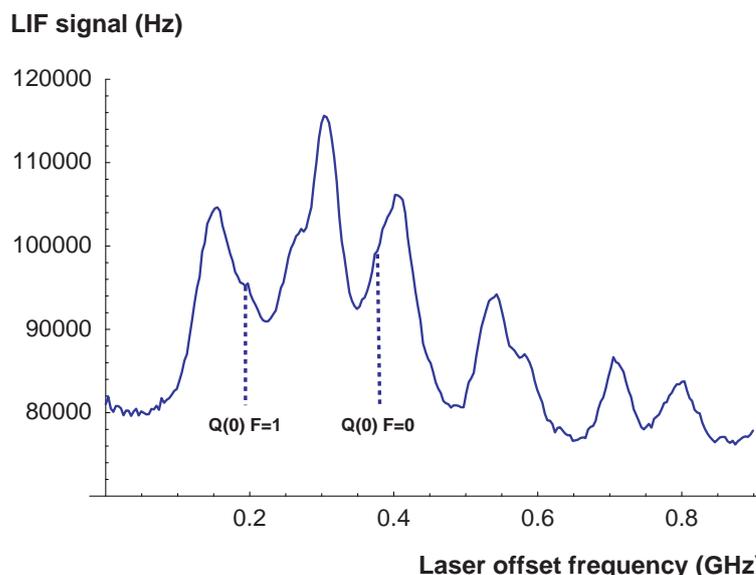


Figure 2.7: An LIF scan around the $Q(0)$ transitions.

transition. Although the two peaks corresponding to $Q(0)$ $F=0$ and $Q(0)$ $F=1$ are not resolved, their intensity can be inferred from the intensity of the other lines. In this scan the intensities of the $Q(0)$ $F=0$ and $Q(0)$ $F=1$ transitions are around 8 kHz .

With the LIF detector I am able to monitor the fluorescence rate in real time (using another Labview program, `Monitor`). This proves to be a very useful diagnostic for getting the interferometer running: the molecular beam's condition can be assessed; the alignment of the molecular and laser beams can be optimised; and the laser frequency can be quickly adjusted using the YbF spectrum as a reference.

It should be mentioned that this LIF detection scheme has also been used to carry out a detailed spectroscopic study of the YbF molecule — details can be found in [26], [27] and [25].

2.4 Pump and probe beams — locking the first laser

Once I'm happy that the molecular beam is operating correctly it's time to lock the first laser to the I_2 spectrometer.

The first laser is responsible for generating the pump and probe beams. If the interferometer is to be run continuously for many hours, as is necessary to measure the EDM, these beams must stay resonant with the transitions they drive. To this end, the first laser's frequency is locked to a resonance line in I_2 . The spectrometer and associated locking electronics that make this possible are the main subject of this section.

The frequencies of the pump and probe beams differ by 170 MHz, the hyperfine interval in the $X^2\Sigma^+(\nu = 0, N = 0)$ state. This frequency offset is produced by an acousto-optic modulator (AOM), which is described below. Usually the first laser is locked to the $Q(0) F=0$ transition (probe) and the AOM is used to generate light for the $Q(0) F=1$ transition (pump).

2.4.1 The I_2 spectrometer

The first laser needs to be locked to within a few MHz of the centre of the $Q(0) F=0$ transition. Correspondingly, if the interferometer is to be run for 10 hours, a frequency stability of better than around $0.5 \text{ MHz}\cdot\text{hr}^{-1}$ is required.

A saturated absorption scheme is employed in the I_2 spectrometer [30]. This scheme has two salutary features :

- The scheme is nearly Doppler-free — features of width $\sim 20 \text{ MHz}$ can be resolved. This considerably eases the technical challenges that would be associated with locking to a Doppler broadened feature (width $\sim 400 \text{ MHz}$).
- A frequency modulation technique, described below, gives the spectrometer lines dispersion type lineshapes. Locking to a zero-crossing point is advantageous.¹²

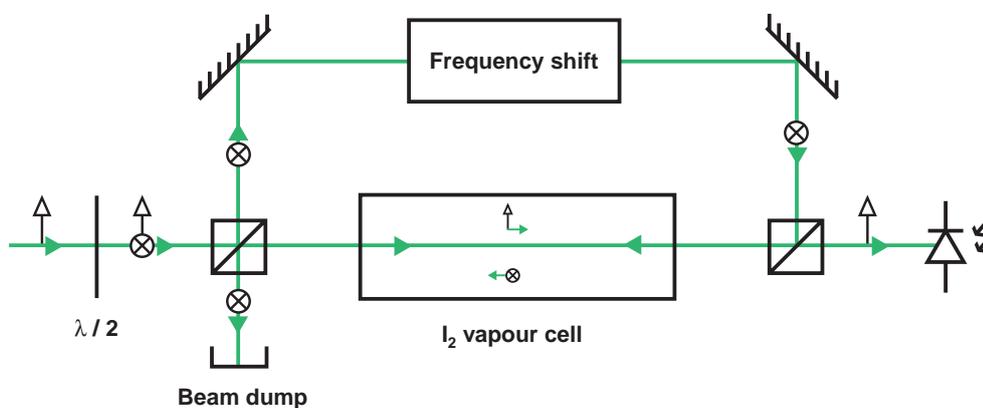


Figure 2.8: Layout of the I_2 spectrometer's optical components showing the laser polarisation in each section.

The optical layout of the spectrometer differs from the standard shallow crossing angle pump-probe layout [30]. The layout of the optical components is shown in figure 2.8. In the figure the black and white arrows indicate the polarisation of the laser beam — the beam with two arrows can be adjusted to have any linear polarisation.

¹²Zero doesn't drift with time !

The spectrometer is fed light picked off from the first laser — about 150 mW. This horizontally polarized light is incident on a $\lambda/2$ plate whose orientation can be adjusted, giving light of arbitrary linear polarisation. The beam is then split by a polarising beam splitter (beam cube). The horizontally polarized component passes through the beam cube unhindered. This beam then passes through the I_2 cell and a second beam cube. Finally the intensity of this beam is measured by a reverse biased photodiode. This horizontally polarized beam is the probe beam. Meanwhile, the vertically polarized component that was reflected from the first beam cube is taken to the frequency shifter, an AOM used in a double-pass configuration [22], which shifts the laser frequency by ~ 254 MHz. This shifted light is sent to the second beam cube such that after reflection it fully overlaps the probe beam, which is travelling in the opposite direction. This is the pump beam.

This optical arrangement has two advantages over the conventional shallow crossing scheme. Firstly, the beams can be made to overlap completely, including more molecules in the interaction. Secondly, the $\lambda/2$ plate allows the ratio of pump to probe intensity to be varied continuously — this ratio can be adjusted to give maximum signal:noise (the optimum ratio is around 25:1 pump:probe).

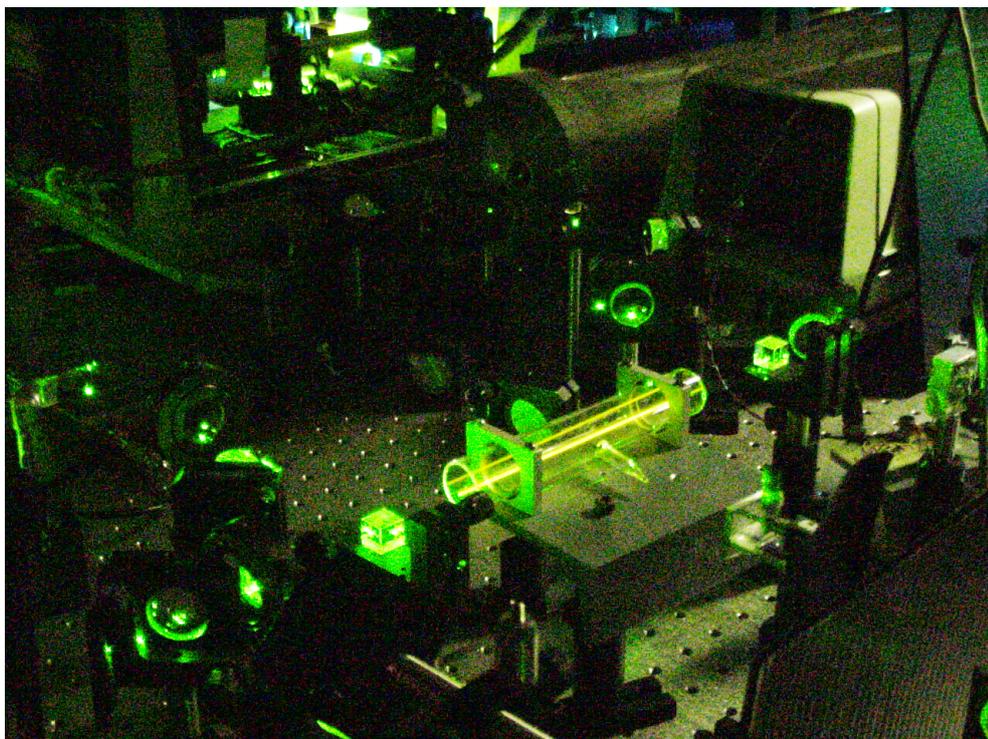


Figure 2.9: A picture of the I_2 spectrometer.

The introduction of the frequency shifter in the pump arm of the spectrometer allows the output spectrum to be ‘tuned’ in frequency space. Standard saturated absorption spec-

troscopy makes use of the fact that stationary molecules will interact strongly with both the pump and probe beams when the laser frequency is resonant with a molecular transition. By using a pump beam shifted in frequency by Δf , the spectrometer selects a different velocity class to interact with both beams, that with Doppler shift $\frac{\Delta f}{2}$ along the beam direction.¹³ This shifts the output of the spectrometer by $\frac{\Delta f}{2}$. Modulating the frequency shift at an intermediate frequency (100 kHz) facilitates phase sensitive detection and yields dispersion type lineshapes.

The control and locking electronics are shown schematically in figure 2.10. The frequency shifter is fed rf from a voltage controlled oscillator¹⁴, running at around 127 MHz¹⁵. The control voltage, and hence frequency, of this oscillator is modulated by a second oscillator running at 100 kHz. The modulation depth is 12 MHz. The phase of the modulation is continuously adjustable and discrete π or $\pi/2$ phase shifts can also be switched in. This phase is adjusted to maximise the output of the spectrometer. After amplification and low-pass filtering, the component of the spectrometer's photodiode output at the modulation frequency is extracted with a mixer. This output signal is filtered and sent to the computer (line ISPEC in figure 2.1). The output is also sent to an integrator, with time constant 2 ms. The output of the integrator is the laser feedback signal (line L1L in figure 2.1).

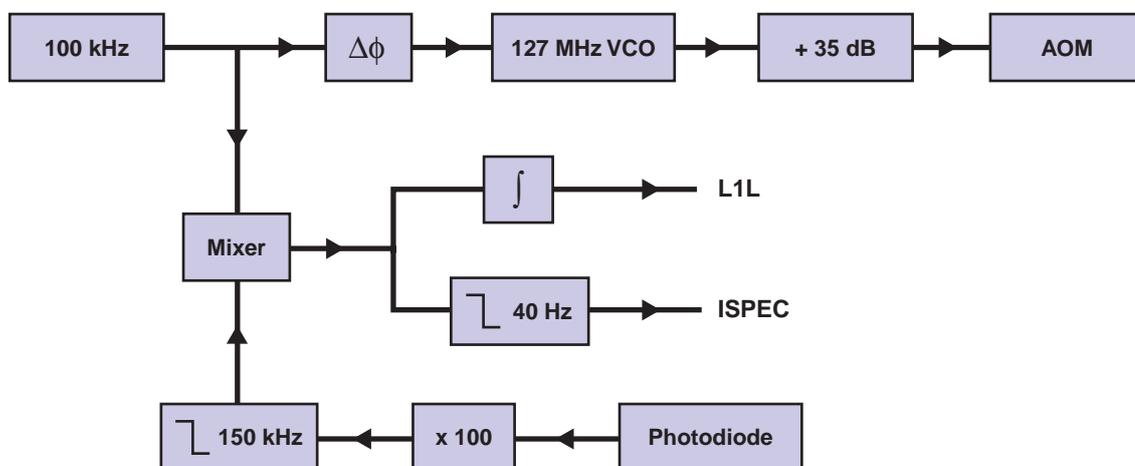


Figure 2.10: The I_2 spectrometer's control electronics.

Figure 2.11 shows the output from the spectrometer, recorded by the Laser scan

¹³This means the frequency shift, $\frac{\Delta f}{2}$, must be less than the Doppler width of I_2 .

¹⁴Some care was taken to ensure that the VCO's output power does not depend on frequency and that the output frequency is linearly related to the control voltage, over the relevant frequency range. This is necessary to get symmetric lines from the spectrometer.

¹⁵As noted above, the AOM is used in a double-pass configuration, so the laser frequency is shifted by twice the rf frequency. The output of the spectrometer is shifted by half of the laser frequency shift — this is just the rf frequency.

program, as the laser frequency was scanned. The LIF detector signal is shown for reference. The scan range is approximately 1 GHz. Notice that the zero-crossing point of a strong feature is aligned with the $Q(0) F=0$ transition — choosing the frequency shift of $\frac{\Delta f}{2} = 127$ MHz ensures this.

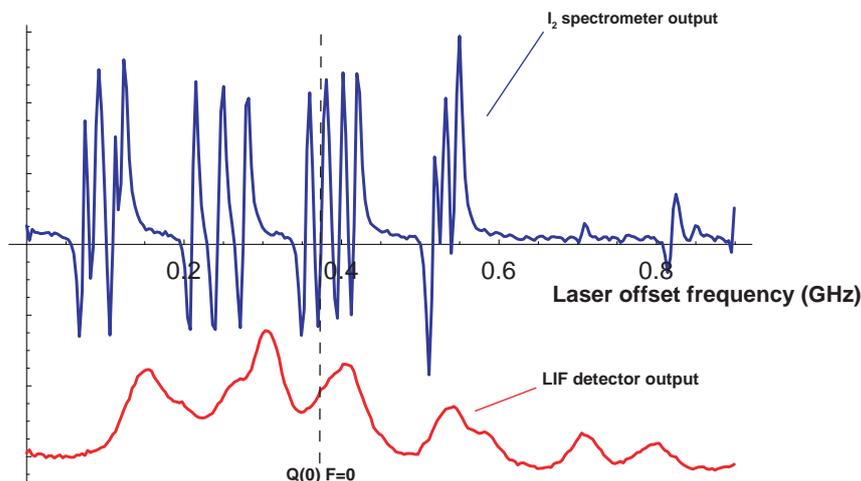


Figure 2.11: A scan over some I_2 spectrometer lines. The zero-crossing that coincides with the $Q(0) F=0$ transition is marked.

Locking the first laser is reasonably straightforward. First thing each day, some tweaking of the spectrometer alignment is usually needed to maximise the output signal. Then I can simply scan the laser by hand to the correct frequency, monitoring the LIF and spectrometer signals with `MONITOR`, and throw the lock switch. If the laser is aligned and operating correctly the laser can stay locked for up to 10 hours. The lock is sufficiently robust that it takes a fair jolt to the optical table to make it unlock.

2.4.2 Pump beam

The pump beam, resonant with the $Q(0) F=1$ transition, is derived from the first laser's output. About 20 mW of light is picked off from the laser output. This light is sent through an AOM, again in the double-pass configuration, which shifts the frequency by 85 MHz each pass. The efficiency of the AOM is about 50% for two passes, giving ~ 10 mW of light 170 MHz red detuned from the first laser frequency.

The pump beam passes through a solenoid actuated shutter — this shutter is primarily for convenience, as it saves having to lean into the optical table¹⁶, but also allows the computer to block and unblock the pump beam, although this feature is rarely used. The

¹⁶A procedure notorious for unlocking lasers !

beam enters and leaves the vacuum chamber through glass windows set a few centimeters below the bottom of the field region (see figure 2.2 and section 2.7).

2.5 Radio frequency transitions

Two phases of the interferometer, the **Pump** and **Probe** phases, are working. Now it's time to make the **Split** and **Recombine** phases work.

The **Split** and **Recombine** phases are both effected by driving the same transition in the molecule — an rf oscillating magnetic field drives an M1 transition. The transition is between the $|F = 1\rangle$ and $|F = 0\rangle$ levels of the $X^2\Sigma^+(\nu = 0, N = 0)$ state (from here on in just referred to as the ground state). The theory and implementation of these rf transitions will be described below.

In order to observe narrow rf resonances, stray magnetic fields must be excluded from the region between the splitter and recombiner. Passive magnetic shielding is employed to achieve this. The implementation and performance of the magnetic shields will also be described in this section.

2.5.1 Theory

The theory of radio frequency transitions, of the type described here, has been well understood since the middle of the last century, growing from the work of Bloch, Rabi and Ramsey (see [31] and [32]). In this light, the presentation below may seem unnecessary. Its value lies in establishing the framework within which the operation of the interferometer will be described (§2.8).

Thus far I have described the hyperfine levels of the ground state in only one basis, which I call the z-basis, aligned with the electric field direction.¹⁷ States in the z-basis can be represented as 4-component column vectors, with the components representing the amplitudes of the states $|0, 0\rangle$, $|1, 1\rangle$, $|1, 0\rangle$ and $|1, -1\rangle$.

When considering the action of the rf field, it is useful to use a second basis. This basis, the Cartesian basis, is made up of states $|0, 0\rangle$, $|y\rangle$, $|1, 0\rangle$ and $|x\rangle$. The basis vectors $|x\rangle$ and $|y\rangle$ are the $F = 1, m_F = 0$ states of the x- and y-bases. Explicitly, the $|x\rangle$ and $|y\rangle$ states are given in terms of z-basis states by,

$$|y\rangle = \frac{1}{\sqrt{2}} (|1, 1\rangle + |1, -1\rangle) , \quad |x\rangle = \frac{1}{\sqrt{2}} (|1, 1\rangle - |1, -1\rangle) .$$

Vector representations in this basis, written with a subscript c , are related to those in the

¹⁷In the electric and magnetic fields typically used the Stark interaction is much larger than the Zeeman interaction — this justifies the choice of basis.

z-basis by the self-inverse transformation,

$$\mathbb{CB} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & \frac{1}{\sqrt{2}} & 0 & \frac{1}{\sqrt{2}} \\ 0 & 0 & 1 & 0 \\ 0 & \frac{1}{\sqrt{2}} & 0 & -\frac{1}{\sqrt{2}} \end{pmatrix}. \quad (2.2)$$

The rf transitions are driven by an oscillating magnetic field. The magnetic field is linearly polarized along the direction of the molecular beam — this direction is labelled y . Being a linearly polarised field it drives a $\Delta m = 0$ transition between the $|F = 0\rangle$ and $|F = 1\rangle$ states, when viewed in the y -basis. The action of the rf magnetic field is simply described in the Cartesian basis — it is a coupling between the $|0, 0\rangle$ and $|y\rangle$ states. The interaction Hamiltonian is,

$$\hat{H}_{int} = -\vec{\mu} \cdot \vec{B} \cos(\omega t + \phi).$$

Making the rotating wave approximation, this Hamiltonian is written, *in the Cartesian basis* as,

$$\hat{H}_{int} = \begin{pmatrix} 0 & \hbar b e^{i(\omega t + \phi)} & 0 & 0 \\ \hbar b e^{-i(\omega t + \phi)} & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}_c,$$

where b , the Rabi frequency, is a measure of the interaction strength given by

$$b = -\frac{B}{2\hbar} \langle 0, 0 | \mu_y | y \rangle = \frac{\mu_B B}{2\hbar}.$$

By integrating this Hamiltonian we can write a matrix that describes the time evolution of the Cartesian vector representation of an arbitrary initial state in the rf magnetic field¹⁸,

$$\mathbb{RF} = \begin{pmatrix} e^{\frac{i\omega\tau}{2}} \left(i\theta_c \sin\left(\frac{a\tau}{2}\right) + \cos\left(\frac{a\tau}{2}\right) \right) & i\theta_s e^{\frac{i\omega\tau}{2}} e^{i(\omega t + \phi)} \sin\left(\frac{a\tau}{2}\right) & 0 & 0 \\ i\theta_s e^{-\frac{i\omega\tau}{2}} e^{-i(\omega t + \phi)} \sin\left(\frac{a\tau}{2}\right) & e^{-\frac{i\omega\tau}{2}} \left(-i\theta_c \sin\left(\frac{a\tau}{2}\right) + \cos\left(\frac{a\tau}{2}\right) \right) & 0 & 0 \\ 0 & 0 & e^{-\frac{i\Omega\tau}{2}} & 0 \\ 0 & 0 & 0 & e^{-\frac{i\Omega\tau}{2}} \end{pmatrix}_c, \quad (2.3)$$

where ω is the rf frequency, Ω is the zero-field transition frequency¹⁹, t is the time at which the interaction starts, τ is the duration of the interaction and

$$a = \sqrt{\delta^2 + 4b^2}, \quad \delta = \Omega - \omega, \quad \theta_c = \frac{\delta}{a}, \quad \theta_s = -\frac{2b}{a}.$$

¹⁸In deriving this expression, several important assumptions have been made. The interaction is supposed to turn on and off suddenly, and have constant strength during the interaction period. Incoherent processes have not been included.

¹⁹The zero of energy has been chosen to lie half way between the $|F = 1\rangle$ and $|F = 0\rangle$ zero-field energies.

Now we're in a position to analyse a 'pump-probe' experiment. The molecular beam first interacts with the pump beam, leaving the molecules in the state $|0, 0\rangle$ — just as in the **Prepare** phase of the interferometer. It then interacts with a single rf magnetic field. Finally the population remaining in the $|0, 0\rangle$ state is measured — the **Probe** phase of the interferometer. Pump-probe experiments are a valuable diagnostic for both the rf system and the **Prepare** phase — I routinely carry out such experiments. Some results from pump-probe experiments are presented in section 2.5.3.

