COLD ATOM GRAVITY GRADIOMETER FOR FIELD APPLICATIONS

by

Andrew Aneurin Lamb

A thesis submitted to
The University of Birmingham
for the degree of
DOCTOR OF PHILOSOPHY

Ultracold Atoms Group
School of Physics and Astronomy
College of Engineering and Physical Sciences
The University of Birmingham

January 2019

“There is a time for many words, and there is also a time for sleep."

- The Odyssey

….what does Homer know anyway?

Abstract

This thesis presents the first efficient implementation of a ^{87}Rb prism MOT gravity gradiometer. The resulting portable system (weight <50 kg and volume <0.1 m^3) is capable of producing atom numbers up to 5×10^7 with typical cloud temperatures of 7.8 μK. Combining with a robust control system and a 5 m delivery umbilical the effectiveness has been demonstrated in a non-laboratory environment.

A measurement campaign to detect a buried maintenance tunnel, with an expected signal size of 230 E, was performed during March 2018. An ambient temperature range of -0.5 to 10.2 °C and wind speeds in excess of 10 m/s provided adverse operating conditions that imposed no loss of functionality on the sensor. Resulting data shows a trend indicative of the tunnel with long term drift limiting fit error to 140 E.

A second campaign was performed to detect a signal of 280 E generated by a lead mass of 450 kg in a laboratory environment, with ambient temperature fluctuations of 0.5 °C recorded. Sensor drift of ~220 E/hour was seen making it impossible to discern the expected signal. Correlation between the Raman sideband-to-carrier amplitude ratio and measured gradient signal was identified and linked to ambient temperature fluctuations.

ACKNOWLEDGEMENTS

I’ve read my fair share of soppy acknowledgements sections and never thought I’d do the same, but the truth is you don’t understand what it’s like until you’ve been through the gruelling thesis writing process. It’s an emotional experience and there will always be so many people to thank, I hope I remember them all.

Thank you to Kai Bongs for giving me the opportunity to study here, to Dstl (specifically Stephen Till, Gareth Brown and Jonathan Pritchard) for providing the funding and the Quantum Technology Hub for providing the facilities and equipment. A deep thank you to Mike without whom I would have been adrift in a sea of physics, he was somehow always able to make time to talk through even the smallest problem despite his time being worth more than my weight in gold. To Aisha who has been an invaluable source of scientific wisdom and guidance, always available to talk through a physics problem or just chat when I needed a break. Jonathan “JoJo" Jones and Yu-Hung for their willingness to answer questions from the trivial to the complex. The rest of the Gravity Imager team past and present: guar-Anthony Rodgers, Ben, Jamie “8-inches" Vovrosh, Geoff, Kevin, Artur and Alex, without whom this system could not have been constructed. A big thank you to the talented members of the workshop team, particularly Steve, Dave and Warren, for pushing their manufacturing capabilities to the limit yet still producing high quality results and meeting challenging deadlines.

Clemens, Lingxiao and Andy who taught me the ropes when I was but a humble theorist making his forays into experimental physics and the whole UoB team for their ongoing support. Jonny who was always ready to talk, drink and when needed, hug, my problems away and the rest of “Pub Club" not mentioned above: Rustin, Georgina, Luuk, Hester, Dana. There are times that you need to talk and times you need to work but sometimes you just need a pint, thanks for not leaving me to drink alone.

Last but by no means least all those who have supported me from outside the group. My family: the

mother and father who raised me and the brother who supported me not to mention the grandfather who inspired me. My friends old and new: Holly, Connie, Nick, Lucy, Tash and Lynette you were there for what I needed through thick and thin. Which brings us finally to Ellie, you were there whilst you could be and helped while you were, for that I will be eternally grateful.

CONTENTS

1 Introduction **1**
- 1.1 State of the Art in Commercial Gravimetry . . . 3
  - 1.1.1 Scintrex CG-5 . . . 3
  - 1.1.2 GWR Instruments iGrav . . . 5
  - 1.1.3 Atom Interferometry . . . 6
- 1.2 Our Goal . . . 7
- 1.3 Sensitivity Requirement . . . 11
- 1.4 Thesis Breakdown and a Note on the Plotting Convention . . . 14
- 1.5 Statement of Contributions . . . 14

2 Building a Portable Quantum Device **16**
- 2.1 Magneto-Optical Trapping . . . 16
  - 2.1.1 Capture Velocity and Volume . . . 20
- 2.2 Prism MOT Prototype . . . 21
  - 2.2.1 Prism MOT Stability . . . 25
- 2.3 Prototype Gravimeter . . . 26
  - 2.3.1 Pressure Calculation . . . 28
  - 2.3.2 Cloud Temperature . . . 31
  - 2.3.3 Atom Interferometry . . . 34
- 2.4 Implementation in Gradiometry . . . 41
  - 2.4.1 Prototype Gravimeter Capture Volume . . . 42
  - 2.4.2 Gradiometer Geometry . . . 44

3 Sensor Package **49**
- 3.1 The Vacuum System . . . 53
  - 3.1.1 Prism Mount . . . 55
  - 3.1.2 Indium Sealing . . . 57
  - 3.1.3 Cleaning and Bakeout Procedure . . . 59
- 3.2 Magnetic Field Control . . . 60
  - 3.2.1 MOT Coils . . . 61
  - 3.2.2 Bias Coils . . . 62
- 3.3 Optical Delivery . . . 64
- 3.4 Integration . . . 67