Assuming the **Prepare** phase is ideal, it's straightforward to show, using (2.3), that a molecule with velocity v will have a probability

$$P_0 = \cos^2 \left(a \frac{l_{rf}}{v} \right) + \theta_c^2 \sin^2 \left(a \frac{l_{rf}}{v} \right) ,$$

of being found in the $|0, 0\rangle$ state before the **Probe** phase, where l_{rf} is the length of the rf interaction region. If the rf field is resonant with the transition frequency ($\theta_c = 0$), then for molecules of a given velocity, v , it is possible to adjust the effective Rabi frequency, a , to achieve complete population transfer into the $|y\rangle$ state. The condition that must be satisfied is,

$$a \frac{l_{rf}}{v} = (n + 1/2) \pi , \quad (2.4)$$

where n is an integer.

In a molecular beam, where the molecules have a velocity distribution, it is not possible to satisfy this condition simultaneously for all of them. The question arises, "What is the optimum effective Rabi frequency for a given velocity distribution, such as that belonging to the YbF beam?". Averaging over the velocity distribution (2.1), the probability of a molecule in the beam being found in the $|0, 0\rangle$ state before the **Probe** phase is given formally by

$$\bar{P}_0 = \int_0^\infty \frac{2}{\alpha^4} v^3 \exp \left(\frac{-v^2}{\alpha^2} \right) \cos^2 \left(a \frac{l_{rf}}{v} \right) dv ,$$

where I have assumed that the rf field is on resonance for all molecules. This integral is readily evaluated numerically. Figure 2.12 shows the population transferred to the $|y\rangle$ state, $1 - \bar{P}_0$, vs. the Rabi frequency a . The a -axis is scaled to the useful quantity

$$a_0 = \frac{1}{2} \frac{v_p}{l_{rf}} \pi ,$$

that is, the Rabi frequency such that molecules with the most probable velocity are completely transferred to the $|y\rangle$ state (a π -pulse). Points to note from the figure are: the Rabi oscillations are quickly damped, due to the rapid dephasing of molecules with dif-

ferent velocities; the maximum transfer occurs at a Rabi frequency very close to a_0 ²⁰; the maximum possible population transfer is 76%.

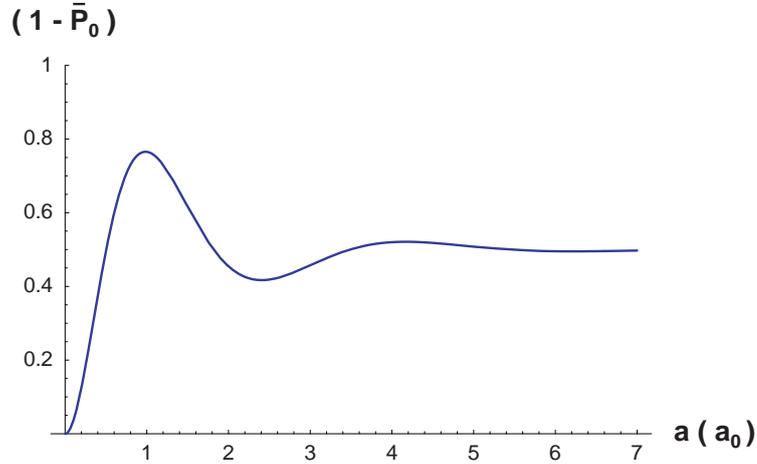


Figure 2.12: Population transfer vs. Rabi frequency.

2.5.2 Implementation

rf systems

There are two independent rf generation systems, one for the interferometer’s splitter transition and one for its recombiner transition.

The ‘red’²¹ system is usually used to drive the recombiner transition. The heart of this system is an HP3325B frequency synthesizer. The output frequency of the synthesizer can be controlled by the computer through a GPIB interface. The synthesizer’s maximum output frequency is 61 MHz — in order to reach the hyperfine interval frequency of 170 MHz the output from this synthesizer is frequency quadrupled. After being appropriately amplified and filtered, output from this synthesizer is fed to the recombiner rf antenna.

The other rf system, known as ‘green’, is somewhat simpler. The HP8657A synthesizer used in this system has a higher maximum output frequency, 1040 MHz. This synthesizer is used to generate the 170 MHz rf directly. The frequency and amplitude of this synthesizer are under computer control, again through GPIB. The output from this synthesizer is simply amplified and fed to the splitter antenna.

Two Labview programs assist in setting up and characterising the rf systems. The first, `function generator scan`, steps the frequency of either the red or green synthe-

²⁰It will be shown in section 2.8.1 that this is also very close to the condition for an optimal interference signal.

²¹So called because of the colour of its front panel LED display.

sizer and records the LIF detector signal at each increment. `Function generator scan` is useful for measuring the rf transition centre frequencies and lineshapes. The second program, `point by point`, repeatedly measures the difference in LIF signal as the rf frequency of one of the synthesizers is stepped back and forth between two values (usually on- and off-resonance). `Point by point` calculates statistics from the data accrued which provide an effective means to measure the rf transition intensities quickly — more quickly than using `function generator scan`.

The output from each rf generator is fed to an rf antenna. The antennae are cylindrical copper loops, diameter 8 cm, thickness 1 mm, length 1 cm. They are held at either end of the field region, inside the magnetic shielding. Their position relative to the vacuum chamber is shown in figure 2.13. The distance between the rf antennae is 62.5 cm.

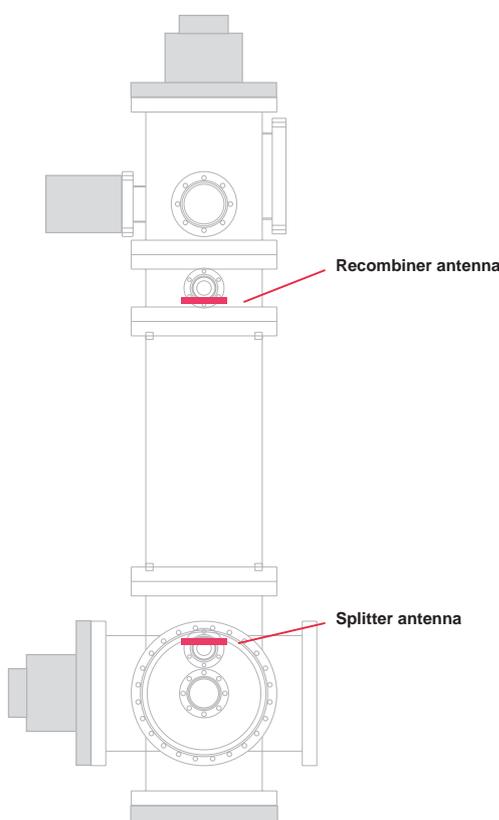


Figure 2.13: The position of the rf antennae.

Magnetic shielding

The magnetic shielding is absolutely essential. In order to see unbroadened rf resonances the ambient magnetic field in the interaction region must be less than $10 \mu\text{T}$. Furthermore, as will be shown in section 2.8.1, for the proper operation of the interferometer the average magnetic field over the molecules path between the splitter and recombiner must be less

than 20 nT. Finally, it should be noted that the EDM interaction energy is equivalent to a magnetic interaction with at most a few pT. Any noise in the ambient magnetic field must be taken very seriously as it can be a large source of noise in the EDM measurement [9].²²

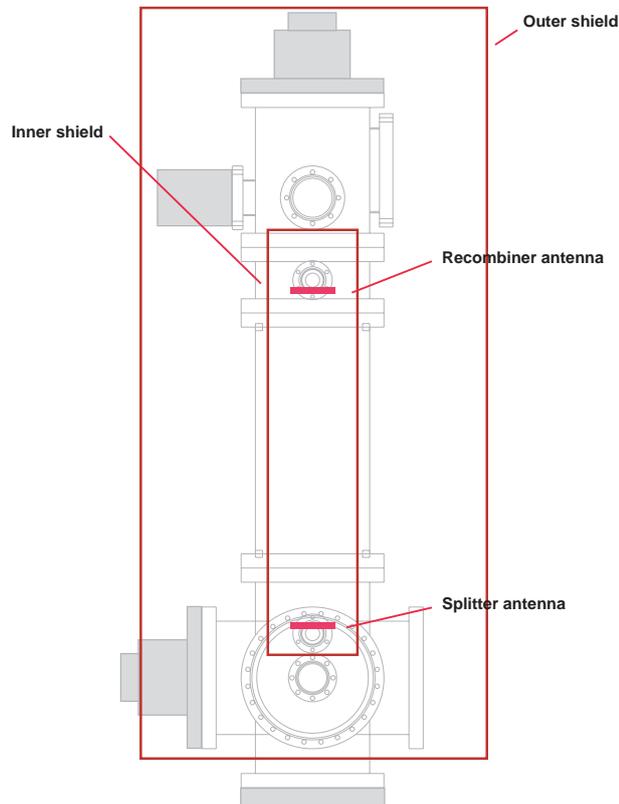


Figure 2.14: The position of the magnetic shields.

The YbF beam is shielded from magnetic fields by two layers of mu-metal. The inner shield is a cylindrical can, diameter 17 cm, length 79 cm, which sits inside the vacuum chamber. The field region is contained within this can — the pump and probe beams are outside the shield. The inner shield has small holes cut into it for electrical connections and laser access²³. The outer shield is also cylindrical, diameter 64 cm, length 140 cm, and sits around the vacuum chamber. Holes are cut into the outer shield to accommodate the lasers ports, turbo pump etc. The position of the shields relative to the vacuum chamber and the rf antennae is shown in figure 2.14.

Copper wires are wrapped around both the inner and outer shields — alternating current can be passed through these wires to degauss the shields.

²²The sensitivity to magnetic fields depends strongly on their direction — this has not been taken into account here.

²³Unused in the current configuration.

The shielding factor has been calculated to be around 1000. The ambient field inside the inner shield is typically < 5 nT immediately after degaussing the shields.

2.5.3 Operation

Before attempting to drive any rf transitions I usually degauss the magnetic shields. Although this isn't necessary, it *does* have to be done before any interference experiments are attempted, and it only takes a few minutes — doing it before checking the rf system always seems convenient.

The day to day running of the rf system is trouble-free — it just needs switching on. Below I present three sets of measurements that were carried out to characterise the rf system. First, the power dependence of the rf transition intensities is shown. Second, the lineshapes of the rf transitions. Finally, the pressure dependence of the rf transition intensities.

Power dependence of rf transition intensities

These measurements only have to be performed after some part of the rf system has been changed. The method is essentially that of a pump-probe experiment. `point` by `point` was used to measure the difference in LIF signal with the rf on- and off-resonance. This was repeated for several rf drive powers.

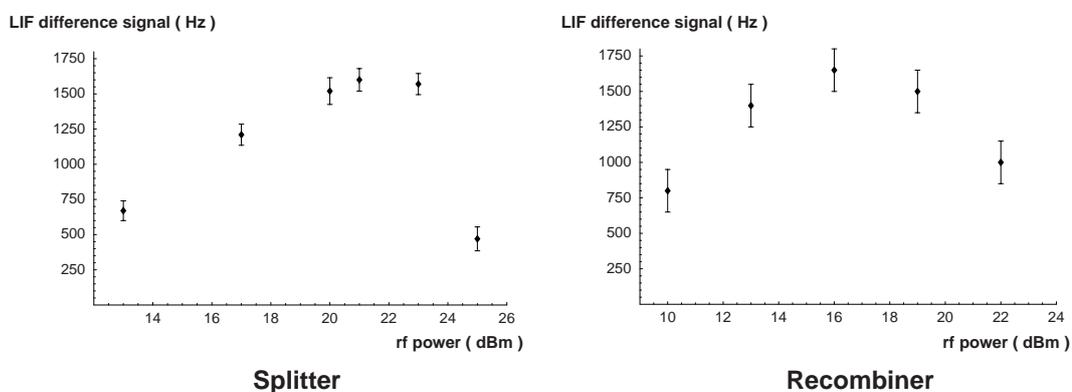


Figure 2.15: rf transition intensity vs. rf drive power.

Plotted in figure 2.15 is the forward power to the antennae, measured in dBm, vs. the LIF difference signal. It is straightforward to extract the optimum rf power for each loop from figure 2.15. Remembering that dBm are logarithmic, it can be seen that these curves have the gross form of figure 2.12. One feature to note is that the splitter displays very poor transfer efficiency for high powers — much worse than would be expected from the theory presented above. This effect remains unexplained.

Lineshapes of the rf transitions

Again the method is that of pump-probe. For these measurements `function generator scan` was used to measure the LIF signal over a range of rf frequencies. The rf is generated at the optimal power levels determined above. These measurements are also performed infrequently — however, as will be detailed in section 2.7, similar measurements are regularly used to characterise the electric field.

Figure 2.16 shows the results. The points each represent 1.1 s of integration. The blue Lorentzian curves have been fitted to the points. The linewidths of the splitter and recombiner transitions are 3.0 kHz and 4.5 kHz respectively. Transit-time broadening, due to the finite interaction time, likely accounts for some of this linewidth, with power-broadening making up the rest.

It can be noted from figure 2.16 that the splitter is significantly less effective than the recombiner. Measuring the rf transition efficiencies is not easy, but a crude estimate can be made by comparing the intensity of the rf line with the intensity of an LIF spectrum taken under the same conditions. An estimate made this way of the *combined* efficiency of the pump laser and the rf transition gives $\sim 9\%$ for the splitter and $\sim 12\%$ for the recombiner; these efficiencies are much lower than the theoretical maximum (76%).²⁴ It would be worthwhile to devise a method that could accurately measure the transition efficiencies, as there may be the possibility for a significant increase in the interferometer's signal.²⁵

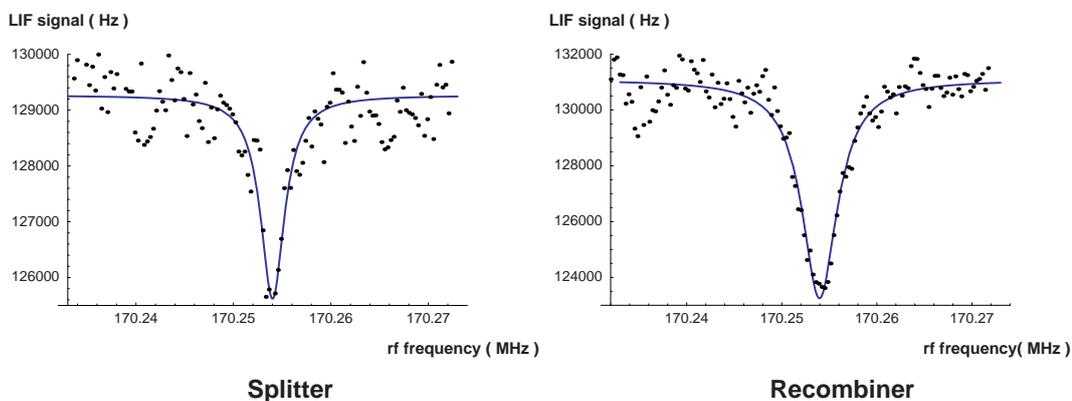


Figure 2.16: The rf transition lineshapes.

²⁴I don't believe that the rf transition efficiencies are really this poor; I think it likely that the measurement is flawed. The procedure is reliant on a theoretical model of the LIF spectrum which relates the relative intensities of different transitions. This model is probably the source of the error.

²⁵A possible method would be to switch both the splitter and recombiner on- and off-resonance, measuring the signal in all four combinations. It may be possible to extract the absolute transition efficiencies from these data.

Pressure dependence of rf transition intensities

These measurements justify the comment made in section 2.1.1, that the pressure in the vacuum chamber is low enough that the rate of collisions with background gas is negligible.

Again, a pump-probe experiment was carried out. The background gas pressure in the chamber was rather crudely varied by slowly spinning down the turbo pumps. Point by point was used to measure the intensity of the recombiner rf resonance at several different pressures. The data are presented in figure 2.17.

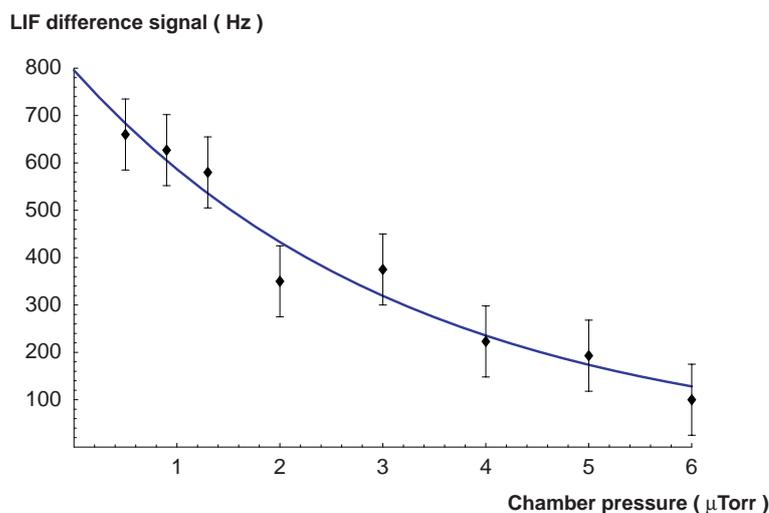


Figure 2.17: rf transition intensity vs. chamber pressure.

The data can be interpreted with a simple model. Consider a population difference between the $|F = 0\rangle$ and $|F = 1\rangle$ states produced by the pump beam. Collisions reduce this population difference as the molecules travel to the recombiner antenna. The rf transition intensity will be proportional to the population difference arriving at the antenna. Making some assumptions about everything being uniform, the pressure dependence of the rf transition intensity is

$$N = N_0 e^{-\alpha p},$$

where N_0 is the rf transition intensity in the absence of collisions, p is the pressure and α is proportional to the collision cross section. This model was fitted to the data, and is shown as the blue curve in figure 2.17. The interesting parameter $\alpha = 0.30 \mu\text{Torr}^{-1}$. This value implies that at the operating background pressure of $0.4 \mu\text{Torr}$ the interference contrast is reduced by less than 5% by hyperfine-state changing collisions — this is considered negligible. The collision cross section can be easily extracted from α — the result, $\sigma = 1.3 \cdot 10^{-14} \text{ cm}^{-2}$, seems reasonable.

2.6 Repump — locking the second laser

Now we've almost got a working interferometer — only the **E & B** phase (see figure 1.4) is yet to be implemented. Before moving on to describe that, I will return to the **Pump** phase and describe a significant enhancement that can be made.²⁶

2.6.1 Repump transition

The **Pump** phase, as described so far, is the interaction of the molecules with a laser beam resonant with the Q(0) F=1 transition. As described previously, this laser beam, the pump beam, creates a population imbalance between the ground state hyperfine levels, favouring $|F = 0\rangle$. Molecules initially in the $|F = 1\rangle$ state are pumped to either the $|F = 0\rangle$ state or leak out to other molecular states. The **Pump** phase is improved by plugging the largest of these leaks.

The leak is plugged by letting the molecules interact with a second laser beam during the **Pump** phase. Consideration of angular momentum conservation and parity show that the “other molecular states” mentioned above must be either the $X^2\Sigma^+(\nu = 0, N = 2)$ state or states that are vibrationally excited. Little is known about the Frank-Condon factors for YbF, but it is assumed that the dominant decay mode from an excited state is to a state with the same vibrational quantum number.²⁷ If this is the case, then the majority of the molecules pumped from the $|F = 1\rangle$ state that don't decay to the $|F = 0\rangle$ state will end up in the $X^2\Sigma^+(\nu = 0, N = 2)$ state. Pumping molecules back out of the $X^2\Sigma^+(\nu = 0, N = 2)$ state will plug the biggest leak.

The transition of interest from the $X^2\Sigma^+(\nu = 0, N = 2)$ state to the $A^2\Pi_{1/2}(\nu = 0, N = 0)$ state, the ${}^{\circ}P_{12}(2)$ transition, is 40 GHz red detuned from the Q(0) transition. This detuning necessitates the use of two lasers, being a little too far to conveniently add a sideband to the first laser. The second dye laser is frequency stabilised by locking it to a Fabry-Perot cavity that is itself stabilised with respect to the first laser.

2.6.2 Cavity lock

The optical layout of the cavity locking system is shown in figure 2.18. As in figure 2.8, the black and white arrows represent polarisation. Beams from the two lasers are first incident on beam splitters which reflect a small fraction of the light onto photodiodes (N1 and N2), used to normalise away laser intensity noise.²⁸ The two beams are then combined on a beam cube and sent into the optical cavity. The transmitted light from the

²⁶This is described here, not in section 2.4, as the rf system is needed to make it work.

²⁷This assumption is experimentally vindicated by the success of the repump scheme.

²⁸This normalisation is crucially important, as the lock scheme is a ‘level-lock’.

cavity is separated again by another beam cube. The transmitted intensity of each laser is measured independently by photodiodes (C1 and C2).

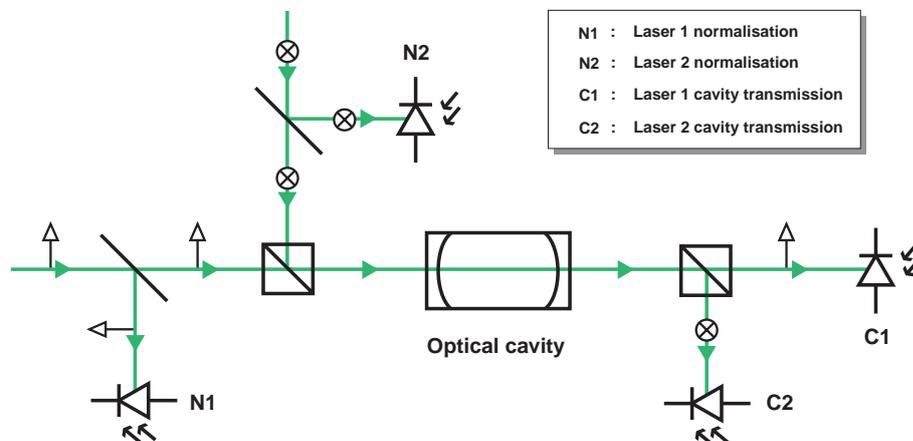


Figure 2.18: The cavity lock optical layout showing the laser polarisation in each section.