4 Control Package 70
4.1 Cooling Laser 72
4.2 Raman Laser 74
4.3 Intensity Stabilisation 75
4.4 Computer Control 79

5 Experimental Procedure and Characterisation 80
5.1 Experimental Sequence 81
    5.1.1 Atom Preparation 81
    5.1.2 Raman Pulses 85
    5.1.3 Detection 88
5.2 Temporal Variation 90
    5.2.1 MOT Stability 91
    5.2.2 Frequency 94
    5.2.3 Raman Interactions 96
5.3 Chirp Rate 100
    5.3.1 Choice of Raman Beams 101
5.4 Spatial Variation 103
    5.4.1 Magnetic Field Profile 103
    5.4.2 Spatially Dependent Rabi Frequency 109

6 Measurement Campaigns 113
6.1 Systematic Offsets 113
6.2 Ellipse Analysis 116
    6.2.1 Fitting the Ellipse 116
    6.2.2 Phase Corrections 119
6.3 Tunnel Survey 123
    6.3.1 The Survey Site 123
    6.3.2 Methodology 129
    6.3.3 Results 132
6.4 Lead Detection 137
    6.4.1 Results 139

7 Conclusion and Outlook 147

Appendix A I

Bibliography II

LIST OF FIGURES

1.1 Two bodies of diameter d_1 and d_2 with uniform mass distributions, separated by a distance r. 1
1.2 The situation in Fig. 1.1 has been adapted for the case in which one of the bodies no longer has a uniform mass distribution and $d_2 \approx r > d_1$. 2
1.3 Scintrex CG-5 gravity sensor. 3
1.4 Left: A relative gravimeter in the case that the capacitor plates lie perpendicular to gravity (adapted from [2]). Right: Introducing a tilt θ between the spring and the normal to the capacitor plates measures only the signal projection of gravity $g \cos\theta$. 5
1.5 For an interferometer the interrogation light is input from one end and retro-reflected from a mirror at the base of the system. This leads to both atomic clouds being simultaneously interrogated by the same light. 8
1.6 The subject of this thesis, the portable gravity gradiometer constructed at the University of Birmingham in its intended environment. 9
1.7 The toy model for two gravimeters at points S_1 and S_2 above a buried point mass. 12

2.1 Simple energy level diagram for a two level MOT system. In an active MOT system the lasers L_1 and L_2 would have σ_+ and σ_- polarisations respectively. The frequency of the laser is ν_{laser} and Δ is the detuning from the atomic resonance. 17
2.2 (a) A traditional six-beam MOT with are individual delivery beams for all six co-ordinate axes from six independent delivery telescopes. (b) In a prism MOT a single delivery beam provides the light for all six capture directions. 19
2.3 Diagram of the prism MOT prototype. 22
2.4 Loading curve for the prism MOT prototype demonstrating 2.7×10^8 atoms required for sufficient sensitivity (Eq. 1.11). 23
2.5 Left: The prototype system housed in its transport cage, below the vacuum chamber a rack mounted fibre based laser (similar to that in 4.1) is mounted. Right: The first picture of the MOT produced in the prism system with a very satisfying red glow. 24
2.6 Position drift comparison between prism MOT and six-beam MOT. The data sets are both over approximately one hour going from blue to red in the six beam case and yellow to green for the prism MOT. The data for the six-beam MOT is reproduced from [26]. 25
2.7 Schematic of the gravimeter vacuum system. The ion pump and getter are now two separate components mounted directly onto the chamber which has been lengthened to allow for the drop. 26
2.8 Loading curve data (blue) for the prototype gravimeter. By fitting to this data (black) a longer loading time of 1.34 s and lower saturated atom number of 2.2×10^8 is measured when compared to the MOT prototype. 27
2.9 Gravimeter light sheet measurement. 33

2.10 Boltzmann distribution of a 2.82 μK cloud. The shaded region shows the atoms that would be transferred by a 50 μs Raman pulse. 33
2.11 Level structure of the D2 transition of ^{87}Rb (adapted from [65]). The cooling and repump beams are used in the MOT and optical molasses stages to cool the atoms and keep them out of the dark state. R_1 and R_2 are the two frequencies of light required to drive the two photon Raman transition. 36
2.12 Rabi oscillation for the prototype gravimeter during system tests outside of the lab. 37
2.13 The motion of the two atomic clouds in the absence of an external phase shift, in the ideal the final $|2, p+2\hbar k_{eff}\rangle$ state will be empty. 38
2.14 Optical Mach-Zehnder Interferometer, $\Delta\phi$ can be seen as analogous to gravity in the atomic case. 38
2.15 Comparison of fringes generated both inside (upper plot) and outside (lower plot) of the lab environment for a T-time of 1 ms. 40
2.16 The gravimeter prototype in its first foray outdoors to what has become the preliminary testing ground for all portable systems. 41
2.17 Gaussian intensity profile overlaid with the prisms in the relative locations used for the prism MOT. 43
2.18 Capture velocity with beam size comparing the real gravimeter prototype system to an idealised infinite intensity case. 44
2.19 Potential Gradiometer implementations. 45
2.20 Capture velocity versus capture volume for the designs shown in 2.19a, 2.19b and 2.19c . . . 46

3.1 Final implementation of the gradiometer, two separate sensor heads (colour coded as 1.4) coupled together by a common mechanical structure, common light input and a central double sided mirror. 50
3.2 Cross section of the Top sensor head, the region enclosed by the black dashed region represents the location of the prisms in the Bottom sensor head. 52
3.3 Two dispensers mounted onto a single four-pin power feedthrough with barrel connectors that have had a slit cut into them. Ceramic beads are used to ensure that the barrel connectors do not come into contact with the chamber wall, causing a short. 54
3.4 Left: Birds eye view of the chamber front showing the location of the positioning grooves relative to the prisms. Right: A cross sectional view of the Top science chamber where the depth of the positioning groove has been circled. 56
3.5 Left: Prism mount for the gravimeter system showing the parts used for mechanical alignment. Right: The tool used to lower the mount into position. 57
3.6 Left: The uncompressed indium ring under the window with a clear gap around its edge. Right: The compressed indium seal post bake, under compression it has spread into the chamber opening reducing the clear aperture. 58
3.7 Design and use of compression flanges for indium sealing windows to the chamber. 59
3.8 MOT magnetic field (3.8a) and field gradient (3.8b) at 3 A operating current. 62
3.9 Measured and simulated bias field profile of the gravimeter coil at a current of 0.2 A. 63
3.10 Diagram of the main experiment telescope. Both cooling and Raman beams are delivered to the telescope from the same direction and reflected to ensure that delivery into the chamber is near coaxial. 64
3.11 The chamber based optical delivery for the Top chamber 66

3.12 The fibre distribution system for the experiment light. The whole system, along with electrical distribution, fits into the base of the system with a volume of ~8 dm^3. 67
3.13 Left: The forty piece mounting structure before assembly. Right: The four modules the sensor package can be easily broken down into after assembly. 68
3.14 Pole and pockets pairings of the mounting structure, the end of each set of uprights features one of these features allowing for easy modification of the system modules. 68
3.15 The detailed CAD model of the system presents a fairly accurate representation of the final product, as can be seen here. 69

4.1 To use a biological analogy, the control package is the “brain" and the sensor package is the “body" with the umbilical acting as the nervous system. 71
4.2 Schematic breakdown of the cooling laser, this is all contained within a single 2 U rack unit with the exception of the spectroscopy and lock box. 72
4.3 Schematic of the spectroscopy box, contained in a single 2 U rack box with the MTS module contained within its own smaller enclosure. 74
4.4 Raman laser for the system, it is a modification of the same fibre laser used for the cooling light. 74
4.5 Applying a variable gain to an input diode signal allows control of the light level as measured at the telescope whilst also facilitating stabilisation. 76
4.6 MOT light level variation for both the case of stabilised and free running. Each of the datasets has been normalised separately to its mean in order to provide a better comparison. The few rogue points visible near the end of the stabilised data set arise due to the duty cycle of the device, a sequence is regularly resubmitted leading to a downtime and thermal changes to the AOMs in the lasers. 77
4.7 Drift of the Raman light level with time for both the stablised and free running cases. The rogue points in the stabilised data set are similar to those in Fig. 4.6b and likely arise due to the same thermal cycling of the AOMs. 78

5.1 Optical molasses cooling sequence and state preparation. 82
5.2 Effect of pre-molasses dark time on cloud at detection. 83
5.3 Scan of state preparation time, 0% corresponds to no atoms detected above the background pulse but 100% is just the maximum of the dataset. See 5.1.3 for the detection process used to gather these results. 85
5.4 Typical Rabi oscillation for a Raman pulse at t =68 ms. 87
5.5 The result of performing a Rabi oscillation at each of the three pulse locations. The increased damping effect arises from the thermal cloud spreading out over a larger distribution of Raman intensities. 88
5.6 Detection sequence for the system. The peaks are labelled by their position in the sequence, i.e. from left to right they take labels 1 – 6. 90
5.7 Top: The variation in cloud temperature with time with the central retro-reflection mirror. Bottom: The same result after the central mirror has been removed. There is a distinct increase in both the average temperature and the noise on the temperature in the case that the mirror has been removed, this is to be expected as variations in cooling light will no longer be common within a given module. 92
5.8 Variation in the arrival time. 93

5.9 (a) Measurement of the long term stability of the experimental lasers as performed with a wavelength metre that has a noise level of ~2 MHz (represented by the dashed black line). (b) The impact this frequency noise would have on the scattering force of the cooling laser as a fraction of the force at the standard experimental detuning of 13.2 MHz. 95
5.10 Variation of π-time without intensity stabilisation overlayed with the average value of I_6 for the 15 oscillation. There is a clear inverse correlation between the two, as expected. 97
5.11 (a) Shows the relative optical power of the sidebands for optimised IQ voltages (see Fig. 5.12 for the peak definitions). (b) Shows how these peak sizes vary over the course of 14 hours. . . 99
5.12 When an EOM is driven with modulation frequency ν_M it creates an infinite set of sidebands separated by ν_M. 100
5.13 Chirp rate scans performed on the two clouds, the true chirp rate $\alpha = -25.17 \times 10^3 + \delta\alpha$ kHz/s with the offset applied to make the data more legible. This is a historic data set (December 2017) and uses the older normalisation style as mentioned in 5.1.3. 101
5.14 Frequency shift depending on chirp rate choice 102
5.15 Simulated Zeeman splitting in 190 mG magnetic field an a Raman pulse of τ =5 μs. 104
5.16 Scan of the Raman sideband frequency difference after ~100 ms of free fall. The larger sets of triplets on the far left/right extremes of the plot arise from probing the Zeeman sensitive transitions with counter-propagating linearly polarised light. The central triplet is the same Zeeman sensitive transitions probed with the small amount of Doppler insensitive circularly polarised light present in the input beams. 106
5.17 Magnetic field profile inside the science chamber as measured by the atoms, the black dashed lines represent the locations of the Raman pulses used in the current experimental sequence. 107
5.18 Background field profile. 108
5.19 Rabi variation with height using an EOM, the cloud size is set by assuming an initial diameter of 2 mm and a velocity selective Raman pulse of 5 μs. 110
5.20 The results of Rabi oscillations for the two chambers first using and EOM to generate the Raman sideband and then and IQ modulator. 111
5.21 Rabi variation with height using IQ modulation and common input beams for both clouds. . 112