The cavity is constructed from two mirrors bonded to a quartz spacer. The cavity is 20.67 mm long, giving a free spectral range of 7.3 GHz – this is very close to one sixth of the required offset between the lasers. The mirrors are coated to 65% reflectance, resulting in a cavity finesse of around 7. The cavity is wrapped first with a ‘tape’ heater and then with spun alumina insulation (FibreFrax) to provide thermal insulation. This bundle is tightly wedged inside a 2” diameter aluminium tube, fitted with windowed aluminium endcaps. A small thermistor, in thermal contact with the cavity, measures the cavity’s temperature.

The electronic control system comprises two independent lock circuits. The first lock is responsible for stabilising the length of the cavity to the transmission from the first laser, which has a stable frequency because it is locked to the I_2 spectrometer. The transmission of this light is monitored and the cavity’s length is changed, thermally, to keep the transmitted light intensity constant. The second lock monitors the transmission of the second laser through the cavity, which is now length-stabilised. The second laser’s frequency is adjusted to maintain a constant transmitted intensity. This locks the two lasers’ frequencies together.

A block diagram of the control electronics is shown in figure 2.19. The first lock circuit is simpler than the second. Following the diagram, the cavity transmission photodiode signal, C1, is first normalised to the first laser intensity, N1. The divider circuit’s gain control is used to match the output range of the photodiodes to the input range of the differencing amplifier. The normalised signal is compared to an adjustable voltage reference, BIAS, and the difference integrated and sent to the cavity heater, CL. This lock

circuit stabilises the cavity over times longer than around 100 ms — the limit imposed by the cavity heating rate²⁹ — on shorter timescales than this the passive stability of the cavity is more than adequate.

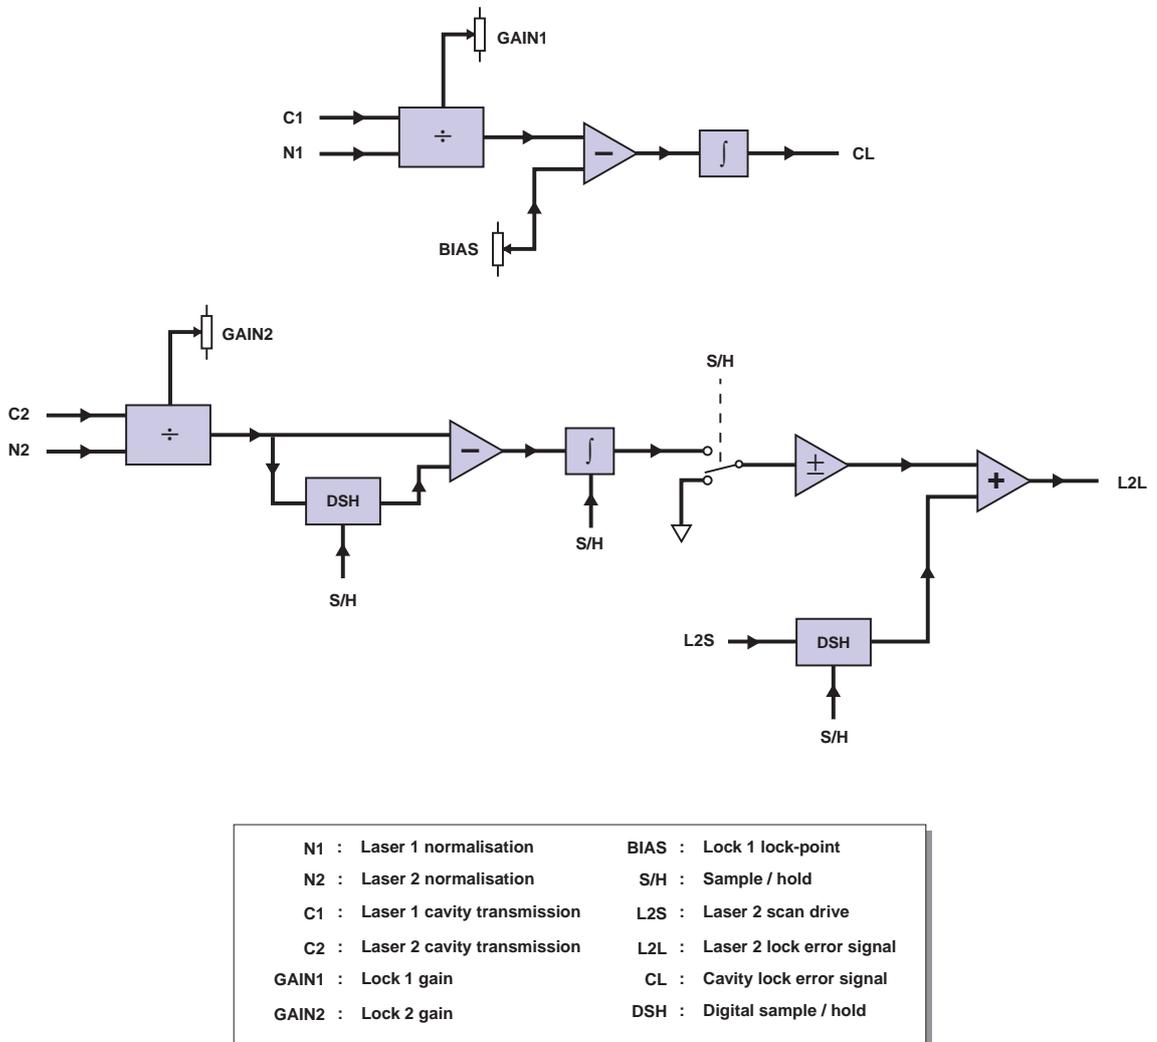


Figure 2.19: The cavity lock control electronics.

The second lock circuit is, in principle, very similar — the difference is in the way in which the lock-point is set. The first generation control electronics had two lock circuits identical to the lock circuit described above — this system proved almost impossible to lock to the correct cavity fringe; many a frustrating hour was spent watching the laser jump 7.3 GHz ! The improved lock circuit has an input for a scan voltage. When the circuit is in its scan state, the integrator is disconnected and the output simply follows the scan input (L2S). When the circuit is switched into its lock state (the control line is

²⁹I'm still amazed the cavity reacts that fast !

marked S/H), the output voltage (= scan voltage) and the normalised cavity transmission are ‘recorded’ by two digital sample and hold circuits (DSH). The lock tries to maintain the normalised cavity transmission at the recorded level, and it does this by varying the output voltage (L2L) around its recorded level. The bandwidth of this lock circuit is approximately 1 kHz. The improved lock circuit is much easier to lock to the correct cavity fringe; the laser is simply scanned to the correct frequency (described below) and the scan/lock switch thrown.

The lock circuitry is designed with low noise and long term stability in mind. One aspect of the design is interesting enough to merit mention: the digital sample and hold circuitry. The DSH’s input voltage is digitised by an ADC and the data is sent straight to a DAC, which produces the DSH’s output. The ADC and DAC are clocked rapidly, and the output simply follows the input. However, by inhibiting the DAC from loading new data, the output of the circuit can be held constant. The DAC output, and hence the DSH, has extremely good long term stability.

2.6.3 Operation

Before attempting to lock the lasers the cavity needs to be heated to its operating temperature, around 120°C, which takes approximately 30 minutes. Some adjustment of the photodiode and cavity alignment is usually necessary; both laser beams travel around 2 m before reaching the cavity, making them sensitive to small changes in the output beam pointing of the lasers.

When the cavity has reached operating temperature, it is locked to the first laser (which is already locked to the I₂ spectrometer) by simply turning on the control electronics. Locking the cavity to the first laser is worry free, as it doesn’t matter exactly which cavity fringe the system is locked to. Locking the second laser to the cavity is a little more involved, but nonetheless straightforward. The technique I have adopted is to use `laser scan` to scan the second laser frequency, at each frequency step recording the difference in LIF signals between rf on- and off-resonance. This scan of repump laser frequency versus rf transition intensity clearly shows when the repump laser is on resonance. Once the correct frequency has been identified `laser scan` is paused at that frequency, and the lock circuit switched from scan to lock. The repump is locked. A recorded scan of repump frequency versus rf transition intensity is shown in figure 2.20, the frequency axis is not calibrated but the total scan width was near to 1.25 GHz. The repump transition is clearly visible as a peak in the rf transition intensity; the ratio of intensities with the repump on- and off-resonance is seen to be around 2.8, which is typical. The smaller peak to the blue side (~200 MHz) of the main transition is due to a weaker transition (O₁(2)) that also connects X²Σ⁺(ν = 0, N = 2) to A²Π_{1/2}(ν = 0, N = 0) — introducing

a third frequency of light to the **Pump** phase to excite this transition could be a future improvement to the repump scheme.

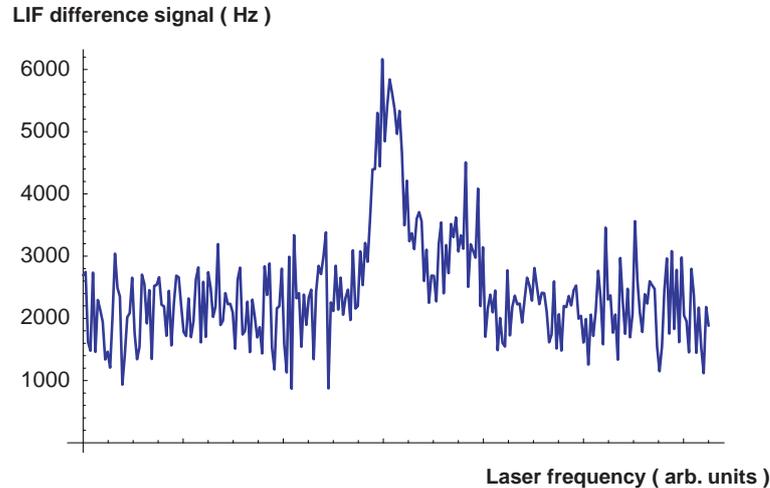


Figure 2.20: Scan of the second laser offset frequency over the repump transition.

The performance of the cavity lock has been well characterised. Measurements of the fluctuation spectrum of the normalised transmitted intensities can be related to the fluctuations in the laser offset frequency. These measurements showed that over the range $0.1 < 10$ Hz the rms amplitude of the fluctuations was less than 0.25 MHz. The lower frequency bound on these measurements is determined by the longest sampling time of the spectrum analyser. A much more stringent limit on the cavity lock's performance at low frequency is determined from the fact that the repump laser can stay locked on resonance for several hours at a time.³⁰ The long term stability of the cavity lock is limited by two independent factors:

- locking a laser to a cavity to a laser to a spectrometer is akin to balancing a stack of pencils on top of each other. Often the second laser unlocks because the first laser has unlocked from the I_2 spectrometer;
- the alignment of the cavity changes slightly over the course of a few hours which can result in a slow drift of the relative frequency of the two lasers. I think that problem this could be remedied by moving the cavity lock closer to the laser outputs to minimise the effect of changing beam pointing.

I can't emphasise enough how crucial the repump scheme has been to the EDM measurement. Its effect is far beyond that of a simple $2.5\times$ increase in EDM sensitivity. The

³⁰The proof of the pudding is in the eating, as they say.

additional factor in signal to noise of 2.5, and the corresponding decrease in integration time required to see a given signal of 2.5^2 , took the apparatus above a threshold where it became possible to operate by permitting many optimisations of the interferometer that were previously too time consuming to be practical.

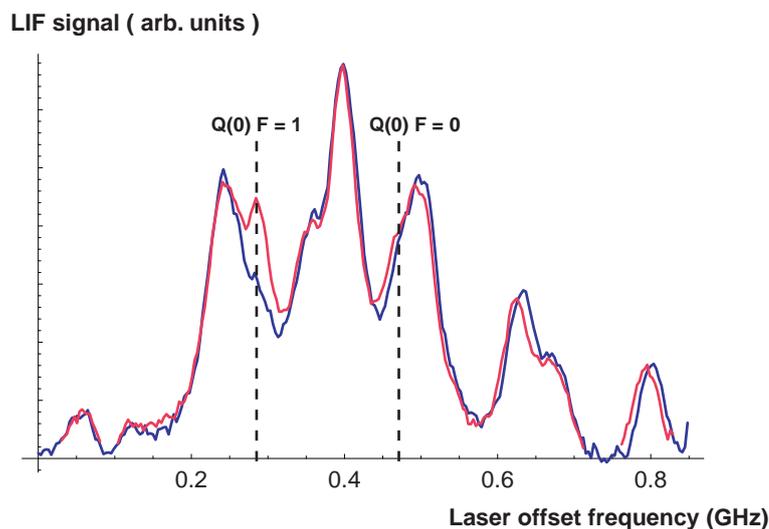


Figure 2.21: The repump really works !

Finally, I present some measurements that graphically illustrate the effect of the repump laser. Shown in figure 2.21 are two LIF scans over the Q(0) region of the spectrum. The blue curve was taken with the repump laser blocked, and the red curve with the repump present — no pump laser was used for either scan. The enhancement of the population in the $|F = 1\rangle$ state is clearly visible. There should also be some enhancement of the $|F = 0\rangle$ population. This is not visible in the scans; this is because the enhancement for $|F = 0\rangle$ is not as large as for $|F = 1\rangle$, and the population of $|F = 0\rangle$ is smaller to start with.

2.7 Electric field

So far we've put together all the parts to make a complete interferometer — now we need something to measure with it. We need to let the molecules interact with some fields — the **E & B** phase (see figure 1.4). In this section I describe the method of application and effect of electric fields.

An electric field is, of course, necessary to measure the EDM but the electric field plays a second important rôle. Without a field present the molecules will tend to 'follow' stray magnetic fields in the interaction region, resulting in uncontrolled evolution that

destroys the coherence between the $|F = 1\rangle$ Zeeman sub-levels so carefully prepared by the splitter. The electric field pins the molecules to a well defined axis, restricting their interaction with magnetic fields that are not along that axis. This electric constraint stabilizes the fringe contrast of the interferometer in the presence of stray magnetic fields. This point will be discussed later, in much more detail, in relation to systematic errors (§3.6).

2.7.1 The Stark shift

The EDM interaction is not the only one between a YbF molecule and an electric field. A much larger interaction of more mundane origin is the Stark interaction — the interaction of the induced molecular electric dipole moment with the electric field.³¹ The rf transitions driven in the **Split** and **Recombine** phases take place inside the electric field region and in order to drive these transitions effectively, account must be taken of the Stark shifting of the transition frequencies. The Stark shift of the hyperfine levels of the ground state was extensively studied during the early days of the YbF EDM project [33],[26]. Some of the results are presented here in order that the effect on the molecule of applying an electric field can be appreciated.

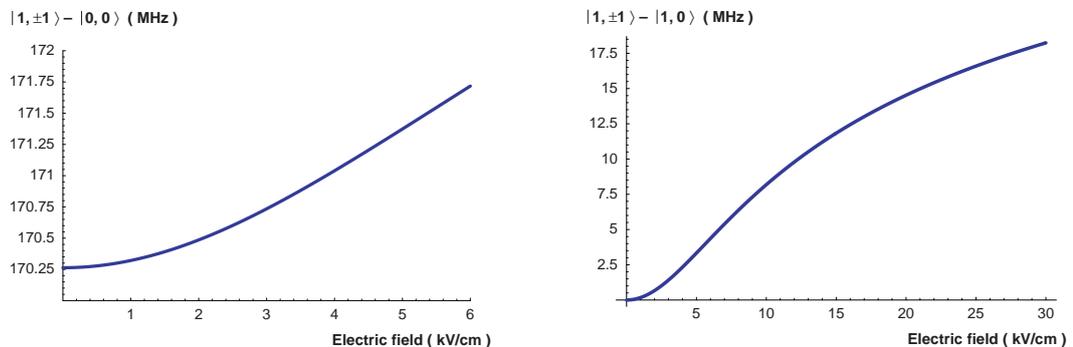


Figure 2.22: Stark shifts of the groundstate hyperfine levels.

Figure 2.22 summarises the Stark shifts of the ground state hyperfine levels over the range of electric fields of interest. The electric field in the region where the rf transitions are driven is approximately 3.3 kV/cm. The left hand graph in figure 2.22 shows that in this field the transition frequency is shifted by around 500 kHz from the zero field

³¹There exists a point of potential confusion here. Molecules are often described as “possessing a permanent electric dipole moment”. This term is a little misleading for small molecules like YbF. The “permanent” dipole moment is defined in the molecule fixed frame — for a quantum object such as a diatomic molecule it is not clear what this means (think of defining a permanent electric dipole moment of the hydrogen atom in the frame aligned with the electron-proton axis — clearly preposterous). In the absence of an external field the expectation value of the electric dipole moment of the YbF molecule is zero, neglecting T-violating effects.

frequency. The rate of change of transition frequency with field is 300 Hz/V.cm^{-1} around the operating field. Figure 2.22 also shows the $|1, \pm 1\rangle - |1, 0\rangle$ splitting for completeness — the large magnitude of this “tensor” splitting indicates the strong anisotropy of the electrons’ molecular environment (§3.6).

2.7.2 Implementation

A high electric field is desirable, as the sensitivity to the EDM increases with increasing electric field. However, it is advantageous to drive the rf transitions in a lower field because field inhomogeneity broadens the rf transitions and this effect worsens with increasing electric field. To satisfy both of these competing requirements, the electric field plates are split into three parts that define three distinct regions, two ‘guard-regions’ and a ‘C-region’. The guard regions are at the ends of the field region and it is in these regions that the rf transitions are driven. A lower voltage is applied to the ‘guard plates’ than the ‘C plates’. The guard field, 3.3 kV/cm , is chosen such that the molecules interact more strongly with the electric field than any magnetic fields they are likely to encounter — this ensures that the molecules remain ‘aligned’ to the electric field throughout the interferometer. A higher electric field is generated in the C region, 8 kV/cm . This field is limited only by the breakdown voltage of the field plates.

The field plates were made by laminating thin aluminium sheet (1 mm) onto glass plates. Each glass plate carries three separate aluminium plates (two guard plates and a C plate). The two glass plates are bonded to precision cut glass spacers which define the inter-plate spacing (12.9 mm). The plate assembly is shown in close-up in figure 2.23 — this picture shows the recombiner guard region. The plate assembly is mounted on an aluminium and titanium support structure which slots into the vacuum chamber. The plates and support assembly are made entirely of non-magnetic materials — measurements with a fluxgate magnetometer showed that the assembly is non-magnetic down to the sensitivity of the magnetometer, around 0.5 nT .³² The size and location of the field plates relative to the vacuum chamber are shown in figure 2.24; the rf antennae are included for comparison.

As mentioned above, variation of the plate spacing in the region where the rf transitions are driven can translate into a broadening of the rf lines, through the Stark shift. Mechanical measurements indicate that the variation in the plate spacing is approximately 0.5% for the recombiner guard and C regions and 1% for the splitter guard region. Using the slope of the Stark shift curve, figure 2.22, this level of inhomogeneity can be shown to correspond to a broadening of the rf transitions by about 5 kHz for the recombiner and 10 kHz for the splitter. The measured linewidths will be presented in the next section.

³²The sensitivity of the magnetometer is limited by fluctuations in the ambient field in the lab.



Figure 2.23: A close-up view of the electric field plate assembly.

Four independent power supplies generate the potentials on the field plates. Two 50 kV (maximum) Bertan high-voltage supplies charge the C plates and two 20 kV (maximum) Bertan supplies charge the guard plates (the two sets of guard plates are wired in parallel). The supplies are connected to the plates through two switching systems, one for the C region and one for the guard regions. Optical digital signals are sent from the computer to the switching system controllers to control field polarity and to switch the field on and off. The controllers drive two sets of high voltage relays: the C region relays are solenoid-actuated, air-insulated relays; the guard region relays are reed relays. The controllers and relays are designed to draw the same current regardless of their state, an important precaution against systematic effects. The relay system is in a corner of the lab some metres from the beam machine, again to minimise any systematic effects associated with relay-state dependent magnetic fields (see section 3.6).

2.7.3 Operation

I usually turn the electric field on after the lasers have been locked and the rf systems checked, as the rf transitions provide a useful diagnostic. If all is well, the only effect of the electric field is to shift and broaden the rf transitions.³³ The new centre frequencies of the rf transitions must then be located. This is accomplished, as in section 2.5.3, by monitoring the LIF signal as the rf transition frequency is scanned, using `function gener-`

³³Symptoms of all not being well include: an increased count rate from the LIF detector when the field is turned on, indicating a glow discharge or sparking; or malformed lineshapes, usually indicating a gross leakage current flowing.

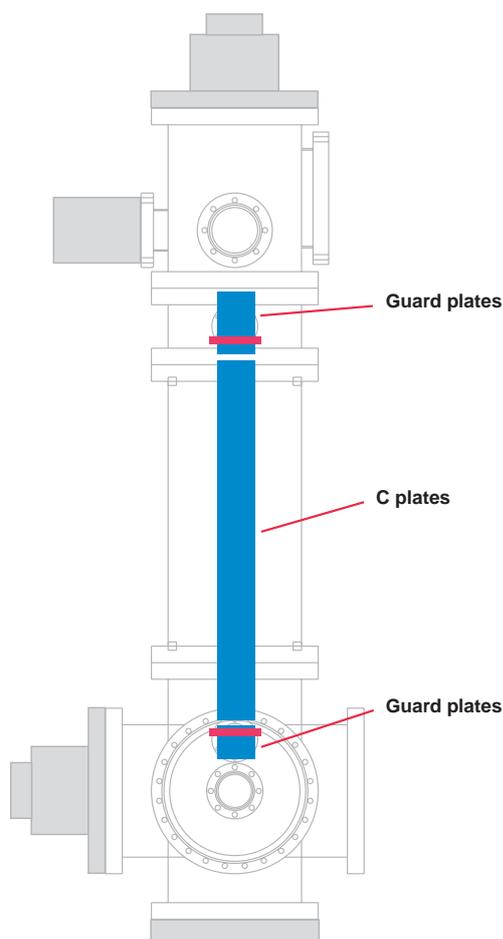


Figure 2.24: The location of the electric field plates.

ator scan. Figure 2.25 shows the rf lineshapes, measured with the plates charged to their operating voltages. The recombiner transition looks much as before, with a linewidth of 6 kHz, only slightly broadened from the zero field linewidth of 4.5 kHz — this is consistent with the broadening expected from the plate inhomogeneity. The splitter transition is, however, very different from before, being neither narrow nor particularly attractive³⁴ — the lineshape is seen to be highly asymmetric and very broad, almost 30 kHz. The splitter guard plates are known to be less parallel than the recombiner guard plates, but not enough to account for the very broad splitter rf lineshape. If figure 2.24 is re-examined, it can be seen that the splitter rf antenna lies very close to the upper end of the lower guard plates. I think that the rf transitions are probably being driven in the field gradient between the guard and C regions, broadening the lineshape.³⁵ This broadening of the

³⁴Thanks to M. R. Tarbutt for that phrase !