6.1 The evolution of the relative state populations with laser phase: 6.1a in the absence of any phase noise, 6.1b in the presense of 20% common phase noise. 6.1c the Lissajou figure produced by the two sinusoids demonstrating the effect of the common mode noise suppression.115
6.2 Effect of normalisation style on ellipses fitted to sets of 90 measurement shots over the course of 10 hours. 119
6.3 Left: Data normalised with Eq. 5.1 that shows a good correlation with the ellipse. Right: Data normalised to the average background for all the data points that shows a much larger spread and a clearly worse match to the data. 120
6.4 The site before setup on the first day of the survey. The red rectangle shows a region in which the distance to the surface drops to 30 cm, leading to an increase in the detected signal.124
6.5 Given the physical parameters of the anomaly and those of the sensor, the CoG relative to each chamber is labelled after integration of the signal across the anomaly, the integrated values of Eq. 6.17a and Eq. 6.17b are the dashed (blue) and solid (red) lines respectively. The top of the anomaly is located at Δz =0 m. The values of R_B and R_T include systematic offsets of the system: the thickness of the tarmac above the anomaly, the distance from the ground to the bottom sensor and the separation between the two sensors. 125

6.6 Sketch of the impact of taking a background measurement. Assuming everything else is constant it simply means there is a comparison between an anomaly with or without material present. . . . 126
6.7 An ellipse and fit representative of the state of the system approximately two weeks before the survey. . . . 128
6.8 Results of a preliminary survey of the site performed with a Scintrex CG-5, 0 m denotes the location of the centre of the anomaly. The linear component correction is the drift in the CG-5 measurement as referenced against a reference base station. . . . 129
6.9 The author providing indisputable proof that the sensor is person portable if appropriate care is taken. . . . 130
6.10 Weather data courtesy of the UK Met Office weather station at Coleshill. . . . 131
6.11 The gazebo which acted as home over the course of the survey. Through the sodden window it is possible to see the glow of the computer screen and in the background is what remained from the recent snow storm. A true “We’re not in Kansas anymore." moment. . . . 132
6.12 The tilt data from Survey 2, during the breaks in the data the sensor is being moved between measurement points. . . . 133
6.13 . . . 135
6.14 The phase data arising from Survey 1. To compare this result to Fig. 6.8 point 1 will be located at ~4.5 m, point 6 at ~0 m and point 11 at ~-4 m. The expected value is that determined using the CG-5 (Fig. 6.8) modified using Eq. 6.22. . . . 136
6.15 Left: The lead placed around the sensor in the “lead" position, the Top sensor is visible and the Bottom is below the optical tables. Right: Schematic of the experimental setup, placing lead at the midpoint of the two sensors it is possible to maximise the gradient signal by producing opposite phase shifts on the two clouds. . . . 138
6.16 The three orientations of lead shown would have one of three effects (from left to right): maximise the gradient signal, minimise the gradient signal or produce half maximum gradient signal on the ellipse. By modulating between these three regularly and in a random order it would allow the effect of the lead to be seen above any drifts the system may experience. . . . 139
6.17 The result of the two day lead modulation survey. Each of the phase value is the result of 1620 measurements and the sideband ratio is recorded after every 90, the drift in the phase shows a good correlation with the sideband ratio. . . . 140
6.18 The sideband ratio was manually changed whilst maintaining constant total Raman power. A 90 shot ellipse measurement was performed at each value, the result is a clear correlation to which a linear relationship was fit. . . . 141
6.19 The change of primary sideband ratio against time as compared to the temperature variation measured at the sensor head. Discretisation of the temperature data is due to the limited resolution of the temperature sensor and makes it impossible to see any rapid variations. A correlation is clear despite some short term deviations around the 2 hour and 20 hour marks. . . . 142
6.20 The variation of the primary sideband ratio over the course of the 48 hour lead survey. The simulated effect is clearly much smaller (roughly an order of magnitude) than the measured impact, however it can be seen that it follows the temperature variation closely. The lower plot has reduced the y-range of the other to better display the variation in the simulation. . . . 145

1 Complete breakdown of the RF and light distribution system from generation till exiting the flight case. I

LIST OF TABLES

2.1 List of materials used in the vacuum system along with the surface areas as determined by the CAD model of the system and the outgassing rates from literature [52][53][54][55]. The numbers shown in this table are only a rough approximation as the exact value will depend on heat treatment and pre-baking procedure. The area of TorrSeal has been approximated with a likely value as it is difficult to determine the exact surface area covered. 29

2.2 Rate of Helium permeation through the windows in the sensor head. 30

3.1 Outline of key features of the experimental system. 51

CHAPTER 1

INTRODUCTION

The force of gravity is the attraction between two massive bodies; the more massive these bodies and the closer they are the greater this attraction. Mathematically this relationship is represented by Newton's law of gravitation

$$\mathbf{F} = \frac{G m_1 m_2}{r^2} \hat{\mathbf{r}}, \tag{1.1}$$

For point masses m_1 and m_2, with a distance r between their centres of mass in the $\hat{\mathbf{r}}$ direction. Treating the two bodies as point masses allows us to assume that all the gravitational force acts through the centre of mass, this also implies that there is complete uniformity in their mass distribution. This is valid only under the condition that the diameters of the two masses (d_1 and d_2) are much smaller than the distance between them (Fig. 1.1)

$$d_1, d_2 \ll r. \tag{1.2}$$

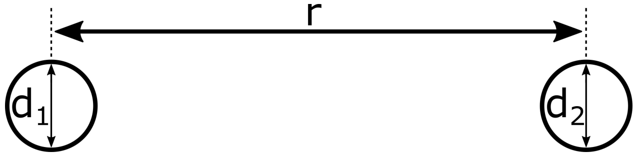

Figure 1.1: Two bodies of diameter d_1 and d_2 with uniform mass distributions, separated by a distance r.

Under this condition it would be impossible to see local variation in the mass distribution and thus they can be assumed homogeneous. That is to say if the gravitational force at m_1 were measured as it orbits m_2 at fixed distance r then by the symmetry of the system it would be impossible to determine the relative position of the two masses.