³⁵The field gradient results in a spatially varying resonance frequency; for a given oscillator frequency this can be thought of as a spatially varying detuning. As the molecules fly through the ‘detuning gradient’

splitter transition does not interfere with the operation of the interferometer, but it is very likely that if the transition could be narrowed the efficiency of the interferometer would be increased.

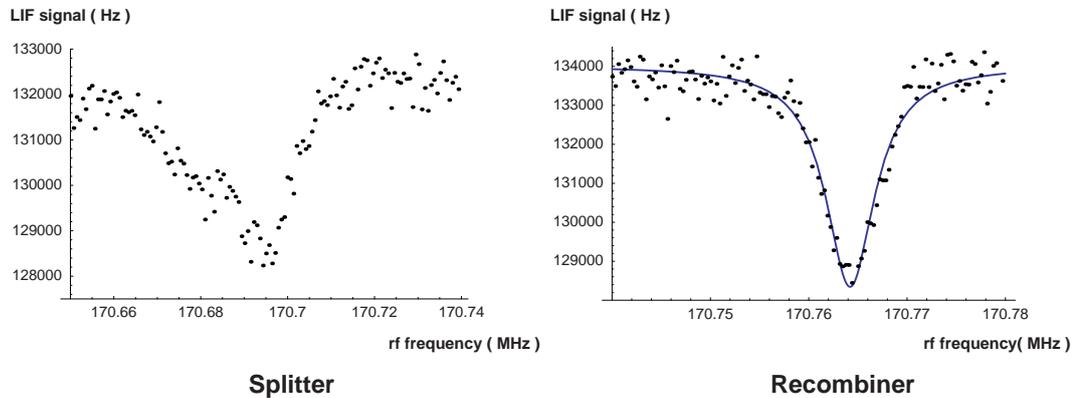


Figure 2.25: rf lineshapes with the field plates charged to their operating voltages.

The electric field also suffers from a problem of long term instability. Over the course of a day, the rf transition centre frequencies will drift, often by several linewidths. This necessitates periodic re-scanning and re-tuning of the rf systems, which is rather disruptive. The drifting problem is much worse for the splitter guard region than the recombiner guard region. My feeling is that the problem is again related to driving transitions in the fringe field of the guard plates. The area of glass between the recombiner guard and C plates does, over the course of several months, develop a thin coating, presumably of Yb metal. I think it quite possible that the very slow deposition of crud, possibly charged crud, on the plates and dielectric between them could be responsible for changing the field distribution enough to account for the slow drift of the rf transition frequencies. The solution is a complete rebuild of the field plates — new plates are under construction for the next generation of the EDM experiment that should eliminate both the drifting and broadening problems. The glass backing plates have also been eliminated from the new design, which should enable the plates to support a much higher maximum field.

A very important measurement to be made on the electric field plates is the amount of current that they leak when charged to operating voltage. It is clear that any current flowing may generate a magnetic field, and a magnetic field that reverses when the direction of the electric field is reversed can mimic the EDM. A sensitive measurement of the leakage current has been made using a technique suggested by Chris Regan [9]. The field plates are charged to their operating voltages and disconnected from the power supplies. Then, by monitoring the centre frequency of the rf transitions, the time variation of the Ramsey-type effects can occur, giving rise to complicated lineshapes.

slowly decaying field can be inferred. This is readily related to a discharge current. Figure 2.26 shows a typical set of results. Each small cell on the figure corresponds to 50 ms of counting the LIF signal, the brighter the cell, the higher the LIF signal. The figure was ‘filled in’ left to right, bottom to top, and the whole figure represents around 6 minutes of data acquisition. The bright rf transition is clearly seen decaying from its starting point, 170.372 MHz, towards the zero field frequency of 170.254 MHz. Analysis of similar data taken with various field plate connections showed that the leakage current from any field plate is less than 10 nA at the operating voltages.³⁶ It will be shown that this level of leakage is incapable of generating a significant false EDM (§3.6).

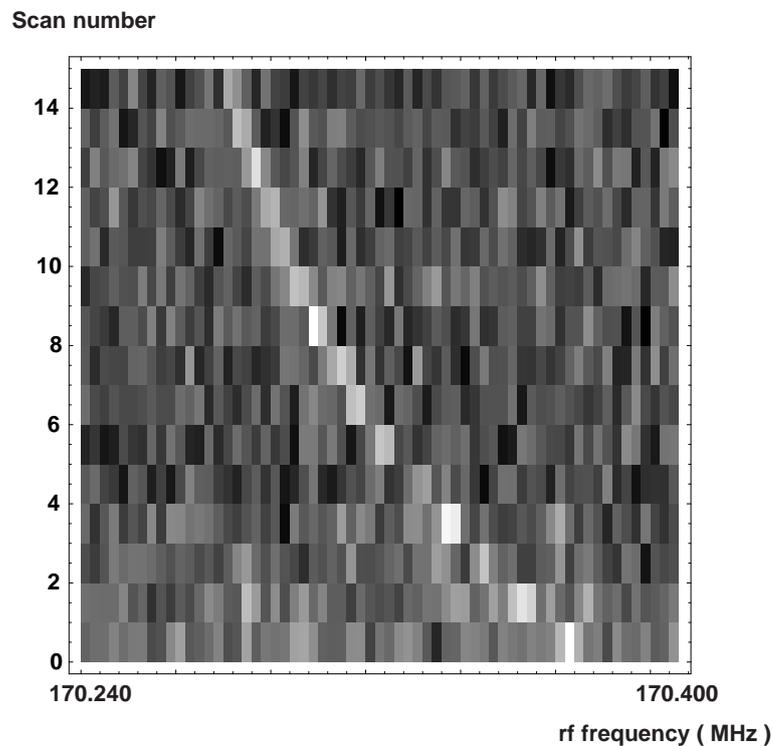


Figure 2.26: Watching the electric field decay using rf transitions.

2.8 Interference

The final ingredient of the interferometer is a magnetic field, completing the **E & B** phase. By changing the applied magnetic field we hope to observe interference between the two arms of the interferometer.

³⁶A significant fraction of this leakage is probably external to the machine, in the cables and vacuum feedthroughs — it was not feasible to disconnect these to make the measurements.

In this section I will first present a theoretical model of the interferometer. This model serves as a useful guide, indicating how experimental parameters may be adjusted to maximise the interferometer’s output. The model also generates a lineshape that can be fitted to measured data to extract experimentally relevant quantities e.g. the residual magnetic field inside the magnetic shields. This will be followed by a description of the hardware used to produce the magnetic field. Finally, some interference measurements will be presented.

2.8.1 Theory

During their passage through the interferometer the molecules interact simultaneously with static magnetic and electric fields, rf oscillating magnetic fields and optical frequency electric fields. Analytic solution of this problem is not possible, moreover a numerical attack on the problem is impeded by the broad frequency spectrum of the various fields. In order to progress the problem must be simplified — the approach I will take is to split the molecules’ interactions into discrete stages that can be solved analytically, corresponding closely to the various phases of the interferometer described previously.

The **Pump** phase is assumed to be perfect. As such, its action is represented by the projection matrix,

$$\text{PUMP} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix},$$

which is the same in both the Cartesian- and z-bases. This is an adequate model — making the **Pump** phase imperfect simply reduces the contrast of the interference signal.

The **Split** and **Recombine** phases are considered to be identical. In the interferometer the rf transitions take place in both the static electric and magnetic fields. In the model the molecules are assumed to interact only with the rf magnetic field in the splitter and recombiner — neglecting the static electric and magnetic fields in the rf regions doesn’t change the gross character of the results. The rf magnetic field is modelled as switching on and off discontinuously, as in section 2.5.1.³⁷ Under these approximations the result can be carried directly from equation (2.3), the matrix \mathbb{RF} describes the molecules’ time evolution in the splitter and recombiner.

After the splitter the molecules interact with the static electric and magnetic fields, the **E & B** phase. To first order, the Stark interaction shifts the $|1, \pm 1\rangle$ states equally — the interferometer is not sensitive to this ‘common-mode’ shift, consequently the Stark

³⁷This may neglect some interesting effects occurring in the wings of the rf field — as the experiment becomes more sensitive these may become relevant.

interaction is not included. The EDM interaction is included explicitly. The interaction Hamiltonian is,

$$\hat{H} = (\mu_B B_{ext} - \eta d_e E_{eff}) \hat{\sigma}_z .$$

I assume, again, that the interactions turn on and off discontinuously. The magnetic field and electric fields are taken to both lie along the z-axis. The Hamiltonian is readily integrated, giving a time evolution operator,

$$\mathbb{E}\mathbb{B} = \begin{pmatrix} e^{i\frac{\Omega}{2}\tau} & 0 & 0 & 0 \\ 0 & e^{-i(\frac{\Omega}{2}+\xi)\tau} & 0 & 0 \\ 0 & 0 & e^{-i\frac{\Omega}{2}\tau} & 0 \\ 0 & 0 & 0 & e^{-i(\frac{\Omega}{2}-\xi)\tau} \end{pmatrix}_z ,$$

where $\xi = \frac{1}{\hbar}(\mu_B B_{ext} - \eta d_e E_{eff})$ and τ is the interaction time.

The final phase of the interferometer, the **Probe** phase, is taken to simply measure the population in the state $|0, 0\rangle$.

Compositing these stages in the correct order gives an operator that describes the evolution of an arbitrary initial state through the interferometer,

$$\text{INT} = \text{CB} \cdot \text{RF}(t = T + \tau \rightarrow 2T + \tau) \cdot \text{CB} \cdot$$

$$\mathbb{E}\mathbb{B}(t = T \rightarrow T + \tau) \cdot \text{CB} \cdot \text{RF}(t = 0 \rightarrow T) \cdot \text{CB} \cdot \text{PUMP}$$

Applying this operator to an initial state, chosen for simplicity to be $|0, 0\rangle$, the probability for the final state to be measured as $|0, 0\rangle$ is

$$P_0 = \left[\cos^2 \left(\frac{aT}{2} \right) + \theta_c^2 \sin^2 \left(\frac{aT}{2} \right) \right]^2 + \theta_s^4 \sin^4 \left(\frac{aT}{2} \right) \cos^2 (\xi\tau) \\ - 2\theta_s^2 \sin^2 \left(\frac{aT}{2} \right) \cos (\xi\tau) \left\{ \left[\cos^2 \left(\frac{aT}{2} \right) - \theta_c^2 \sin^2 \left(\frac{aT}{2} \right) \right] \cos (\delta\tau - \phi) \right. \\ \left. + 2\theta_c \cos \left(\frac{aT}{2} \right) \sin \left(\frac{aT}{2} \right) \sin (\delta\tau - \phi) \right\} , \quad (2.5)$$

where ϕ is the phase difference between the two rf magnetic fields. This expression is quite formidable, and it is primarily presented here to save anyone else the wrist ache of re-deriving it. It is however, worth pausing to look a little more closely. For a given rf field strength, the first term is simply a constant offset. The second term is the ‘interference’ term that we are interested in and the third term is a ‘Ramsey’ term that depends on the relative phase of the rf magnetic fields.

Under normal operating conditions the electric field in the splitter guard region is somewhat different to that in the recombiner guard region, due to a slight difference in

the field plate spacing. As a result, the splitter and recombiner transition frequencies differ by about 75 kHz. The two rf oscillators are run at different frequencies, ensuring that both the splitter and recombiner are on resonance. The frequency difference between the oscillators is so much larger than one over the integration time³⁸ that the Ramsey term in (2.5) is averaged away. This leads to the much simpler expression,

$$\bar{P}_0 = \left[\cos^2 \left(\frac{aT}{2} \right) + \theta_c^2 \sin^2 \left(\frac{aT}{2} \right) \right]^2 + \theta_s^4 \sin^4 \left(\frac{aT}{2} \right) \cos^2(\xi\tau), \quad (2.6)$$

which describes the interferometer when the difference between the oscillator frequencies is large and both oscillators are equally detuned from their respective resonance frequencies,

The next step in modelling the interferometer is to take account of the molecules' velocity distribution. As in section 2.5.1, the procedure is to average (2.6), reparameterised in terms of velocity, over the probability distribution (2.1),

$$\langle \bar{P}_0 \rangle = \int_0^\infty \frac{2}{\alpha^4} v^3 \exp\left(\frac{-v^2}{\alpha^2}\right) \bar{P}_0(v) dv. \quad (2.7)$$

The integration must again be performed numerically.

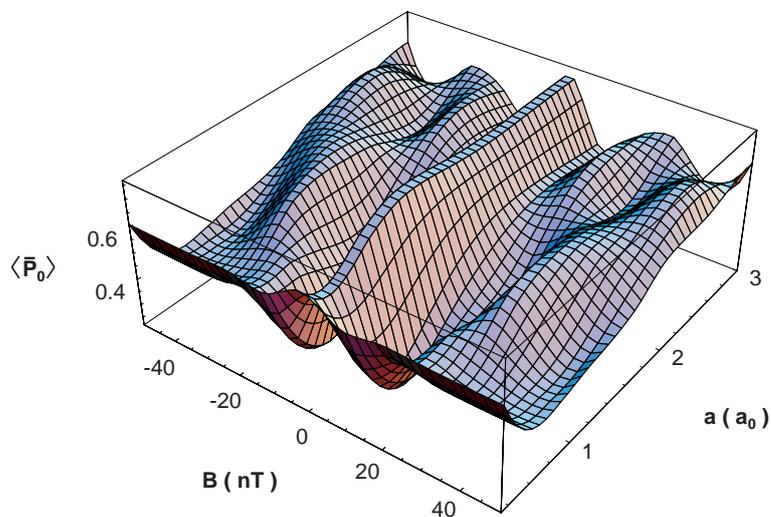


Figure 2.27: Power dependence of the interference lineshape.

An important question that the model can answer is, “What rf magnetic field strength optimises the interference signal ?” To answer this, the integral (2.7) was evaluated for a range of static magnetic fields and rf interaction strengths. Figure 2.27 shows the result, where the rf interaction strength is again measured in terms of the Rabi frequency that

³⁸Which is about 100 s

gives the molecules with the most probable velocity a π -pulse, a_0 . The greater the slope of the interference lineshape in the B-direction, the greater the sensitivity to the EDM. Figure 2.28 shows the maximum value of this slope as a function of the rf interaction strength. The peak corresponds to an rf interaction strength of 92% a_0 . The conclusion is that optimising the splitter and recombiner transition separately using pump-probe experiments very nearly optimises the interference signal. The lineshape calculated at this rf interaction strength is shown in figure 2.29.

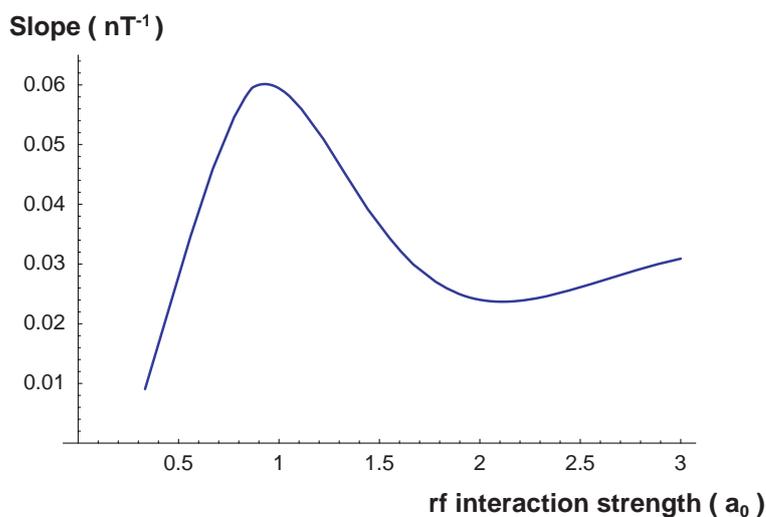


Figure 2.28: Lineshape slope vs. rf interaction strength.

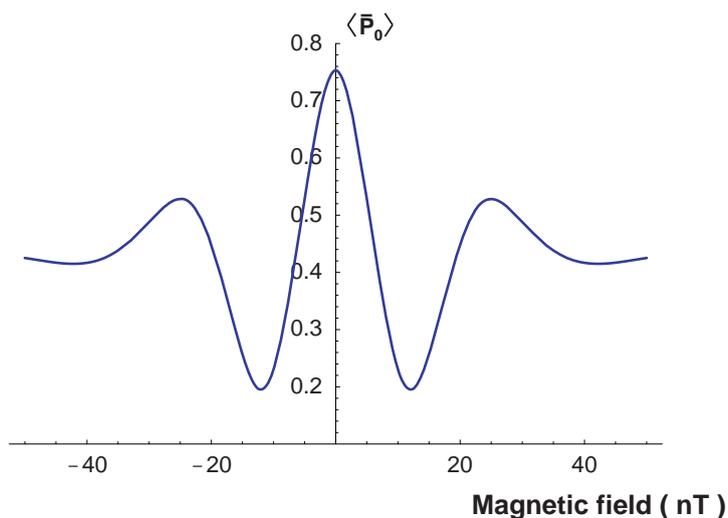


Figure 2.29: The optimal interference lineshape.

There is one more interesting thing to learn from the model. The above results are

all derived with the assumption that the two oscillator frequencies are different and as such the phase coherence between them quickly averages away. It would be technically feasible to adjust the voltages supplied to the guard regions such that the electric fields, and hence the transition frequencies, were equal. In this case the two rf oscillators could be phase-locked together. The question is, “is there any advantage in doing this ?” The answer is, “yes”. By integrating (2.5) in the fashion of (2.7), it is found that the slope of the interference lineshape is increased by 30% if the two oscillators are locked in antiphase. This can be understood qualitatively as follows: a molecule that has a velocity not equal to the most probable velocity will receive a $q\pi$ pulse. Even if there is no magnetic field or EDM interaction, after a second $q\pi$ pulse it will not in general return to the $|0, 0\rangle$ state.³⁹ However, if the phase between the oscillators is locked at π , then each molecule receives a $q\pi$ and a $-q\pi$ pulse. In the absence of a phase difference from the **E & B** phase, all molecules, regardless of velocity, will return to the $|0, 0\rangle$ state — this is the essence of the improvement.

It should be noted that phase-locking the oscillators together also has some disadvantages. The main disadvantage is an increased susceptibility to systematic errors; under the phase averaged scheme, changing the detuning of the rf oscillators simply changes the lineshape intensity (for small detunings) whereas under the phase-locked scheme changing the detuning can change and shift the interference lineshape. It is through this mechanism that imperfect reversals, of the electric field for instance, can introduce significant systematic errors. The new field plates mentioned in 2.7 have been designed to allow the fields in the splitter and recombiner region to be made equal, so that the phase-locked scheme can be investigated.

2.8.2 Magnetic field hardware

Magnetic fields are generated by passing current through copper wire loops glued to the inside of the inner magnetic shield. Three independent coils generate fields in the x-, y- and z-directions — only the z-coil is used routinely, but the other two are useful for performing diagnostics. The field distribution inside the shield produced by the z-coils has been calculated — they give a field which is uniform to 1% over the beam volume [22].

Current is driven through the field coils by a computer-controlled current supply. The supply reads bytes from the computer’s RS232 port, carried to the supply over an optical fibre. These bytes are decoded by a UART and fed to a DAC, the output of which drives a high-stability low-noise voltage to current converter. The supply can drive cur-

³⁹Figure 2.29 shows that, for the phase averaged case, only 75% of the molecules return to the $|0, 0\rangle$ state.

rents between -3 mA and 3 mA with 8-bit resolution. The sensitivity of the z-coil with the magnetic shields in place has been measured to be 16.5 nT.mA $^{-1}$ with a fluxgate magnetometer — the supply can therefore generate fields between -50 nT and 50 nT with a 0.4 nT resolution.

2.8.3 Interference measurements

With both lasers locked, the field plates charged to high voltage and the rf oscillators retuned to their new frequencies we are ready to make an interference measurement. `Laser scan` is used to scan the magnetic field⁴⁰, recording the LIF signal at each step. Figure 2.30 shows a lineshape recorded in this manner. Each point corresponds to 12 s of integration. The blue curve shows the lineshape derived above fitted to the data.⁴¹ The fit parameters are a constant offset, a sloping background, an overall normalisation, the residual magnetic field and the oven temperature.⁴² The fit is certainly convincing. How-

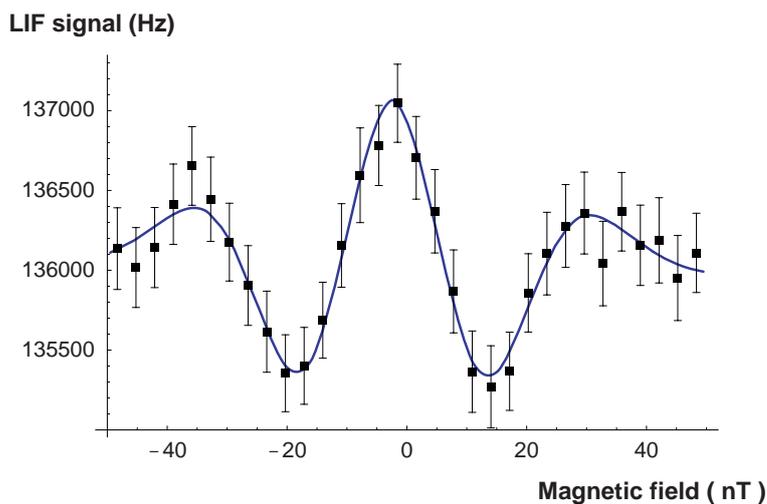


Figure 2.30: A measured interference lineshape.

ever, examining the fitted parameters does raise one concern — the fitted oven temperature is much higher than I can believe. Optical pyrometer measurements made long ago suggested the oven temperature was within a few hundred K of 1500 K; the fitted oven temperature is typically much closer to 2500 K. Although I run the oven hotter now than

⁴⁰Typically, a different current supply than described above is used when scanning for interference patterns — just the voltage to current converter stage of the computer controlled supply.