This condition breaks down when one or both of d_1 and d_2 are the same order or larger than r and there is an asymmetry in the mass distribution. If some small asymmetry were introduced to the case in Fig. 1.1, then the body could still be modelled as a homogenous mass, $m_2^{'}$, so long as condition 1.2 were maintained. The variation in the mass distribution would provide a negligibly small difference as compared to the symmetric case. When this condition no longer holds the contribution from the inhomogeneity will introduce a non-negligible offset to the centre of mass of the body.

Taking the black oval in Fig. 1.2 as a region of increased density, then the system no longer has symmetry of position and as $r^{'} \approx d_2$ then the inhomogeneity is impossible to ignore. Again consider the gravity as measured at m_1 as it orbits m_2; for fixed r there would now be a substantial change in the measured gravitational acceleration as $r^{'}$ varies. One can either use this spatial variation with a known anomaly to determine position or with a known position to determine an unknown anomaly.

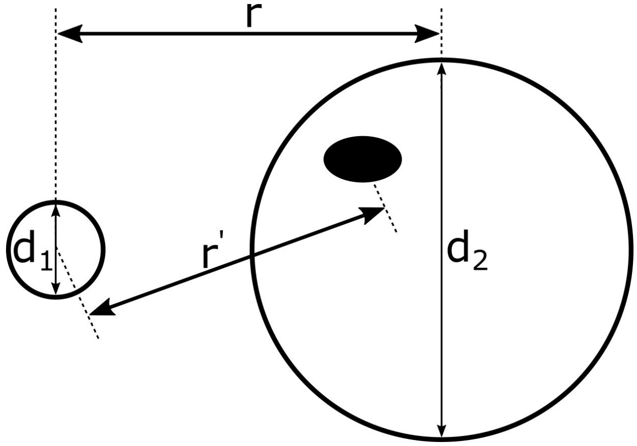

Figure 1.2: The situation in Fig. 1.1 has been adapted for the case in which one of the bodies no longer has a uniform mass distribution and $d_2 \approx r > d_1$.

The second of these is the subject of interest for this thesis. Using a sufficiently precise accellerometer, measurements of local acceleration due to gravity can be taken over the area of interest, matched with known location data the resulting inverse problem can be solved to identify the anomaly. This is a gravity survey; by characterising the local gravity profile it is possible to identify hidden features of the

Figure 1.3: Scintrex CG-5 gravity sensor.

area. Applications of this technique are as varied as the potential anomalies; from resource acquisition by identifying deposits of minerals to mapping out under-road pipes and utilities[1]. Technologies such as ground-penetrating radar are already used en masse to provide this information but suffer from severe attenuation from near surface objects. The infinite range and unblockable nature of gravity means that it does not suffer from these drawbacks and provides the potential for significantly faster surveys with a reduced noise sensitivity.

1.1 State of the Art in Commercial Gravimetry

Gravity surveys are a growing area of interest for both scientific and commercial applications leading a number of companies to try and fill the product space. Some of these use classical means of measuring gravitational acceleration and some use a quantum approach similar to that of the system presented in this thesis. In this section I will give a brief overview of the current state of gravimetry.

1.1.1 Scintrex CG-5

The most widely used commercial gravimeters for survey applications are the Scintrex CG-5 and the new more compact CG-6. They have been used in everything from small scale surveys [2] to large, country-wide measurement campaigns [3][4]. The CG-5 is a sensor that utilises a classical means of detecting relative changes in gravity over the course of a survey; it is a mass on a spring gravimeter, that utilises

Hooke's law to link spring extension to local gravity.

$$F_s = -kx \tag{1.3a}$$

$$F_g = mg \tag{1.3b}$$

$$\Rightarrow g = -\frac{kx}{m} \tag{1.3c}$$

Where F_s is the restoring force of a spring with constant k and extension x and F_g is the gravitational force acting on a mass m in local gravity g. Suspending the mass from the spring equates the two forces leading to Eq. 1.3c, if x can be measured sufficiently preicsely on a well characterised mass and spring system it is then possible determine local gravity. The CG-5 utilises a zero-length spring[5] mounted between the two plates of a capacitor such that they can supply a restoring force proportional to any changes in x. This creates a relationship between the current output from the capacitor[6][7] and the local gravitational acceleration; combining this with an inbuilt GPS receiver to provide a precise measurement of the sensor location the CG-5 providing an integrated device for surveys[8].

Like all gravimeters the CG-5 suffers from high sensitivity to environmental instabilities such as tilt, thermal and vibration. As the sensor is only sensitive to changes in spring length normal to the two plates then there is a reduction in sensitivity if the sensor is tilted as demonstrated in Fig. 1.4. An angle θ between the local gravity vector and the normal to the plates will result in the sensor measuring the projection of gravity, $g\cos\theta$, rather than true local gravity, g, since the spring will hang along the true gravity vector.

Thermal variation will induce changes to the extension of the spring as it expands or contracts. Vibrational noise however is the main problem area; vibration causes the spring to shake up and down meaning there will be a variety of noise frequencies in the final result. To suppress this, long integration times are required[9].