⁴¹The error bars shown on the data points are the standard errors of each point. It should be noted that the distribution each point is drawn from is very far from Gaussian — this explains why so many of the points are within 1σ of the model curve. The error bars are included as they give an indication of the signal:noise ratio.

⁴²This is a particularly indirect way to measure a temperature.

when the pyrometer measurements were made, I don't believe I run it 1000 K hotter. The most plausible explanation for the discrepant fitted temperature is that the assumption that the source is purely effusive is flawed. Recalculating the model lineshape shows that the measured lineshape can be modelled well if the molecular velocity distribution in the beam is changed to that from a source that shows some 'jetting'. Inside a jet source, the vapour pressure is high enough that molecules are 'pushed' out of the source. The most significant effect of the jetting is that some of the molecules' thermal energy is converted to centre of mass energy; the molecular velocity distribution is narrowed and pushed to higher velocity. The measured lineshape can be fitted with a velocity distribution corresponding to a 1500 K source if some jetting is included. It should be noted that an incorrect current \rightarrow magnetic field calibration would also show up as an unrealistic oven temperature in the fit. I think the calibration is sound — it was confirmed by Zeeman shift measurements — but it such an important parameter that it should be checked again.⁴³ Notwithstanding this problem, fitting to a model lineshape is very useful — the fit reliably extracts the signal:noise ratio and the residual magnetic field inside the inner magnetic shield.⁴⁴ The signal:noise ratio of the interference curve is directly related to the EDM sensitivity, making it the ultimate diagnostic, and the residual magnetic field must be known before an EDM measurement can be made. A convenient fitting package, written in Mathematica, allows me to fit data and extract these important parameters easily.

2.9 Comments

The EDM measuring tool, the interferometer, is complete. Before moving on to measuring the EDM with this interferometer, the subject of the next chapter, I'd like to give some impression of how the interferometer actually runs. Beginning in the morning, starting the interferometer goes as follows :

1. The oven, which has been at baking temperature overnight, is turned up to its operating temperature. It needs a couple of hours to heat up.
2. The lasers need to be started. They take around one hour to stabilise before they can be used.
3. Some of the optics usually need a little tweaking. This can be done while the laser and oven warm up.
4. Once the oven is up to temperature, the LIF signal is looked at. Sitting on a molecular fluorescence peak, the oven alignment and probe and pump beam alignment are

⁴³Unfortunately at the time of writing the interferometer is not operational so the check cannot be made.

⁴⁴Of course, so would fitting to just about any function with a hump.

optimised. The repump lock cavity is warmed up.

5. The first laser is locked to the I_2 spectrometer.
6. The magnetic shields are degaussed.
7. The rf system is checked, then the second laser is locked to the repump transition.
8. The electric field is turned on. The rf transitions are checked and the new centre frequencies are measured.
9. An interference measurement is made, and the residual magnetic field and signal:noise ratio measured.

The interferometer is now ready to use. The above procedure takes 3–4 hours in all — the interferometer is usually ready to use around lunchtime.

Once the interferometer is running the main inconveniences are the drifting rf transition frequencies and the laser system unlocking. If everything is going well, the interferometer can stay running for 10 hours before the oven starts to fade.⁴⁵ Now, we should get on and measure the EDM before something breaks !

⁴⁵and the experimenters start to fade too.

Chapter 3

Measuring the EDM

Now, with the details of the interferometer out of the way, we can think about measuring the EDM. The method of measuring the EDM described in chapter 1 — applying a magnetic bias field to move to the steepest part of the interferometer fringe and then reversing a large electric field — whilst in principle adequate, is not the most efficient way to conduct an EDM measurement. The problem that has to be combatted is non-statistical¹ noise in the interferometer's output. The optimal method for making a measurement is dictated by the properties of this non-statistical noise. In the first part of this chapter I present a measurement of the properties of the noise on the interferometer's output and describe the measurement scheme implemented to minimize the effects of this noise.

Once a measurement has been made the data must be analysed. A description is given below of how the final value of the EDM is derived from the data, along with many useful diagnostics. Some of the possible sources of systematic error are discussed after that. Particular attention will be paid to the few systematic problems that I have so far encountered, including details of how these were overcome. The chapter ends with a presentation of the complete EDM measurement.

3.1 Noise performance

Statistical limit

There is a fundamental limit imposed upon the sensitivity of the interferometer. The output of the interferometer's LIF detector is a stream of pulses, corresponding to photons arriving at the photomultiplier. This stochastic process is inherently noisy.² It is well

¹Throughout this chapter non-statistical will be used to describe noise that is not counting noise — it isn't meant to imply that the noise does not permit statistical treatment.

²If there are suitable quantum correlations between the particles this noise can be squeezed, but there is no source of such correlations in the interferometer.

known for this process, the shot noise process, that if N counts are expected in one counting period the actual count will belong to a distribution of mean N and standard deviation

$$\sigma = \sqrt{N} . \quad (3.1)$$

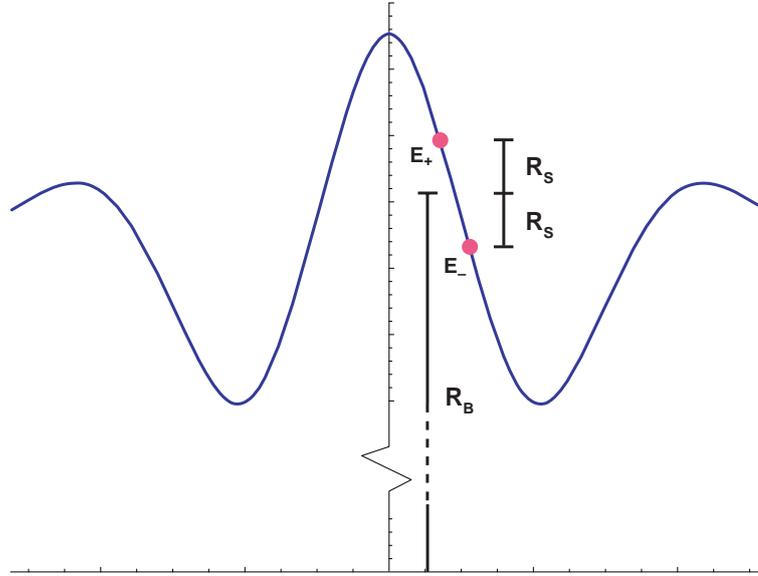


Figure 3.1: A simple EDM experiment.

It is interesting to predict what limit this shot noise puts on the interferometer's sensitivity to the EDM. Consider a primitive EDM experiment, as chapter 1. Figure 3.1 shows the scheme; R_S is the rate of LIF pulses attributable to the EDM interaction and R_B is the total 'background' rate. Given that we have a time T to make the measurement, we decide to spend $T/2$ counting at the higher point (E_+) and $T/2$ at the lower point (E_-).³ The signal, S , is the difference between the two counts,

$$S = \frac{T}{2}(R_B + R_S) - \frac{T}{2}(R_B - R_S) = R_S T = \frac{\eta d_e E_{\text{eff}}}{\mu_B} \cdot m \cdot T ,$$

where the fringe shift due to the EDM has been rewritten in terms of an effective magnetic field $\eta d_e E_{\text{eff}} / \mu_B$ multiplying the slope of the interference lineshape with respect to the magnetic field, m . In the limit of large background, small signal, the uncertainty on the signal is

$$\sigma_s \simeq \sqrt{R_B T} ,$$

³This scheme is optimal if the signal rate, R_S , is much smaller than the background rate, R_B . If this is not the case then it becomes advantageous to spend more time counting at the higher point, αT , and less at the lower point, $(1 - \alpha)T$, where $\alpha = \frac{(1+\gamma)+\sqrt{1-\gamma^2}}{2\gamma}$ and $\gamma = R_S/R_B$.

using (3.1). Equating the signal and the uncertainty, we can extract the minimum EDM that we could differentiate from zero with 1σ certainty in time T ,

$$d_e = \frac{\mu_B}{m \eta E_{\text{eff}}} \sqrt{\frac{R_B}{T}}. \quad (3.2)$$

Inserting some realistic values for m (150 Hz·nT⁻¹), η (0.5), E_{eff} (25 GV·cm⁻¹), R_B (125 kHz) and T (23.5 hours, about the size of the complete dataset) gives

$$\sigma_{d_e} = 3.6 \times 10^{-26} \text{ e} \cdot \text{cm}$$

for the 1σ uncertainty in the measured EDM. This value will be compared to the measured value presented at the end of this chapter.

In the above discussion we were careful to count for equal times at the E_+ point and the E_- point, as it was stated, without proof, that this was the optimum scheme. As the interferometer, sadly, runs in the régime where $R_S \ll R_B$ we will always want to divide the integration time up this way. However, no detailed consideration was given to how the division should be made — is it best to count for $T/2$ at E_+ , and then switch to E_- and count for the remaining $T/2$, or is it better to switch rapidly between E_+ and E_- , such that in total each point is visited for $T/2$? The answer is that it doesn't matter, as long as the interferometer's output is simply shot noise. It makes no difference when the measurements are made if the pulses are truly uncorrelated — in other words, the Poisson distribution has a flat frequency spectrum.

The real interferometer's output is not purely shot noise, though. Factors such as electrical noise, wandering beam intensity, laser instability etc. all contribute extra fluctuations to the measured signal, and these fluctuations display strong time correlations. Under these conditions all divisions of the integration time are certainly not equally good. In order to design a good data collection scheme, the correlated fluctuations in the interferometer's output must be measured.

The noise power spectrum

Conceptually, the simplest way to measure the frequency spectrum of the fluctuations in the interferometer's output is to measure the count rate as a function of time and take the Fourier transform. A small C++ program was written to count the number of LIF detector pulses arriving in consecutive time bins. The noise power spectrum of these data was then extracted by calculating the discrete Fourier transform. Figure 3.2 shows the spectrum, taken with the interferometer fully operational — oven hot, laser locked etc.⁴ — the bin period was 8 ms. The inset is a smoothed and expanded view of the low-frequency region. The most striking feature of the figure is the strong harmonic series with fundamental

⁴The repump laser was not used.

frequency 3.1 Hz. This series is due to a flaw in the data acquisition scheme. It is found that the period for which the computer holds the LIF counter gate open is not constant; it is modulated by ~ 0.5 ms at a frequency of 3.1 Hz. The origin of this 3.1 Hz modulation is unknown, but I have a strong suspicion that the data acquisition board's driver software is at fault. This is a topic for future investigation. The gate period modulation — and hence the peaks in the noise spectrum — is also a 'feature' of EDM2, the program which takes EDM data (§3.3). Fortunately, as is shown in section 3.2, the noise associated with these peaks can be rejected by a good data collection scheme.

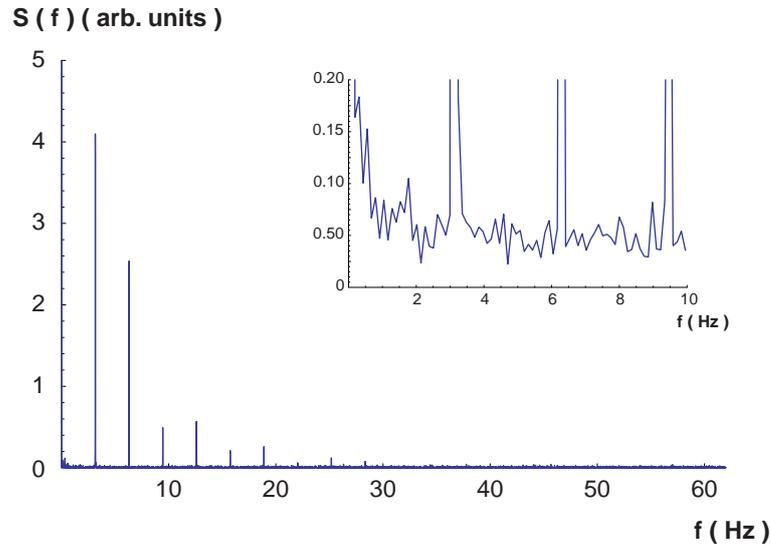


Figure 3.2: The measured frequency spectrum of the noise in the interferometer's output.

As well as knowing the frequencies at which there are peaks in the noise spectrum, it is important to know the shape of the smooth part of the noise spectrum; particularly the low-frequency deviations from pure shot-noise. The inset of figure 3.2 clearly shows the low frequency fluctuations for frequencies < 3 Hz. A particularly direct way to investigate these low frequency correlations is the Allan variance, often used to quantify oscillator stability. Consider counting pulses for a time T , then waiting for a time τ , then again counting pulses for a time T . Writing the number of pulses counted between t and $t + T$ as $M_T(t)$, the delayed Allan variance is defined as,

$$\sigma_A^2(T, \tau) = \frac{1}{2T^2} \langle [M_T(t + T + \tau) - M_T(t)]^2 \rangle,$$

where the average is over the starting time, t . An efficient way to measure the delayed Allan variance is to count for T and wait for τ many times in a row, recording the count for each gate period T — the variance is then computed as the square difference between neighbouring count rates, averaged over all neighbouring pairs.

For our purposes, the delayed Allan variance is valuable as, apart from being easy to measure, it is related to the noise power spectrum of the pulse train. The relationship is [34],

$$\sigma_A^2(T, \tau) = \frac{4}{\pi T^2} \int_0^\infty \frac{d\omega}{\omega^2} S(\omega) \sin^2\left(\frac{\omega T}{2}\right) \sin^2\left(\frac{\omega(T + \tau)}{2}\right), \quad (3.3)$$

where $S(\omega)$ is the noise power spectrum of the pulse train. The shot noise power spectrum is readily transformed [34],

$$S(\omega) = 2R_0 \quad \Leftrightarrow \quad \sigma_A^2(T, \tau) = \frac{R_0}{T}, \quad (3.4)$$

where R_0 is the average count rate. Any deviation of the Allan variance from this form is indicative of non-statistical noise. Low-frequency fluctuations are often well modelled by the '1/f' power spectrum. This spectrum can also be transformed [34],

$$S(\omega) = \frac{2\pi C}{|\omega|} \quad \Leftrightarrow \quad \sigma_A^2(T, \tau) = 2C \left(1 + \frac{\tau}{T}\right) \log 2, \quad (3.5)$$

where C is the '1/f constant'.

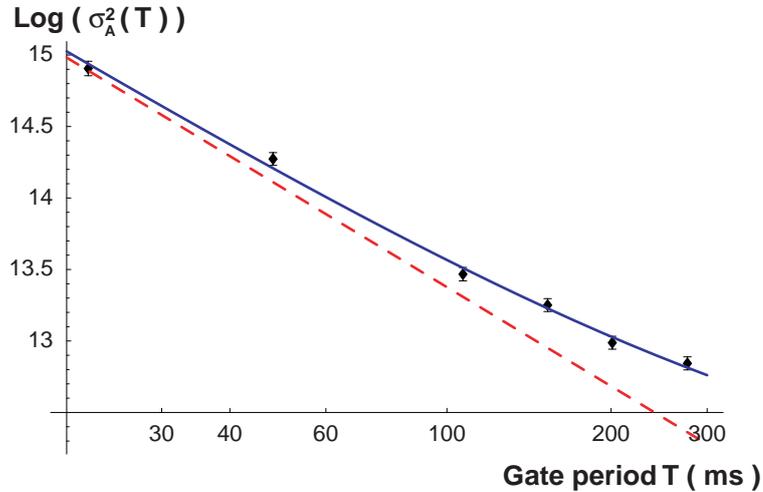


Figure 3.3: The measured delayed Allan variance.

The EDM data acquisition program, EDM2, which will be described fully later in this chapter, was configured to measure the sequence of counts described above. The gate period, T , was varied over the range 20–300 ms, the delay time, τ , was fixed at 11 ms — for each different gate period the program recorded the number of counts in 1024 consecutive periods. The data were taken with the machine fully operational. Figure 3.3 shows the results plotted logarithmically. The blue curve is a fit to the sum of the Allan variances (3.4) and (3.5). The dashed red line shows the variance that would be expected

for pure shot-noise (3.4) — notice how graphically the Allan variance shows the low-frequency excess noise. Points to note from the figure: for high frequencies (short gate periods) the measured Allan variance converges to the shot-noise limit; the low-frequency noise starts to take hold at a gate period ~ 150 ms, consistent with the observation from figure 3.2 that the low-frequency fluctuations are important for frequencies < 3 Hz.

Now, with a knowledge of the power spectrum, we can design a scheme that is resistant to the low-frequency, non-statistical noise that is present in the interferometer's output — this is the subject of the next section.

3.2 Switching and noise suppression

A commonly used technique for rejecting low frequency noise in a system's output is that of phase-sensitive detection — some parameter of the system is modulated at a well-defined frequency, and the signal is derived from the component of the output at that frequency. The essential result is that only noise in the output *around the modulation frequency* can contribute to the noise in the signal. For systems that exhibit $1/f$ noise spectra the reduction in noise can be very significant.

The EDM measurement uses a generalisation of this technique [35]. The differences are that more than one parameter is modulated, and that the modulation is done in a binary fashion i.e. the parameters are switched back and forth between two values rather than continuously varied. The details of the phase-sensitive detection will be described shortly, but first some more words are needed about the measurement method.

The most primitive method of measuring the EDM, switching the electric field in the presence of a constant magnetic bias field B_0 , has a serious shortcoming — there are many spurious effects that could mimic the EDM signal. The problem is that only one 'channel', the electric field direction, is switched. A much richer dataset is obtained if more than one channel is switched; sophisticated diagnostics can be extracted from the data that give both a detailed characterisation of the interferometer's performance and an invaluable discriminant against phoney EDMs. The first enhancement to be made is to not only switch the direction of the electric field, but also (independently) the direction of the magnetic bias field. This additional channel enables the discrimination of many spurious effects — e.g. a background signal that depends on the electric field direction — from the EDM signal. The second enhancement is to switch between two different magnetic field magnitudes differing by a small increment, ΔB . This channel allows a measure of the interference lineshape slope to be extracted from the data, as is necessary if the EDM is to be measured in physical units. The magnitude of ΔB is chosen to be small enough that the interferometer's response is linear, but large enough that the change in signal

can quickly be resolved from the noise. In total there are eight different ‘states’ that the interferometer can be in. Figure 3.4 shows these states pictorially — as usual, the EDM interaction has been greatly exaggerated.

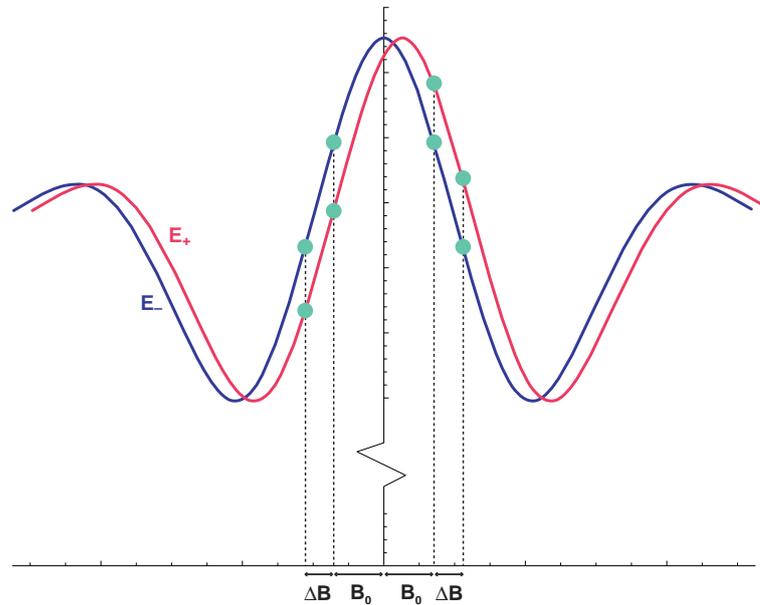


Figure 3.4: A more sophisticated EDM experiment.

To describe the way in which the EDM data is taken a terminology must be introduced. The fundamental unit of data is a ‘point’. The point corresponds to a single gate period of the computer’s counter (usually 48 ms), and all points are of equal duration — the interferometer is in one fixed state for the duration of a point. By convention points are recorded as rates, not as a number of counts. Points are taken in sequences (usually of 1024 points), and the interferometer’s state is changed between points — each of these sets of points is known as a ‘block’. In each block the interferometer’s state changes according to three ‘waveforms’, one for each channel i.e. for each switched parameter: E , B and ΔB . The waveforms are binary strings with as many bits as the block has points — each bit in the waveform represents the state of its channel for the corresponding point in the block. Figure 3.5 may help to make this more clear — the block shown is much shorter than a real block (8 vs. 1024 points). Blocks taken contiguously are called a ‘cluster’.

A compact notation is adopted for the waveforms. The waveforms are decomposed into a basis of square-waves of period 2^n and length equal to that of a block, 2^N .⁵ A waveform can then be written simply as an N -bit binary number, a ‘waveform code’,

⁵By convention, each basis waveform is assumed to start high. Also, note that this basis doesn’t span all possible 2^N bit sequences.