In addition to these environmental effects there are some systematic noise sources on the CG-5. A drift in the measured gravitational signal arises from the gradual stretching of the spring due to the con-

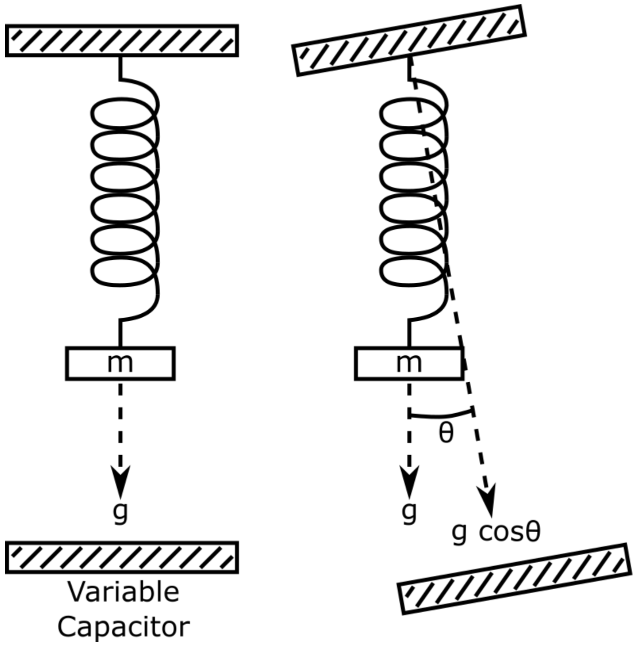

Figure 1.4: Left: A relative gravimeter in the case that the capacitor plates lie perpendicular to gravity (adapted from [2]). Right: Introducing a tilt θ between the spring and the normal to the capacitor plates measures only the signal projection of gravity $g\cos\theta$.

stant tension produced by the mass. This drift has to be constantly characterised by taking measurements at a reference base station over the course of a survey to correct for and remove it. This is also why the CG-5 is only produces a relative gravity measurement, the value from each point must be compared to a reference value or the output result will be contain a bias. Finally the mechanical components of the system are susceptible to mechanical wear, such as the aforementioned spring drift. The CG-5, although it is the workhorse of gravity surveys has limitations in measurement time and stability that can stand to be improved.

1.1.2 GWR Instruments iGrav

The iGrav operates in a way simlar to the CG-5, but replaces the mass on a spring with a superconducting spherical test mass suspended between two superconducting coils. As the gravitational field changes the sphere moves up and down. This change is fed back to the current driver and the current is adjusted to move the sphere to the centre of the trap, this difference in current can be converted into a gravitational measurement[10].

This system suffers from some issues that are fairly unique to this design. Given the whole principle of operation is based around a carefully tuned magnetic balance, a high level of magnetic isolation is required to remove environmental fields. This shielding requirement increases the size and weight of the system relative to the CG-5. Like the CG-5 it is also highly susceptible to both tilt and vibration noise meaning that in any real world environment long interrogation times and precise tilt alignment are required to remove this noise.

1.1.3 Atom Interferometry

At the time of writing this thesis two commercial quantum gravity sensors widely are currently available, μQuans[11][12] and AOSense[13] gravimeters. Both apply the same underlying technique but take slightly different approaches, though both suffer from the same limitations. They follow the same principle of operation that will be described in this thesis, outlined in more detail in Chapter 2, but with an important difference. The sensor outlined in this thesis is a gravity gradiometer that can take advantage of common mode rejection to reduce the noise floor; gravimeters do not have this capability.

With no mechanical parts involved in the operation of an atom interferometer there is no degradation or resulting drift as for classical sensors. Instead, a cloud of atoms at micro-Kelvin temperatures is used as an ensemble of test masses, manipulated with a laser and uses relative atomic state population readout as a means of measuring the gravitationally induced Doppler shift. As this is achieved by making a measurement of an object in free fall it provides a true absolute measurement of gravitational acceleration.

Quantum gravimeters solve a number of the problems faced by classical systems, however they still suffer from the two most devastating noise sources: sensor tilt and vibration. Tilt offset is problematic in the same way as for the CG-5. In a quantum gravimeter the interferometry beams define the direction in which gravity is measured, like the capacitor plates in the CG-5 (Fig.1.4), as such an angle between this and the local gravity vector will lead to a projection and signal reduction. This means that a small angular offset can very quickly lead to a large reduction in measured signal (see 1.2).

High susceptibility to vibration arises from the fact that the interferometry beams measure the phase relative to a fixed reference, i.e. the input optic for the interferometry beams or a retro-reflection mirror depending on the implementation. Vibration of this reference is an acceleration that couples to the atoms in a way which is indistinguishable from gravitational acceleration, this can be averaged out with long integration times. Compensating for these noise sources is possible but doing so usually requires very large, heavy stages that adversely affect the size, weight and power (SWAP) of the system. Removing this vibration noise requires an important design choice between two important factors for a survey system: measurement time and SWAP constraints.

Long measurement times remove one noise source for sensitive gravimeters but introduce another. Gravimeters that are capable of measuring the small changes to local gravity in a survey are also sensitve to the motion of large masses, such as the tidal effect of the ocean or the moon[6]. These are a serious noise source for gravimeters as over the course of a day they can cause a change to the measurement of gravity of ~300 μGal[1] [14]. Models have been developed to correct for this drift[15] but the precision of the measurement is then linked to the accuracy of the model.

1.2 Our Goal

The approach considered in this thesis solves the most common problems suffered by gravity sensors. Instead of taking a single spatial measurement of gravity, two are taken and a difference is determined. Mechanically coupling together two classical sensors[16] with a separation between them allows a gradiometer to be constructed that benefits from similar common mode vibration suppression as a quantum system. They still suffer from some of the same limitations as their gravimeter counterparts[17] as they are subject to drift from mechanical wear and will only produce a relative gradient measurement. Similarly the vibration suppression is limited by the mounting structure, for example any damping of the vibration along path between the sensors will lead to a reduction in this common mode suppression. A cold atom gradiometer removes these limitations, directly coupling together the two test masses with a common reference beam placing a time-stamp on the two clouds which is inherently synchronous. This common interrogation beam leads to a direct coupling of any systematic vibration directly onto both clouds.