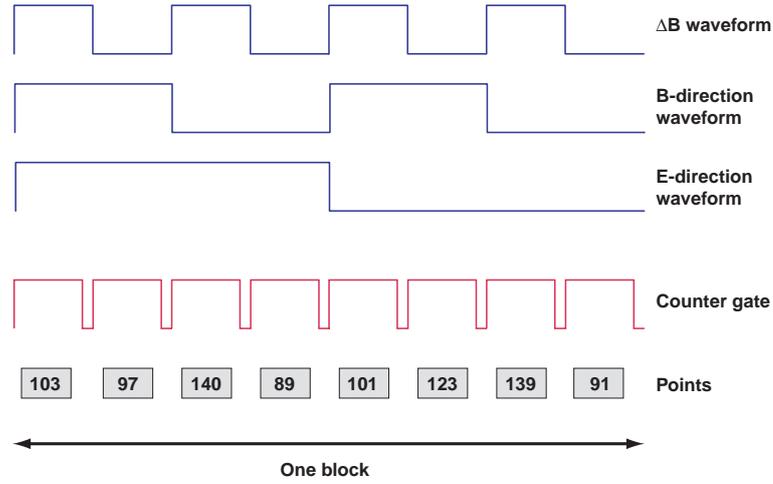


Figure 3.5: Points, blocks and waveforms.

which specifies which basis waveforms are needed to generate it, where the generation operation is the XOR. Some examples should make this clear; the waveform codes are shown on the left, and the waveforms they generate on the right,

$$\begin{aligned}
 001 &\rightarrow 10101010, \\
 010 &\rightarrow 11001100, \\
 100 &\rightarrow 11110000, \\
 101 &= 100 \otimes 001 \rightarrow 11110000 \otimes 10101010 = 01011010.
 \end{aligned}$$

Data collected as detailed above can be analysed to give the part of the interferometer's output correlated with any of the modulation channels, or any product of the modulation channels. The analysis procedure used to extract the EDM and the software written to implement the analysis will be described fully in section 3.4, but the principle will be given here. The waveforms used to switch the electric field direction, magnetic field direction and magnetic field magnitude are labelled S_E , S_B and $S_{\Delta B}$ respectively. To extract the part Q_S of the interferometer's output correlated with the waveform S , a block, D , is simply 'dotted' with the waveform. Explicitly,

$$Q_S = \frac{1}{2^N} \sum_{i=1}^{2^N} (-1)^{S(i)} D(i), \quad (3.6)$$

where $S(i)$ is the i -th bit of the waveform, and $D(i)$ is the i -th point in the block. Given the three channel waveforms, there are eight interesting, linearly-independent, 'analysis channels' that can be extracted. The waveforms for these channels are shown in table 3.1 — some of these channels are interesting enough to warrant their own names and this is

indicated in the table. The significance of each of these analysis channels will be made clear in section 3.4 — for now, let us just view these as signals that we are interested in and return to the discussion of non-statistical noise.

Waveform	Name
0	<i>Total</i>
S_E	<i>E shift</i>
S_B	<i>B shift</i>
$S_{\Delta B}$	<i>Cal</i>
$S_E \otimes S_B$	<i>EDM</i>
$S_E \otimes S_{\Delta B}$	<i>E Cal</i>
$S_B \otimes S_{\Delta B}$	<i>B Cal</i>
$S_E \otimes S_B \otimes S_{\Delta B}$	—

Table 3.1: The analysis channels.

It is through careful choice of the waveforms that the low-frequency non-statistical noise is rejected. Knowing the noise properties of the interferometer’s output, the fluctuations on the signals defined by (3.6) can be calculated. It is useful to divide the fluctuations on the interferometer’s output, rather arbitrarily, into two parts. Fluctuations at frequencies well above the frequency of blocks (about 0.5 Hz) will be called ‘noise’ — these fluctuations can be treated with statistical techniques. Other fluctuations, those at frequencies much lower than the phase-sensitive detection frequency, will be called ‘drift’ — a different approach will be required to deal with these. The division does have a vaguely physical nature, as often the low frequency fluctuations can be identified with a cause, such as a changing oven temperature or waning laser intensity.

A comprehensive analysis of the relationship between the fluctuation properties of the points and the resulting fluctuation properties of the analysis channels has been performed by Harrison, Player and Sandars [35]. The results will be reproduced here, with only the briefest outline of the proofs.

The first result concerns the fluctuations in the analysis channels arising from the ‘noise’ in the points. The approximate, but very useful result, derived by evaluating the Wiener-Khinchine transformation [36] for a simple model noise power spectrum, is that the noise in an analysis channel is simply equal to the noise in the points at the frequency of the fastest component square-wave of the analysis waveform. This is in agreement with one’s intuitive expectations — the same as the phase-sensitive detection result.

The second result concerns drifts in the points. It can be shown that, if the drifting output is expressed in terms of its derivatives (i.e. a Taylor series), that each component

square-wave in the analysis waveform will cancel out ‘one order’ of the drift. For example, it is straightforward to see that one component square-wave will remove the d.c. background. A few moments with a pencil and paper will show that adding a second square-wave will remove any drift that is linear in time. For higher orders the reader is referred to [35].

Now the choice of channel waveforms is clear. The analysis of the interferometer’s output in section 3.1 showed that for frequencies above ~ 10 Hz the noise was only a little worse than the shot-noise limit. The duration of a point has been chosen to be 48 ms, so the channel waveforms can have component square-waves at frequencies up to ~ 10 Hz. Any analysis channel that has a component at this frequency should show fluctuations very close to the shot-noise limit. This gate period also ensures the channel waveforms have no harmonics that are coincident with the peaks in the noise spectrum (figure 3.2). To remove the effect of drifts, it is clear that each channel waveform should have as many component waveforms as possible. The set of waveform codes used in the experiment is,

$$\begin{aligned} S_B &= 1111111011 , \\ S_{\Delta B} &= 1111101100 , \\ S_E &= 1100000000 . \end{aligned}$$

Some words are in order to justify the choice of S_E which seems to ignore the design constraints detailed above. One detail not mentioned so far, which will be expanded upon in the next section, is that every time the interferometer changes state some dead time is introduced to allow transient effects to die away. The dead time that needs to be introduced when the electric field direction is changed is much larger than the other two, consequently to keep the duty cycle high this switch should be performed infrequently. Re-examination of table 3.1 will show that the only analysis channel to be adversely affected by the ‘slow’ S_E is *E shift*. Keeping the duty cycle as high as possible is usually considered more important than having precise measurements of *E shift*.

3.3 The control software

Data collection is coordinated by EDM2. The primary function of the program is to administer the switching / measuring sequence. The program also performs limited on-line analysis and uses the analysis products to optimise the interferometer’s running parameters.

The heart of the program is the ‘acquire loop’⁶. The listing shows the structure of the

⁶There is a lot of code dealing with the user interface, plus code for some diagnostic features, that is not discussed here.

acquire loop in C-like pseudocode. When the acquire loop is started, EDM2 first reads in various parameters from the user interface (e.g. the rf transition frequencies, initial magnetic bias field, etc.) and uses these to initialise the interferometer's hardware. After opening a datafile and writing a header — a record of the various experimental parameters e.g. gate period, switching waveforms, etc — the acquire loop proper starts.

The inner loop in the listing is responsible for taking a single block. After setting the machine state appropriately, EDM2 checks which channels have been switched — each channel has a dead time associated with its switching and the longest of these dead times, out of the channels that have switched, is imposed. After the dead time, the LabPC+ counter is triggered and counts are collected for the gate period. This point is stored in an array and the loop repeats until the block is finished.

Listing 3.1: The acquire loop.

```

initialise hardware;
write datafile header;

for ( i = 0 ; i < numberOfBlocksRequired ; i++ )
{
    for ( k = 0 ; k < numberOfPointsPerBlock ; k++ )
    {

        set machine state as dictated by switching waveforms[k];
        impose appropriate dead time;

        read LIF detector and store point in array[k];

        check for emergency exit signal from user;

    }
    read fluxgate magnetometer;

    append block[i] to the datafile;

    analyse new data;
    calculate and apply magnetic bias field correction;
}

```

Immediately after finishing a block, EDM2 records the output of a fluxgate magnetometer, via a GPIB enabled HP33401A multimeter. The magnetometer is attached to the vacuum chamber near the upper rf loop, and sits between the inner and outer magnetic shields. These magnetometer readings can be used to analyse the data for correlations

with fluctuations in the ambient magnetic field. The block is then stored along with a short header that contains the magnetometer reading.

The block that has just been taken is then analysed. The data analysis is rudimentary — values for each of the analysis channels listed in table 3.1 are calculated. The analysis channel values for the most recent block and the mean of each channel over all of the blocks are displayed. These analysis channels provide a quick indicator of the interferometer’s health — a laser unlocking manifests itself as a sharp change in *Total* and the disappearance of *Cal*; a large *B Shift* indicates a magnetic field problem etc. (§3.4)

The dead times imposed after the interferometer’s state changes are 7 ms for the ΔB and B_0 -direction channels, this delay gives the magnetic field time to stabilise and time for the molecules to traverse the machine. The electric field direction channel has a much longer dead time, 4 s. During this time, the electric field supplies are switched off and allowed to bleed down to low voltage, the relays are switched, and then the supplies are re-energised. The duty cycle of the acquisition over a block is 60%.

The eight states shown in figure 3.4 miss one detail of the real experiment — a small magnetic field offset, B_{bias} is applied to all of the states. The purpose of this field, usually a few nT, is to cancel the non-zero residual magnetic field in the interferometer; it will be shown in section 3.6 that a residual field can lead to a systematic error in the EDM measurement. For reasons largely unknown, the residual magnetic field inside the interferometer is not constant in time.⁷ In order to keep the residual magnetic field as small as possible over the whole cluster, after taking each block of data EDM2 uses the analysis channel values to estimate the residual magnetic field and adjusts B_{bias} accordingly. The residual magnetic field is estimated by taking the combination⁸,

$$B_{residual} = \frac{B\ shift}{Cal} \frac{\Delta B}{2}. \quad (3.7)$$

The adjustment to B_{bias} is chosen empirically to be 1/4 of the estimate of $B_{residual}$. As will be seen in section 3.5, this scheme is very effective at keeping the residual magnetic field near zero. A record of the changing bias field is kept by recording the bias field for each block in the block headers.

The implementation of EDM2 is on the whole uninteresting — lots of fighting with Microsoft’s Foundation Classes — but I think one point is worthy of mention. All of the experiment’s timing is controlled by the computer, bar the LIF detector gate period, which is timed by the LabPC+ card. It was found that by bundling the acquire loop into a separate thread and giving this thread a very high priority, the timing sequence could be

⁷No correlations are observed with the magnetometer readings, indicating that the source is internal to the machine. As usual, the oven is the prime suspect . . .

⁸This expression may seem a bit mysterious — it should make more sense after section 3.4 has described the various analysis channels.

made surprisingly accurate (jitter < 0.5 ms)⁹, even when the computer was running other tasks¹⁰. Data analysis and user interface tasks are handled by threads with lower priority than the acquire loop. It is noted that, even with the modest computer used, *much* more on-line data analysis could be carried out without impairing the acquire loop timing.

3.4 Data analysis

The data analysis can be divided into two stages: the first of these is extracting analysis channel values from the data; the second is using these analysis channels to calculate the EDM and the various diagnostics. All analysis is done per block — the analysis channel values are extracted from each block and are then used to calculate an EDM for each block. These block EDMs are then combined to give the final result.

Before presenting the details of the data analysis some time should be spent describing the product of the first stage, the analysis channels.

Analysis channels

The analysis channels encapsulate all of the useful information that can be extracted from the data.¹¹ Figure 3.6 labels the eight interferometer states. The interference lineshapes are shown having different intensities and backgrounds for the opposing electric field directions — they are shown this way to facilitate the discussion of the diagnostic information that the analysis channels provide. The information that each analysis channel provides is listed below.

$$\mathbf{Total} = \frac{1}{8}(A + B + C + D + E + F + G + H).$$

This is simply the average LIF count rate for the block. Rapid changes in this channel from one block to the next usually indicate laser problems.

$$\mathbf{E\ shift} = \frac{1}{8}(A + B + C + D - E - F - G - H).$$

This channel measures the average difference in the count rate between the two electric field directions. A large value for the *E shift* can be indicative of a problem with the electric field plates. This channel is also sensitive to a difference in the interference lineshape intensity between the electric field directions, however, due to the large uncertainty in this channel (recall the discussion of section 3.2) this effect is never observed.

⁹As mentioned before this jitter has a strong 3.1 Hz frequency component. With better driver software I think the timing could be stabilised to the μ s level — experiments with a different data acquisition board support this feeling.

¹⁰except Mathematica !

¹¹The data does contain additional information about things with time scales less than one block (~ 1 minute), but this is very rarely useful.

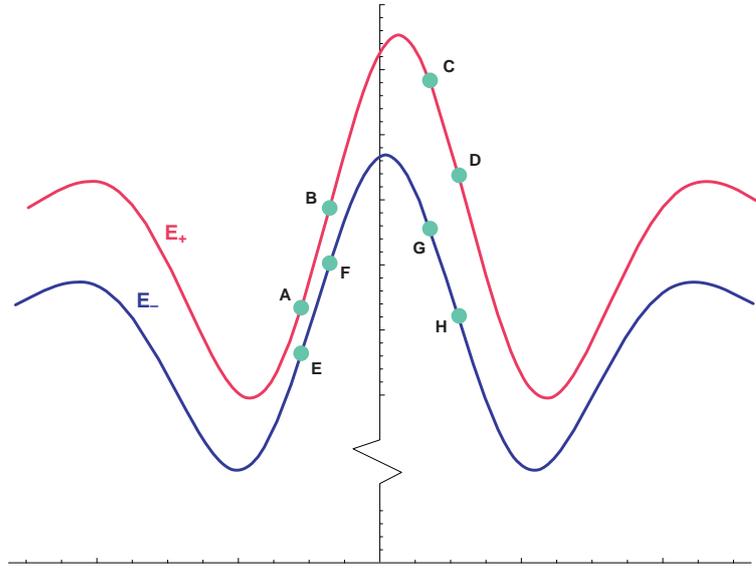


Figure 3.6: Labels for the eight interferometer states.

$$\mathbf{B\ shift} = \frac{1}{8}(A + B - C - D + E + F - G - H).$$

This channel measures the average difference in the count rate between the two magnetic field directions. A non-zero value for this analysis channel results from a non-zero residual magnetic field in the interferometer — it is this channel that is used by EDM2 to track the drifting residual field. A *B shift* could also be generated by an asymmetric interference lineshape but there is no evidence that this actually happens.

$$\mathbf{Cal} = \frac{1}{8}(A - B - C + D + E - F - G + H).$$

This is one of the most important analysis channels. It measures the average slope of the interference lineshape. This is a very useful diagnostic — as equation (3.2) shows, the EDM uncertainty is inversely proportional to *Cal* — and a lot of time is spent in the lab “making *Cal* bigger”. The value of this channel is also needed to extract quantities in physical units; both the EDM and EDM2’s estimate of $B_{residual}$ are normalised to *Cal*.

$$\mathbf{EDM} = \frac{1}{8}(-A - B + C + D + E + F - G - H).$$

This channel, which measures the change in count rate that depends on the relative orientation of the electric and magnetic fields, but not the absolute direction of either, is definitely the most important ! An EDM will give this channel a non-zero value. There *are* some other effects that can make this channel non-zero and these will be discussed in great detail in section 3.6.

$$E\text{Cal} = \frac{1}{8}(A - B - C + D - E + F + G - H).$$

This channel is a measure of the change in the slope of the interference lineshape when the electric field direction is reversed. This provides a sensitive measure of any change in the interference lineshape intensity as the electric field is reversed.

$$B\text{Cal} = \frac{1}{8}(-A + B - C + D - E + F - G + H).$$

This channel measures the difference in the slope of the interference lineshape between the two magnetic field directions. In the absence of an *E shift* or *E Cal* this channel measures the asymmetry of the interference lineshape.

Extracting the analysis channels

The first step in extracting the analysis channels is to separate the points from the block into eight ‘sub-blocks’, each sub-block corresponding to a particular state of the interferometer. These sub-blocks are cleaned to remove rogue points, corresponding to dust falling through the probe laser beam, bursts of electrical noise etc. The cleaning procedure is as follows: the mean and standard deviation of the sub-block are calculated and any points that lie more than 3.5 standard deviations away from the mean are removed. If the sub-blocks have a Gaussian¹² distribution, on average much less than one point per sub-block ($\frac{1024}{8} \times p(x > 3.5\sigma) \simeq 1/20$) should be lost to this cut. It is found that, typically, around 1/8 of a point is lost per sub-block — this excess corresponds to the rogue points.

The mean of each of the eight sub-blocks is then calculated — the appropriate linear combinations of these means are the analysis channels. Extracting the uncertainty on the analysis channels is a little more involved. It is assumed that *within each sub-block* the points are uncorrelated — if the switching scheme is doing what it promised this should be approximately the case. It is certainly not the case that points of different sub-blocks are uncorrelated though — these points are correlated through the action of the interferometer. To extract the uncertainty on the analysis channels, then, it is necessary to calculate the covariance matrix between the eight sub-blocks and use the standard result for linear combinations of correlated variables to transform this into the covariance matrix for the analysis channels [37]. The uncertainties on the analysis channel means are readily extracted from this transformed covariance matrix.

Calculating physical quantities

The next step in the analysis is to take the values for the analysis channels extracted from the data and combine these into the physical quantities of interest. It is rarely interesting

¹²The count rate is high enough that the distribution is close to Gaussian.

to convert the various diagnostic channels to physical units, with the exception of the residual magnetic field in the interferometer, equation (3.7).

Extracting the EDM from the analysis channels is also straightforward. Assuming the interferometer's lineshape is linear around the eight states, the ratio of *EDM* and *Cal* gives the ratio of the electric and magnetic phases induced by the electric field E_{eff} and the magnetic field $\frac{\Delta B}{2}$ — the factor of two accounts for the fact that the 'calibration' magnetic field, added to the bias field B_0 , switches between 0 and ΔB , whereas the electric field switches from E_{eff} to $-E_{eff}$. This leads directly to the expression for the EDM¹³,

$$\frac{-2 \eta d_e E_{eff}}{\mu_B \Delta B} = \frac{EDM}{Cal} \quad \Rightarrow \quad d_e = -\frac{\mu_B \Delta B}{2 \eta E_{eff}} \cdot \frac{EDM}{Cal}. \quad (3.8)$$

Using this expression, the EDM and its uncertainty are calculated for each block.

The final step in the analysis is to take a weighted sum of the individual block EDMs to form the final result.

Analysis software

I have written several Mathematica packages to analyse the EDM data. After a cluster is collected, the first task is to check it by hand, to ensure that everything was operating correctly throughout the entire cluster — this is performed by a package which analyses the data and displays the analysis channel values graphically. This package can be used to identify any blocks that should be removed from the analysis, for instance blocks taken after the laser has unlocked. Samples of this package's output will be shown in section 3.5.

After being checked, clusters are catalogued in a database, again implemented as a Mathematica package. The database records all of the experimental parameters for each cluster, as well as any blocks that are to be excluded and a comment. The database can call the analysis packages and store the analysed channels and the derived EDMs in the database. This ability of the database to automatically analyse the clusters is very useful — this was the reason for writing the database in Mathematica — changes can be made to the analysis packages and then the database told to re-analyse itself unattended. This is in marked contrast to the old scheme where the database was not able to call analysis routines — changing the analysis routines used to be a chore.

A final package queries the database to extract the final EDM and various statistics. Results from this package will be shown in section 3.7.

3.5 Some real clusters

This section should help to make the preceding sections more concrete. A few clusters will be examined to illustrate the points made above. Table 3.2 shows some statistics for

¹³The minus sign accounts for the minus sign in (1.5).

Oven operating current	160 A
LIF background rate	70 kHz
LIF signal rate (centre Q(0) manifold)	60 kHz
rf transition intensity	5.5 kHz
Interference lineshape intensity (peak to trough)	3.5 kHz
B_0	12.16 nT
ΔB	3.89 nT
Number of clusters	6
Number of blocks recorded	353
Number of blocks excluded (see text)	99
Total included blocks	254
Total integration time	3.5 hours
Day's EDM	$3.2 \pm 8.1 \times 10^{-26}$ e.cm

Table 3.2: A typical day of taking EDM data (taken from the 3rd May 2001).

a typical day (3rd May 2001). All of the clusters that will be shown were taken under similar conditions.

Laser unlocking

Figure 3.7 shows the laser unlocking (cluster 03May0101). The figure shows the values of the *Total* and *Cal* channels (the y -axes) for each block in the cluster (the x -axes). In this and all subsequent figures showing analysis channel values, the values are measured in Hz i.e. photon counts per second. The laser is seen to unlock in block 42 — the *Total* drops by 40 kHz and *Cal* drops from a healthy -500 to zero.

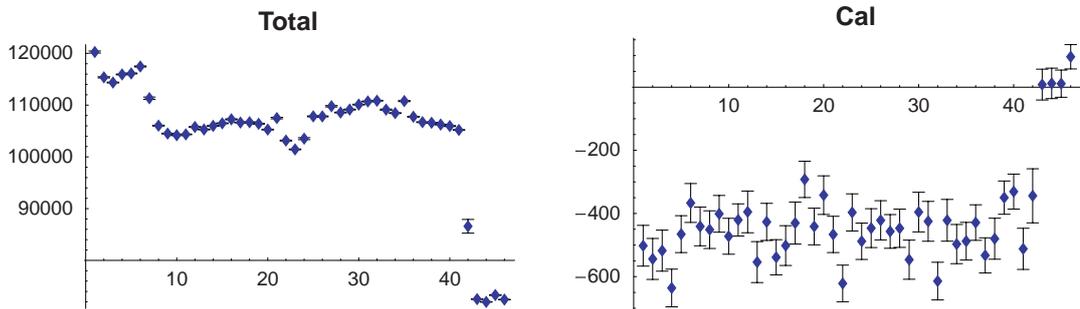


Figure 3.7: The laser unlocking.

B_{bias} tracking

Figure 3.8 shows the residual magnetic field tracking in operation (cluster 06Jun0116). The cluster was started with B_{bias} set slightly incorrectly — this shows up as a large B shift in the first block. EDM2 reduces the B shift to zero by adjusting B_{bias} , stabilising in around 10 blocks.

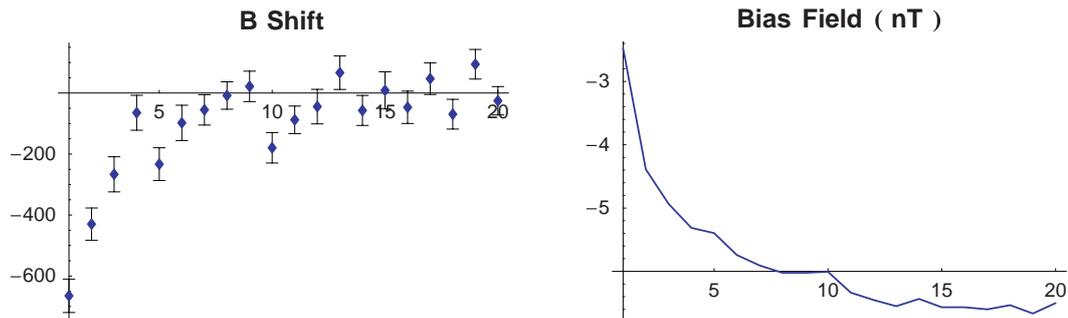


Figure 3.8: The magnetic field tracking in action.

Something broken

Table 3.3 shows the analysis channel values for the cluster of 95 blocks 03May0107. The EDM for this run is enormous, 5σ away from zero. The $E\mathit{Cal}$ channel is also significantly non-zero — this indicates that the interference pattern lineshape is changing as the electric field direction is reversed. A close inspection of the interferometer after this cluster was taken revealed a broken wire in the high voltage relay cabinet — in one electric field state, one of the C-region field plates was disconnected from the supply. This was responsible for the changing interference lineshapes and hence the non-zero $E\mathit{Cal}$ and EDM . This is the reason why table 3.2 shows so many blocks excluded from the analysis — 95 are accounted for by this cluster which had to be discarded. Normally only a few blocks are rejected each day.

3.6 Systematic effects

It is hoped that the reader will not be too disappointed to find, contrary to the precedent set by most EDM theses, that this section does not offer a cornucopia of systematic effects. There is a good reason for this — the EDM measurement scheme described in this thesis is inherently insensitive to many common systematic effects that have troubled other EDM experiments. The first part of this section will discuss this insensitivity.

The measurement scheme is not invulnerable to systematic problems, however. The

Analysis channel	Value (Hz)
<i>Total</i>	112×10^3
<i>E shift</i>	-163 ± 7
<i>B shift</i>	-7 ± 5
<i>Cal</i>	-353 ± 6
<i>EDM</i>	29 ± 6
<i>ECal</i>	21 ± 7
<i>BCal</i>	6 ± 6

Table 3.3: Analysis channel values for cluster 03May0107.

second part of this section will describe the two systematic effects that have been encountered so far and how they have been counteracted.

3.6.1 Magnetic effects

One of the great things about YbF is that it is relatively insensitive to magnetic fields, compared to atomic systems. This can be quantified by examining the ratio of magnetic to EDM interaction energy,

$$R = \frac{\mu_B B}{d_e E_{eff}} .$$

YbF has a much larger effective field than atomic systems; consequently, a magnetic field has to be correspondingly larger to be troublesome. Comparing the YbF effective field, $E_{eff} = 25 \text{ GV} \cdot \text{cm}^{-1}$ to that in the Berkeley Tl experiment, $E_{eff} \simeq 70 \text{ MV} \cdot \text{cm}^{-1}$ shows that YbF is several hundred times less sensitive to magnetic field effects. This reduced sensitivity to magnetic fields makes the experiment much less technically challenging.¹⁴

The YbF molecule has another feature that makes it an attractive environment for an EDM experiment. In an electric field, the highly polar YbF molecules are strongly aligned along the field direction. Because of this strong alignment, the response of the unpaired, interesting electron to magnetic fields is highly anisotropic. In short, the electrons interaction with magnetic fields perpendicular to the electric field direction is strongly suppressed. The systematic effect that has for years limited the most sensitive EDM experiment [9], the motional magnetic field effect, is caused by a magnetic field perpendicular to the electric field direction, and in YbF this effect is very heavily attenuated.

The interaction of the molecule with a magnetic field is important enough to merit a more careful investigation. Consider a simplified Hamiltonian modelling the $F = 1$

¹⁴See the ‘‘mired in this pit of misery’’ quote in the acknowledgements of [9] !

molecules' interaction with a magnetic field,

$$\hat{H} = \hat{H}_0 + \mu_B(B_z \hat{\sigma}_z + B_r \hat{\sigma}_x),$$

in which \hat{H}_0 contains the zero-field molecular Hamiltonian and the Stark interaction Hamiltonian; σ_z and σ_x are the standard $F = 1$ spin matrices. Due to the cylindrical symmetry of the molecule the applied field has been written as the sum of an axial part, B_z , and a radial part, B_r . The eigenstates of \hat{H}_0 are taken to be the z-basis states¹⁵, $(|1, 1\rangle, |1, 0\rangle, |1, -1\rangle)$, with corresponding eigenvalues $(\Delta, 0, \Delta)$. The tensor Stark shift between the $|1, \pm 1\rangle$ and $|1, 0\rangle$ levels, Δ , is introduced as a phenomenological parameter (see figure 2.22), and the contribution of the molecular Hamiltonian to the energy is taken to be zero. It is straightforward to write the matrix elements of the Hamiltonian, \hat{H} , in the z-basis,

$$\hat{H} = \mu_B \begin{pmatrix} \frac{\Delta}{\mu_B} + B_z & \frac{B_r}{\sqrt{2}} & 0 \\ \frac{B_r}{\sqrt{2}} & 0 & \frac{B_r}{\sqrt{2}} \\ 0 & \frac{B_r}{\sqrt{2}} & \frac{\Delta}{\mu_B} - B_z \end{pmatrix}.$$

This Hamiltonian can be diagonalized exactly but the results are complex and rather opaque. Some limiting cases will be considered analytically, then the general case will be investigated both perturbatively and numerically.

In the limit where $\Delta \rightarrow 0$, i.e. without an applied electric field, the energies of the states reduce to

$$E(|\pm\rangle) = \pm \mu_B \sqrt{B_z^2 + B_r^2},$$

$$E(|0\rangle) = 0,$$

where $|+\rangle$, $|-\rangle$ and $|0\rangle$ denote the three eigenstates. The eigenstates are simply the z' -basis states, where z' is along the resultant magnetic field.

Taking the converse limit, where $\Delta \rightarrow \infty$, the energies are

$$E(|\pm\rangle) = \Delta \pm \mu_B B_z,$$

$$E(|0\rangle) = 0,$$

and the corresponding eigenstates are just the original z-basis states $|1, \pm 1\rangle$ and $|1, 0\rangle$. The electric field has effectively 'pinned' the molecule to the z-axis and it is unable to respond to the radial magnetic field. This is the essence of the suppression of the effect of magnetic fields perpendicular to the electric field axis.

¹⁵We only need concern ourselves with the $F = 1$ states — the other states are so energetically distant that they play no part.

The real situation in the interferometer's interaction region is such that $\Delta \gg \mu_B B_z > \mu_B B_r$ and therefore a perturbative calculation can be performed. The unperturbed Hamiltonian is taken to be,

$$\hat{H}' = \hat{H}_0 + \mu_B B_z \hat{\sigma}_z ,$$

with the perturbation,

$$\hat{H}'' = \mu_B B_r \hat{\sigma}_x .$$

The eigenstates of the unperturbed Hamiltonian are simply the z-basis states ($|1, \pm 1\rangle, |1, 0\rangle$) with energies ($E_{\pm}^0 = \Delta \pm \mu_B B_z, E_0^0 = 0$). The perturbed energies, to second order, are,

$$\begin{aligned} E(|\pm\rangle) &= \Delta \pm \mu_B B_z + \frac{|\langle \pm | \hat{H}'' | 0 \rangle|^2}{\Delta \pm \mu_B B_z} + \dots \\ &= \Delta \pm \mu_B B_z + \frac{(\mu_B B_r)^2}{2(\Delta \pm \mu_B B_z)} + \dots , \\ E(|0\rangle) &= 0 . \end{aligned}$$

It is noted that the first order perturbations are zero, and only the $\langle \pm | \hat{H}'' | 0 \rangle$ matrix elements of the second order perturbation are non-zero.

The interferometer is sensitive to the energy difference between the $|+\rangle$ and $|-\rangle$ states,

$$\begin{aligned} (E(|+\rangle) - E(|-\rangle)) &= 2\mu_B B_z + \frac{(\mu_B B_r)^2}{2} \left(\frac{1}{\Delta + \mu_B B_z} - \frac{1}{\Delta - \mu_B B_z} \right) \\ &\simeq 2\mu_B B_z - \frac{(\mu_B B_r)^2 \mu_B B_z}{\Delta^2} . \end{aligned}$$

Taking the derivative of this energy difference with respect to $\mu_B B_r$ gives us a measure of the 'suppression factor'. The suppression factor relates an applied radial magnetic field, expressed in energy units, to the induced energy difference between the $|+\rangle$ and $|-\rangle$ states i.e.

$$S = \frac{\partial}{\partial(\mu_B B_r)} (E(|+\rangle) - E(|-\rangle)) = (-) \frac{2\mu_B^2 B_r B_z}{\Delta^2} . \quad (3.9)$$

This is the key result, showing that the effect of radial magnetic fields is suppressed heavily by the Stark splitting Δ .

The result (3.9) is confirmed to hold very well by the numerical calculations if the radial field is small. A similar perturbative result can be derived for the case $\Delta \gg \mu_B B_r \gg \mu_B B_z$ — it is found that $S \propto 1/\Delta$ — which is again confirmed by the matrix diagonalisation. For axial and radial magnetic fields that are of similar magnitude the numerical calculation is the only approach.

The magnetic fields likely to be found in the interferometer will be discussed at length below, but to get some immediate idea of the size of the suppression factor consider the following typical parameters. At the C-region operating field of $8 \text{ kV}\cdot\text{cm}^{-1}$ the tensor

Stark splitting is about 6 MHz. The applied magnetic field in the z-direction is around 300 nT, and a radial field of 4 nT is a reasonable guess. Under these conditions, the suppression factor $S \simeq 10^{-9}$, a *very* strong suppression.

Now that we understand the effect of magnetic fields on the interferometer output, we should look at some specific systematic effects caused by magnetic fields.

Leakage current induced magnetic field

If a magnetic field is to have an effect that will show up in the *EDM* analysis channel, it must be somehow correlated with the electric field direction. Perhaps the most obvious way for this to happen is if some current flows between the electric field plates — this current will certainly be correlated with the electric field direction. Consider a worst case scenario: a leakage current which generates a magnetic field in the z-direction — it has just been shown that only a magnetic field along the electric field direction has any significant effect. This is rather unnatural as it corresponds to current circulating in the x–y plane rather than flowing from one plate to the other. In a most extreme model imagine that the leakage current flows around the edge of one of the plates. When viewed along the z-direction (i.e. along the electric field) the plates are about 8×60 cm. Such a current loop generates a field $B = 10 I$ fT near the beam, where I is the leakage current in nA. It was reported in section 2.7.3 that the leakage currents through the field plates are less than 10 nA — in this worst-case estimate such a current corresponds to a field of 100 fT. The conversion from magnetic field to false EDM is

$$d_e = 4.4 \times 10^{-30} B \text{ e} \cdot \text{cm} , \quad (3.10)$$

where B is the magnetic field in fT, giving a false EDM from leakage currents,

$$d_e^{\text{leakage}} < 4 \times 10^{-28} \text{ e} \cdot \text{cm} .$$

This false EDM is well below the current statistical uncertainty of the EDM measurement (§3.7). It would be unfortunate if the effect of plate leakage currents were as large as this though, because the experiment will reach a statistical sensitivity better than $10^{-28} \text{ e} \cdot \text{cm}$ soon. I think it is indeed unlikely that any real leakage current could flow in such a disastrous way and, in any case, the new field plate design should be less leaky. This makes me confident that leakage currents will not generate any systematic error in the EDM measurement at the $10^{-28} \text{ e} \cdot \text{cm}$ scale.

Motional magnetic field

In the rest frame of the molecules, the electric field, applied in the lab frame, has a small magnetic component,

$$\vec{B} = \frac{\vec{E} \times \vec{v}}{c^2} .$$

This small magnetic field perpendicular to the electric field is not by itself particularly troublesome. However, if the applied magnetic field, B_z , is not perfectly aligned with the electric field it has a component in the same direction as the motional magnetic field. As a result, the magnitude of the perpendicular magnetic field is correlated to the relative direction of the applied electric and magnetic fields — this looks just like an EDM. This motional magnetic field effect, although extremely serious for atomic beam EDM experiments, is insignificant in this experiment — as shown above, the interferometer’s response to perpendicular magnetic fields is highly suppressed. Taking the beam velocity to be $500 \text{ m}\cdot\text{s}^{-1}$ and the electric field $8 \text{ kV}\cdot\text{cm}^{-1}$ gives a motional magnetic field of 4.5 nT . A gross misalignment in the 20 nT applied field of 5° would give a perpendicular component of 2 nT . These magnetic field parameters can be shown, using (3.10) and (3.9), to give rise to a false EDM,

$$d_e^{\text{motional}} < 8 \times 10^{-33} \text{ e}\cdot\text{cm}.$$

Even for such a gross misalignment this effect is completely negligible.

3.6.2 Measured systematic effects

With the discussion of the insignificant systematic effects over, it’s time to discuss the more worrisome effects — those that are large enough to measure. Two systematic effects have been encountered, and both have been understood and controlled. They are described below. First, however, I will describe a useful technique for isolating the cause of systematic effects.

Manual reversals

The switching scheme implemented by EDM2 enables the part of the interferometer’s output correlated with any combination of the switching channels to be extracted from the data. Additional information about where these correlations originate can be extracted by introducing ‘manual reversals’ to the switching sequence. Two manual reversals are used: the output connections of the high voltage relays can be swapped, connecting the electric field plates to the ‘opposite’ supplies and relays; and the connections from the computer controlled current source to the B_z coil can be swapped, reversing the direction of the applied magnetic field. These manual reversals are carried out between clusters, typically between every cluster I make one manual reversal. An examination of the analysis channels for several clusters with different manual reversal states (‘manual states’) reveals whether the correlation is internal to the interferometer or not — if the correlation is internal, the signal found in the analysis channel in question will be correlated to the manual state; if the correlation is external it will not.

Manual reversals are invaluable in identifying systematic effects — some examples will be shown below. It should be noted that the manual state for each cluster must be recorded in the cluster database; to extract a value for the EDM from the database, the sign of each cluster EDM must be adjusted in accordance with the manual state.¹⁶

The interference lineshape intensity depends on the electric field direction

The first systematic effect was encountered in November 2000; the signal:noise ratio of the interferometer had been steadily increasing and it was noticed that the EDMs from each cluster were no longer consistent with zero. Table 3.4 shows the analysis results from a few clusters. These analysis channel values are derived by finding the mean and standard error in the mean of the analysis channel values for each block. This process is statistically dubious for the channels, such as *Total*, *Cal* etc. whose mean varies on a timescale comparable to the period of a cluster — the errors in these channels should be taken a little less seriously than those of the other channels.

The first thing to notice in the table is that the EDMs are significantly non-zero. Closer inspection shows that the sign of the EDM doesn't seem to be correlated with the E manual state, inspection of more clusters confirms this — this means the effect is probably caused by some part of the electric field hardware external to the interferometer. All of the clusters show large *B shifts* — this indicates that the residual magnetic field inside the interferometer is non-zero¹⁷. It is also noticeable that the clusters have significantly non-zero *ECals* which indicate a change in the interference lineshape intensity when the electric field is reversed. The most likely cause of this dependence is an imperfect reversal of the electric field, probably due to surface charging of the electric field plates. Finally *B**Cal* is non-zero for the clusters. This could indicate asymmetry in the lineshape but I believe it is was a result of the large residual magnetic field biasing the interferometer's data-taking states onto a non-linear part of the lineshape.

Considering the above observations, it was decided that the most likely cause of the systematic EDM was a simple 'geometrical' consequence of having both a residual magnetic field and a change in the interference lineshape intensity when the electric field is reversed. The effect is best described pictorially — see figure 3.9. The left side of the figure shows the interference lineshapes for the case were *ECal* is non-zero, but the *B shift* is zero; there is assumed to be no 'real' EDM interaction. The ΔB switching channel has been ignored for simplicity, giving four interferometer states — in terms of these states the EDM analysis channel is $(B + C - A - D)$. It is clear that if there is no *B shift* then the false EDM is zero, whatever the size of *ECal*. The right hand side of the figure

¹⁶The two states of each manual reversal are labelled + and -. The signs are chosen such that the EDM extracted by (3.8) has its sign adjusted by the product of the two manual state signs to give the correct sign.

¹⁷The residual magnetic field tracking wasn't implemented at that time — it was introduced in response to this problem.

Cluster	Manual state		Analysis channel (Hz)							
	E	B	Total	E shift	B shift	Cal	EDM	E Cal	B Cal	
22Nov0002	-	-	72k ± 350	-25 ± 20	141 ± 7	-247 ± 9	-19 ± 5	36 ± 6	44 ± 5	
24Nov0003	-	-	72k ± 275	-147 ± 59	217 ± 13	-278 ± 9	-37 ± 10	44 ± 9	54 ± 10	
24Nov0004	+	-	72k ± 640	-34 ± 83	193 ± 23	-187 ± 21	-21 ± 13	18 ± 12	33 ± 11	
24Nov0007	+	-	85k ± 189	-106 ± 34	-86 ± 13	-307 ± 13	19 ± 7	24 ± 8	5 ± 8	

Table 3.4: Some clusters from November 2000.

shows the case where both E_{Cal} and B shift are non-zero — the origin of the false EDM is blatant.

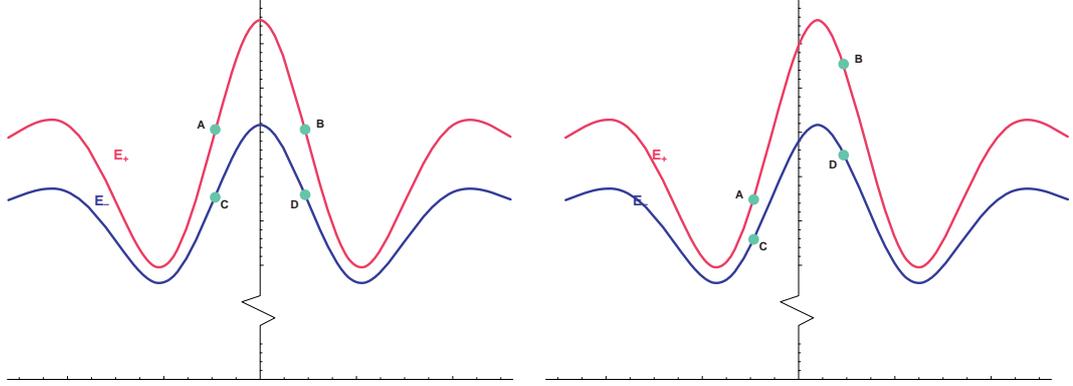


Figure 3.9: Non-zero E_{Cal} and B shift give an EDM .

It is a simple job to calculate the size of the false EDM generated, given the size of the B shift and the change in interference lineshape intensity when the electric field is reversed. If the ratio of the interference lineshape intensities is γ then it can be shown that the false EDM will be given by,

$$EDM^{false} = \frac{1 - \gamma}{1 + \gamma} B \text{ shift} . \quad (3.11)$$

This relationship holds for any interference lineshape, in the derivation it is only assumed that the lineshape does not change ‘shape’ when the electric field is reversed. It turns out to be very easy to extract γ from the data — it is given by,

$$\frac{1 - \gamma}{1 + \gamma} = \frac{E_{Cal}}{Cal} .$$

EDM^{false} can be calculated for each block of a cluster using the above formulae. The measured EDM for each block can then be corrected by subtracting EDM^{false} . The result of carrying out this procedure on the clusters in table 3.4 is shown in table 3.5 — the table shows the mean of the EDM analysis channel over all of the blocks of the cluster, with and without the correction; EDM^* is the corrected EDM . The table shows that the correction is very effective — it also shows the price to be paid for making a correction, an increased uncertainty.

The residual magnetic field tracking system described in section 3.3 was implemented in order to minimise this systematic effect. Reducing the B shift minimises the systematic EDM^{false} . All of the data included in the final analysis was taken with the magnetic field tracking active. The correction (3.11) is still included in the final analysis — the only correction made to the EDM — but its impact is minimal. The correction is small and the uncertainty in the EDM is only increased by 6%.

Cluster	<i>EDM</i>	<i>EDM</i> *
22Nov0002	-19 ± 5	4 ± 7
24Nov0003	-37 ± 10	-3 ± 13
24Nov0004	-21 ± 13	-19 ± 31
24Nov0007	19 ± 7	13 ± 8

Table 3.5: The corrected *EDMs*.*Electric field relay magnetic field*

The second systematic showed up in February 2001, after a two month period of not taking any data. In contrast to the previous effect, this arrived all of a sudden — clearly something had changed while the interferometer was off-line. Table 3.6 shows some clusters.

Again, the first thing to notice are the non-zero *EDM**s. The manual reversals are very informative — they almost uniquely identify the cause of the effect. The signs of the *EDM**s are correlated to the B manual state, as a real *EDM** should be, but not with the E state. The effect must therefore be caused by some part of the electric field system outside the vacuum chamber. The correlation between the B machine state and the false *EDM** indicates that whatever is causing the false *EDM** is coupling to the interferometer magnetically.¹⁸

After these clusters were taken, the first thing that was checked was whether the electric field switching system produced a magnetic field at the interferometer that depended on the direction of the switch. A fluxgate magnetometer inserted between the inner and outer magnetic shields revealed that there was. Repeating the test with the electric field supplies turned off revealed that the high voltage relay controllers were the source of the problem — they were generating a field of 4 nT at the magnetometer, correlated with the switch direction. I don't really know why this magnetic field appeared during the interferometer's offline period — perhaps changing the position of equipment near the relay controllers changed the magnetic field in the machine.