Quantum gravity gradiometers have been used at largest scale for tests of fundamental physics[18]

---

[1] $1\,\mu\text{Gal} = 10^{-8}\,\text{m/s}^2 = 1.019\,\text{ng}$.

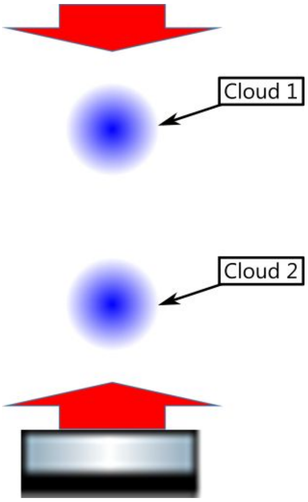

Figure 1.5: For an interferometer the interrogation light is input from one end and retro-reflected from a mirror at the base of the system. This leads to both atomic clouds being simultaneously interrogated by the same light.

and in truck mounted systems that have been able to identify buildings from their horizontal gravity profile[19][20] for almost two decades[21]. In a gradiometer both atom clouds are interrogated with the same laser which is retro-reflected to provide a counterpropagating beam (Fig. 1.5) allowing the probing of a high momentum transfer interaction[22]. As the laser light measures the Doppler shift due to gravity an acceleration of the reference mirror, such as vibration, will be indistinguishable from gravitational acceleration. For a gradiometer this acceleration noise will be imprinted equally onto both clouds allowing a subtraction to remove it.

Performing two simultaneous measurements of gravity at different heights against a common reference it is possible to measure the gradient whilst removing noise sources. Additionally tilt sensitivity is suppressed on the measurement; let us say that we want to see a gravitational signal of 1 ng, then the tilt of the sensor θ_g must then produce at most 1 ng of noise onto the measured value

$$1\,\text{ng} \geq g\left(\cos(0) - \cos\theta_g\right) \tag{1.4}$$

Where this is considering a deviation of the ideal case, i.e. $\cos(0)$. A similar statement can be made for a gradient signal of 10 E[1]

---

[1] 1 Eötvös (E) = $10^{-9}\,\text{s}^{-2} = \frac{g_2 - g_1}{b}$ for two gravity measurements g_1 and g_2 separated by a baseline b. This relation means

Figure 1.6: The subject of this thesis, the portable gravity gradiometer constructed at the University of Birmingham in its intended environment.

$$10\,\mathrm{E} \geq \delta g \left(\cos(0) - \cos\theta_{\delta g}\right) \tag{1.5}$$

Given the acceleration due to gravity $g \approx 9.81\,\mathrm{ms}^{-2}$ and the gravity gradient $\delta g \approx 3000\,\mathrm{E}$[23] then it is possible to identify the acceptable tilt noise for the two cases. Rearranging these equations for the angle sets the maximum tilt alignment variation that is acceptable to still see these signals, $\theta_g \leq 45\,\mu\mathrm{rad}$ and $\theta_{\delta g} \leq 81$ mrad; demonstrating a reduction in tilt sensitivity of almost 2000 times for the gradiometer.

The vision of this thesis was a simple one; to demonstrate that a system based on cold-atom technology was capable of being shrunk and ruggedized whilst maintaining a useful level of sensitivity, suffering minimum reduction in functionality when in challenging environments. With this goal in mind choices were made in how best to deliver a system that can work as a research prototype whilst a providing a design that, once validated, could be used as the inspiration for future iterations. A brief outline of some of the key design and technology choices made to enable this vision will be outlined below.

Traditional Magneto-Optical Trap (MOT) systems use independent six-beam delivery in order to provide three dimensional cooling to the atomic cloud[24][25]. This leads to a large optical delivery system

---

10 E is the gradient signal expected from a 1 ng difference with a 1 m baseline.

increasing the size of a sensor. With multiple distinct inputs there is the possibility for relative alignment, light intensity and polarisation vraiation which will introduce instability onto the cloud. These variations can manifest as a change to the cloud position[26], atom number and temperature that could produce a devastating noise on the gravitational measurement. To limit the impact of these effects a single beam delivery was chosen, this ensures that all noise in the light delivery becomes common whilst facilitating a reduction in sensor volume.

Single-beam delivery allowed the use of a compact design for the ultra-high vacuum (UHV) system whilst improving the stability of the atom cloud (see Fig. 2.6). Reducing the overall footprint of the science chamber led to a streamlining of the system which in turn facilitated a reduction in system weight. The reduction in weight comes from a small UHV system being inherently lighter and that is allows a simpler design for the magnetic shields, often the single largest contribution to sensor weight.

Magnetic compensation is required to ensure that there is no field during the sub-Doppler cooling process and a well defined field direction during the interferometry sequence. For a system that will be frequently moved the most effective way to do this is with magnetic shields to provide passive compensation, unfortunately introducing a substantial weight component. The use of advanced manufacturing techniques is under investigation to allow for more compact shielding in systems with more complex geometries[27][28]. As these techniques are still immature, reducing the system size is the most effective way to limit sensor weight.

Simplifying the chamber and mounting design allows the sensor heads to be made almost identical, facilitating a modular manufacturing process that will provide a substantial benefit if converted into a commercial product. Light generation is simplified by the fortunate synergy between the level structure of ^{87}Rb, with a 780 nm transition and years of development by the telecoms industry in creating quality fibre coupled 1560 nm wavelength lasers. The availability of components for these wavelengths and the ease with which the telecoms lasers can be frequency doubled allowed for the development of robust, compact and agile lasers[29][30].