Investigation of the relay controllers revealed that their current consumption depended on the relay state. The controllers were redesigned such that they draw the same current

¹⁸There are some other, more unlikely, ways for an effect to generate this signature under manual reversals, but they all involve a direct coupling between the electric and magnetic field generating systems — something these system were carefully designed to eliminate.

Cluster	Manual state		Analysis channel (Hz)						
	E	B	Total	E shift	B shift	Cal	EDM*	E Cal	B Cal
02Feb0104	+	-	109k ± 476	227 ± 149	69 ± 23	-279 ± 12	23 ± 11	-4 ± 17	36 ± 11
02Feb0105	-	-	121k ± 899	-212 ± 171	277 ± 25	-518 ± 20	13 ± 11	59 ± 16	-9 ± 15
02Feb0106	-	+	128k ± 369	-54 ± 123	-51 ± 22	-392 ± 18	-34 ± 11	59 ± 16	-9 ± 15
02Feb0108	+	+	114k ± 1.12k	-146 ± 241	26 ± 23	-385 ± 24	-43 ± 17	-22 ± 15	17 ± 15
02Feb0109	+	+	114k ± 674	85 ± 85	-42 ± 17	-299 ± 21	-38 ± 13	-20 ± 16	6 ± 11

Table 3.6: Some clusters from February 2001.

regardless of the phase of the relay. The high voltage rack which houses the high voltage supplies, the relays and their controllers was also moved much further from the beam machine, to a corner of the lab. Repeating the measurements with the fluxgate magnetometer showed that there was no longer any detectable magnetic field correlated with the direction of the electric field switch. Data taken after the relay controllers were redesigned show no sign of this systematic effect. Only data taken with the new controllers are included in the final analysis.

3.7 The final analysis

Data was taken in earnest over the period February 2001 to May 2001. During this period, 10 days of running yielded 58 clusters — a total of 2089 blocks. The final analysis includes 1758 of these blocks, corresponding to 23.5 hours of integration. The majority of the rejected blocks were from 7 clusters that had to be removed entirely, because of problems with the interferometer such as broken wires, defective relays etc. The data run was ended as the electric field plates were becoming dirty, and the leakage current through them was starting to increase.

Analysis channels

Table 3.7 shows the analysis channel values for the entire dataset. The values shown are the mean and standard error in the mean of the analysis channels with the average taken over all of the clusters — the caveat of section 3.6 still stands: the meaning of standard errors in channels whose means are not constant for each cluster are not well defined.¹⁹ The upper half of the table shows the analysis channels, averaged over the clusters in the dataset with a given manual state. The lower half shows linear combinations of the upper values that show which parts of the analysis channel values are correlated with the manual state. The uncertainties shown on the linear combinations are calculated assuming that the analysis channel values for the machine states are uncorrelated, and as such can be considered a worst case uncertainty.

Table 3.7 looks very good. Discounting the *Total*, *E shift*, *B shift* and *Cal* channels, which are not constant from cluster to cluster, all of the analysis channels are consistent with zero except for perhaps *B_{Cal}* in the E–B– manual state — I expect that this value is just a fluctuation and not a real effect, but only more data can give a definite answer. The analysis channels show no correlation with the machine state, except for the slight correlation introduced by the large *B_{Cal}* in the E–B– machine state. The lack of any real features in the analysis channels makes me confident that the interferometer is operating

¹⁹Which isn't to say they're not useful — they do provide a guide to the 'spread' in the values, which is why they are presented.

Manual state		Analysis channel (Hz)						
E	B	Total	E shift	B shift	Cal	EDM	E Cal	B Cal
+	+	119k ± 9k	8.8 ± 30.1	-20.0 ± 11.6	-305 ± 27	2.18 ± 2.69	1.34 ± 2.82	-1.08 ± 3.57
+	-	118k ± 7k	5.6 ± 41.1	19.3 ± 11.8	-283 ± 25	0.71 ± 3.32	-1.07 ± 3.55	-1.08 ± 3.67
-	+	122k ± 8k	-22.4 ± 70.2	-30.8 ± 12.9	-347 ± 28	0.61 ± 4.04	3.12 ± 4.52	-0.11 ± 3.57
-	-	133k ± 9k	-70.4 ± 50.4	0.80 ± 8.49	-276 ± 23	-0.04 ± 3.16	3.67 ± 4.31	-10.3 ± 4.5

Correlated with :		Analysis channel (Hz)						
Nothing		123k	-19.6	7.66	-303	-0.87	1.61	-3.15
E		-4.68k	26.8	7.32	8.77	0.58	-1.48	2.07
B		-2.58k	12.8	-17.7	-22.9	0.53	0.61	2.55
E·B		3.25k	-11.2	-1.92	12.4	0.21	0.59	-2.55
		± 4.07k	± 25.1	± 5.66	± 12.92	± 1.67	± 1.93	± 1.92

Table 3.7: Analysis channel values for the entire dataset.

correctly, and that its operation is well understood.

Per-block statistics

The goal of the switching scheme is to eliminate all noise from the analysis channels other than the unavoidable shot noise. A measure of how well this is working can be made by examining the distribution of the EDMs over the blocks — if the analysis channels are truly shot noise limited each block’s EDM should be drawn from a Gaussian²⁰ distribution. To test this, chi-squared was calculated for the block EDMs normalized to their uncertainties,

$$\chi^2 = \sum_i \frac{(d_{e(i)} - \bar{d}_e)^2}{\sigma_{(i)}^2}.$$

Over the 1758 included blocks, χ^2 sums to 1945. The probability of $\chi^2 \geq 1945$ occurring by chance is only 0.1%. This sounds quite terrible, but it should be noted that with such a large number of blocks only a small change in the $\sigma_{(i)}^2$ ’s is needed to reduce the χ^2 to an acceptable level — they would need to be inflated by 6%. Of course, there is no reason to inflate the $\sigma_{(i)}^2$ ’s; they are correct. The reason for the large χ^2 is that the EDMs aren’t distributed in a Gaussian way. The origin of this non-Gaussian distribution is the low frequency ‘drift’ in the interferometer’s output; even though the switching scheme is very effective at removing the effect of these drifts, some residual effects persist. Monte Carlo simulations show that even a 1% sinusoidal modulation of the interferometer’s output signal with a period of several tens of blocks is sufficient to produce block EDM distributions that are non-Gaussian to a similar extent to the measured dataset. Such drifts in the interferometer’s output are common and are largely due to the oven temperature varying.

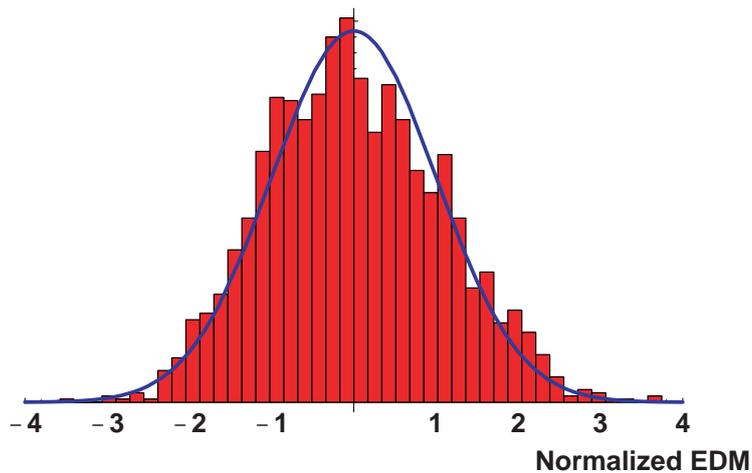


Figure 3.10: A histogram of the block EDMs.

²⁰As the count rate is high.

The non-Gaussian distribution is not such a worry. The drift in the interferometer's output shows no sign of being correlated to the interferometer's state and therefore does not bias the EDM measurement. The uncertainty on the EDM measurement will still be quoted as the 1σ standard deviation, but the reader should be aware that the EDM distribution differs from a Gaussian distribution by a few percent.

The distribution of block EDM values is shown in figure 3.10 as a histogram. The blue curve is a unity width Gaussian distribution, normalized to the total number of blocks.

Per-cluster statistics

Statistics can also be calculated at the cluster level. There is not much to be learned that is not already known from the earlier statistics but figure 3.11 is presented for reasons of tradition. It shows the EDM for each of the 51 clusters in the order that they were measured. Note that each cluster does not have the same number of blocks; this is responsible for most of the variation in the size of the error bars.

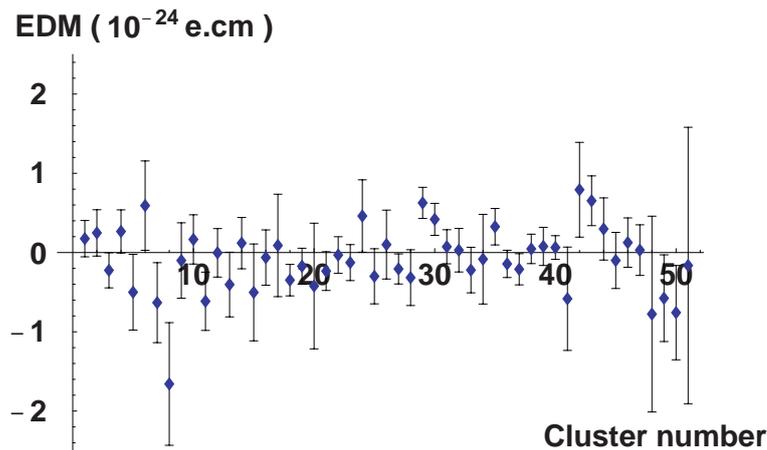


Figure 3.11: EDMs of the 51 clusters.

Result

Taking a weighted sum of the cluster EDMs shown in figure 3.11 gives the final result,

$$d_e = (0.31 \pm 4.01) \times 10^{-26} \text{ e} \cdot \text{cm} .$$

The uncertainty on this measurement can be compared to that predicted using (3.2) — the uncertainty on the measurement is about 10% above the limit imposed by counting statistics. This is a convincing demonstration of the effectiveness of the switching scheme's ability to reject low frequency noise.

Chapter 4

Conclusions

In this chapter I will make some closing remarks. First I will indicate the significance of the results presented in this thesis. After that I will discuss the future prospects of the experiment.

4.1 Result ?

The result of chapter 3 is the most sensitive measurement of the electron EDM made in a molecule, and the second most sensitive measurement in any system. It is an order of magnitude better than the previous molecular limit. The result sets an upper limit on the EDM that is small enough to make a number of proposed extensions to the Standard Model rather uncomfortable. In this respect it is a significant result.

However, the result is a factor of 50 worse than the current upper limit. In some sense, it doesn't tell us anything we didn't already know. I don't think this is a reason to be glum. Rather, whilst the Tl experiment has reached the end of its life, the YbF experiment is just getting started. As shown in the previous chapter, the measurement is completely limited by statistical uncertainty — there are no systematic effects that are significant at the current level of sensitivity. Moreover, calculations of all the known systematic effects give values far below the current sensitivity, with the exception of the plate leakage effect which we hope to reduce with new electric field plates (see below). There is no apparent reason why, given a stronger signal, the measurement can't be made several orders of magnitude more sensitive. This highlights what should be done in the future: make the signal stronger.

4.2 Future prospects

There are a number of plans to increase the sensitivity of the interferometer. They include:

New electric field plates The current electric field plates are not ideal: they suffer from crud build-up on the dielectric connecting the C- and guard-regions, leading to leakage problems; they don't support a very high maximum field; they broaden the splitter rf transition; and they are no doubt responsible for the rf transition frequencies drifting in time. All of these suggest that they be redesigned and replaced. This work is under way and at the time of writing is almost complete. Aside from stabilising the operation of the rf regions, which will increase the experiment's overall duty cycle, it is hoped that the plates will support a much larger maximum field, offering an increase in EDM sensitivity of almost 2 through an improved polarisation of the molecule.

Blue transition Recall that molecules are currently detected by driving the green $A \leftarrow X$ transition and counting spontaneously emitted photons. It might be possible to drive a blue $D \leftarrow X$ transition instead. This would have a number of advantages: the molecular spectrum may be less congested around this transition, reducing the background count rate from overlapping molecular transitions; the oven produces much less black-body light in the blue part of the spectrum than the green — careful filtering could significantly reduce the background count rate due to black-body oven light; and the photomultiplier's photocathode efficiency for blue photons is a factor of 3 higher than for green, increasing the efficiency of the detector. A significant amount of molecular spectroscopy must be performed on the $D \leftarrow X$ transition before it can be employed in the interferometer. Laser systems are currently being prepared for just this task.

Raman transitions In the past, there have been attempts to use STIRAP (stimulated Raman adiabatic passage [38]) for the splitter and recombiner, instead of rf transitions. The idea is to use two, detuned, partially overlapped laser beams. As the molecules fly through the beams the intensities change in such a way that the molecules are adiabatically transferred between ground state hyperfine levels. The advantages of this scheme are: the interaction is confined to a small volume, as large as the laser beam's overlap with the molecular beam (cf. the rf radiation which couples to the whole machine); it is velocity independent, and therefore the theoretical maximum transfer efficiency is 100% (cf. the maximum rf transfer efficiency of 76%, figure 2.12). Confining the interactions to a small region should remove some of the problem currently experienced with transitions being driven in field gradients

— this should make the interferometer more stable. The STIRAP scheme offers a potential increase in signal of 3–4, through improved transfer efficiency. Previous attempts at using STIRAP have been unsuccessful as, due to poor signal:noise, it has proven impossible to optimise the many experimental parameters. The interferometer’s signal:noise has increased greatly since STIRAP was last investigated; I think another attempt might prove successful.

Supersonic beam It was explained in section 2.2 that only a tiny fraction (10^{-4}) of the YbF molecules leaving the oven are in a state that contributes to the interferometer signal; most of the molecules are in rotationally excited states. A widely used technique for producing rotationally cold molecular beams is to use a supersonic jet source. A high-pressure carrier gas is seeded with the molecules to be cooled and allowed to expand into vacuum. The molecules’ thermal and internal energy is converted to forward kinetic energy as the gas expands and accelerates to supersonic velocity. Using these techniques, beams of molecules with rotational temperatures of 0.5 K are routinely produced. Some preliminary work has been carried out in the lab by Mike Tarbutt — an Ar jet has been seeded with YbF molecules produced by laser ablating a solid Yb + AlF₃ target. The results are very encouraging; a rotationally luke-warm (30 K) beam has already been observed. A supersonic beam offers a potentially massive increase in signal:noise. The signal could increase by many orders of magnitude, and the background could be almost eliminated (recall the main source of background counts is the hot oven). I hold great hopes for this scheme.

BaF In the long term, it might prove advantageous to switch from YbF to BaF. Although BaF has a smaller enhancement factor than YbF (§1.2) and is not significantly easier to produce, it has one big advantage: BaF is known to have stable Rydberg states [39]. The existence of stable Rydberg states would allow a much more sensitive detection scheme to be implemented: two-step laser excitation of the molecules to a Rydberg state followed by field ionisation and ion-counting.

I have carried out a comprehensive study of the Rydberg states of YbF. Molecules were first excited to the A electronic state by a single mode dye laser resonant with a transition (P(40)) near the Q(0) transition used in the EDM experiment. The molecules in this intermediate, excited state were then further excited by a UV pulse from a frequency-doubled pulsed dye laser. A large range of UV photon energies around the energy required to ionise the excited molecules were scanned over but no Rydberg molecules were detected. Similar experiments on Yb readily yielded Rydberg spectra. My conclusion was that the Rydberg states of YbF are not long

lived enough to make the Rydberg detection scheme practicable.

The Rydberg detection scheme offers the possibility of near 100% efficiency and negligible background. Changing to Rydberg detection of BaF could increase the signal:noise by several orders of magnitude (the efficiency of the LIF detector is $\sim 10^{-3}$). However, it would require a very substantial investment in time and equipment — it is sensible to see how far YbF can be taken before making the change.

Most of these enhancements are short- to medium-term. If all of the above were implemented the interferometer would likely be sensitive to EDMs of order 10^{-30} e·cm. To progress beyond this level will require a much larger change. The most promising approach involves the rapidly developing field of cooling and trapping molecules. A cold, trapped sample of YbF or BaF, with a spin coherence time of 10 s ($\sim 10^4$ the present coherence time) could be used to make a very sensitive EDM measurement.

Bibliography

- [1] Jeanette Winterson. *Sexing the cherry*. Vintage, 1990.
- [2] Lawrence Sklar. *Physics and chance: philosophical issues in the foundations of statistical mechanics*. Cambridge University Press, 1993.
- [3] J.H. Christensen, J.W. Cronin, V.L. Fitch and R. Turlay. *Evidence for the 2π decay of the K_2^0 meson*. Phys. Rev. Lett. **13**, 138 (1964).
- [4] Y. Fukuda et al. (Super-Kamiokande collaboration). *Evidence for oscillation of atmospheric neutrinos*. Phys. Rev. Lett. **81**, 1562 (1998).
- [5] A.D. Sakharov. Pis'ma ZhETF **5**, 32 (1967). [Sov. Phys. JETP Lett. **5**, 24 (1967).]
- [6] Robert G. Sachs. *The physics of time reversal*. University of Chicago Press, 1987.
- [7] L. Wolfenstein. *Violation of time reversal invariance in K^0 decays*. Phys. Rev. Lett. **83**, 911 (1999).
- [8] I.B. Khriplovich and S.K. Lamoreaux. *CP violation without strangeness*. Springer, 1997.
- [9] Brian Christopher Regan. *A search for violation of time-reversal symmetry in atomic thallium*. Ph.D. thesis, University of California at Berkeley, 2001.
- [10] P.G. Harris et al. *New experimental limit on the electric dipole moment of the neutron*. Phys. Rev. Lett. **82**, 904 (1999).
- [11] D.F. Nelson et al. *Search for an electric dipole moment of the electron*. Phys. Rev. Lett. **2**, 492 (1959).
- [12] L.I. Schiff. *Measurability of nuclear electric dipole moments*. Phys. Rev. **132**, 2194 (1963).
- [13] P.G.H. Sandars. *The electric dipole moment of an atom*. Phys. Lett. **14**, 194 (1965).
- [14] P.G.H. Sandars. *Measurability of the proton electric dipole moment*. Phys. Rev. Lett. **19**, 1396 (1967).

- [15] O.P. Sushkov and V.V. Flambaum. Zh. Eksp. Theo. Fiz 75, 1208 (1978). [Sov. Phys. JETP **48**, 608 (1978).]
- [16] N.S. Mosyagin et al. *Electric dipole moment of the electron in the YbF molecule*. J. Phys. B: At. Mol. Opt. Phys. **31**, L763 (1998).
- [17] M.G. Kozlov et al. *Enhancement of the electric dipole moment of the electron in the BaF molecule*. Phys. Rev. A **56**, R3326 (1997).
- [18] Y.Y. Dmitriev et al. *Calculation of the spin-rotational Hamiltonian including P- and P,T-odd weak interaction terms for HgF and PbF molecules*. Phys. Lett. A **167**, 280 (1992).
- [19] E.D. Commins. *Electric dipole moments of leptons*. Advances in Atomic, Molecular, and Optical Physics **40**, 1 (1999).
- [20] E.D. Commins et al. *Improved experimental limit on the electric-dipole moment of the electron*. Phys. Rev. A **50**, 2960 (1994).
- [21] D.K. Cho, K. Sangster and E.A. Hinds. *Tenfold improvement of limits on T violation in thallium fluoride*. Phys. Rev. Lett. **63**, 2559 (1989).
- [22] Giles D. Redgrave. *Spin spectroscopy of YbF using molecular beam interferometry*. D.Phil. thesis, University of Sussex, 1998.
- [23] G. Herzberg. *The spectra & structure of simple free radicals: an introduction to molecular spectroscopy*. Cornell University Press, 1971.
- [24] H. Lefebvre-Brion and R.W. Field. *Perturbations in the spectra of diatomic molecules*. Academic Press, 1986.
- [25] B.E. Sauer, Jun Wang and E.A. Hinds. *Anomalous spin-rotation coupling in the $X^2\Sigma^+$ state of YbF*. Phys. Rev. Lett. **74**, 1554 (1995).
- [26] B.E. Sauer, Jun 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**, 17 (1996).
- [27] B.E. Sauer et al. *Perturbed hyperfine doubling in the $A^2\Pi_{1/2}$ and $[18.6]0.5$ states of YbF*. J. Chem. Phys. **110**, 8424 (1999).
- [28] K.L. Dunfield et al. *Laser spectroscopy of the lanthanide monofluorides: analysis of the $A^2\Pi-X^2\Sigma^+$ transition of ytterbium monofluoride*. J. Mol. Spec. **174**, 433 (1995).

- [29] R.F. Barrow and A.H. Chojnicki. *Analysis of the optical spectrum of gaseous ytterbium monofluoride*. J. Chem. Soc. **F2**, 71:7128 (1975).
- [30] W. Demtröder. *Laser spectroscopy: basic concepts and instrumentation*. Springer, 1996.
- [31] Rodney Loudon. *The quantum theory of light*. Oxford University Press, 1973.
- [32] Norman F. Ramsey. *Molecular beams*. Oxford University Press, 1956.
- [33] Jun Wang. *Laser and radiofrequency spectroscopy of ytterbium fluoride ground state*. Ph.D. thesis, Yale University, 1996.
- [34] C.M. Van Vliet and P.H. Handel. *A new transform theorem for stochastic processes with special application to counting statistics*. Physica **113A** (1982) 261-276.
- [35] G.E. Harrison, M.A. Player and P.G.H. Sandars. *A multichannel phase-sensitive detection method using orthogonal square waveforms*. J. Phys. E **4**, 750 (1971).
- [36] D.K.C. MacDonald. *Noise and fluctuations: an introduction*. Wiley, 1962.
- [37] R.J. Barlow. *Statistics*. Wiley, 1989.
- [38] K. Bergmann et al. *Coherent population transfer among quantum states of atoms and molecules*. Rev. Mod. Phys. **70**, 1003 (1998).
- [39] Z.J. Jakubek and R.W. Field. *Core-penetrating Rydberg series of BaF — s-similar-to-p-similar-to-d-similar-to-f supercomplexes*. Phys. Rev. Lett. **72**, 2167 (1994).