As the system (pictured in Fig. 1.6) has begun to produce promising results and with the filing of a patent[31] the design has begun to be used in a number of new projects within the University of Birmingham and by industrial partners. Changing the requirements on sensitivity and SWAP it is possible to explore different regions of the parameter space with modified designs to find new solutions to common problems. One such system has placed a focus on improving the sensitivity whilst reducing SWAP constraints to find the operating limits of the design. Another is exploring the other end of the spectrum placing much more stringent constraint on SWAP, to the point that it can be mounted onto a drone, whilst maintaining the ability to perform atomic interferometry with reduced sensitivity. In addition to these local projects new systems are in development in partnership with Teledyne e2V to harness the work done in the context of a research project and create an industrial prototype to feed into an eventual commercial product.

1.3 Sensitivity Requirement

To make a portable sensor, limitations must be placed on the SWAP requirements. As many of the choices made in reducing sensor SWAP impact its potential sensitivity and stability, most laboratory systems are built into a laboratory and stay there for decades. However a system that has to be moved and used in the field has very different priorities; it must be as compact as possible whilst maintaining functionality. In order to arrive at the system shown in Fig. 1.6, these limitations had to be set such that the SWAP is minimised but sensitivity was usefully high.

To set this sensitivity a test case (Fig. 1.7) was considered where two gravimeters at points S_1 and S_2 are aligned directly above a point like anomaly of mass M at distances z_1 and z_2 respectively from the sensors. The gradiometer baseline is defined as the separation between the two sensors, $b = z_2 - z_1$. The gravitational acceleration at the two points is given by

$$g_1 = \frac{GM}{z_1^2} \tag{1.6a}$$

$$g_2 = \frac{GM}{(z_1 + b)^2}. \tag{1.6b}$$

Setting $b \approx z_1$, as is the definition for a near surface anomaly, the difference between the two signals in Eq. 1.6 is given by

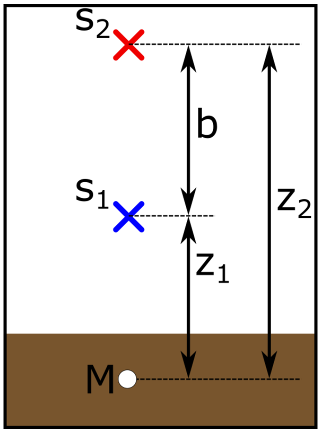

Figure 1.7: The toy model for two gravimeters at points S_1 and S_2 above a buried point mass.

$$g_1 - g_2 \approx \frac{3}{4}\frac{GM}{z_1^2} = \frac{3}{4}g_1. \tag{1.7}$$

For a given sensor if the anomaly of interest produces a signal smaller than this it will not be visible within its resolution.

Now consider from the perspective of the atoms; the phase induced on an atom over the course of the drop is

$$\phi = k_{\text{eff}} g_1 T^2. \tag{1.8}$$

Where k_{eff} is the effective wavevector of the two photon interaction (producing a momentum transfer of $\hbar k_{eff}$) and T is the length of time between the interferometry pulses, giving the “length" of the interferometer arms (see 2.3.3). For N atoms in the cloud there is a single shot sensitivity gain of $1/\sqrt{N}$. This leads to a single shot resolution of [32]

$$\frac{\Delta g}{g_1} = \frac{\Delta\phi}{\phi} = \frac{1}{\sqrt{N} k_{\text{eff}} g_1 T^2}. \tag{1.9}$$

Hence the difference given by Eq. 1.7 must be larger than or equal to the sensitivity given by Eq. 1.9. Assuming that S_1 and S_2 have the same sensitivity this tells us

$$\frac{1}{\sqrt{N}k_{\text{eff}}T^2} \leq \frac{3}{4}g_1. \tag{1.10}$$

There are three clear ways to improve sensitivity and meet this condition: increase atom number, momentum transfer or the drop time. In a two-photon scheme the momentum transfer is limited by the transition which can be probed and as such by the choice of atom. To take advantage of the 1560 nm lasers used in telecoms this limits the choice to the D_2 line of ^{87}Rb, setting the value of $k_{\text{eff}} \approx 2\times 2\pi/780\,\text{nm}^{-1}$. In many photon interferometry schemes it is possible to increase k_{eff} but this introduces additional complexity so only a two-photon scheme is used in this thesis. Increasing the atom number and drop time of the interferometer can be achieved by changing the sensor geometry but from Eq. 1.9 it is obviously favourable to increase the T-time due to the T^2 scaling. The atom number achievable in a system depends on a number of specific factors like beam size, laser power and vapour pressure. Common atom numbers for MOTs using similar technology are on the order of 3×10^8 [26][33][32], with roughly a 1% transfer efficiency from the first interferometry pulse (see 2.3.2) giving a value of $N = 3\times10^6$. Using these values in combination with Eq. 1.6 and Eq. 1.10 it is possible to determine the T-time required to detect a given signal. In this case the signal is taken to be a 220 kg mass 3 m below the bottom sensor, the signal expected from a void in cement (e.g. a sewage pipe) of 120 cm diameter and 2 m long, requiring a T-time of

$$T \geq \sqrt{\frac{4}{3}\frac{z_1^2}{\sqrt{N}k_{\text{eff}}GM}} = 172\,\text{ms}. \tag{1.11}$$

We have set the minimum requirement of the total drop time; the T-time needs to be at least 172 ms ($2T$ =344 ms). Using ballistics this means a total drop of 58 cm, however given that gravity is a conservative field it is possible to gain the same T-time in a quarter of the distance by launching the atoms before letting them again fall under gravity, producing an atomic fountain[34][35]. With this trick the same $2T$ time of 344 ms would only need a launch distance of 14.5 cm, which means that an absolute minimum height for a gradiometer (with a 1 m baseline) is ~1.15 m rather than the ~1.60 m without the launch. This calculation was used as a guideline for the drop distance required in the system; more limitations were imposed when aspects such as manufacturing capabilities and additional SWAP constraints were considered.

1.4 Thesis Breakdown and a Note on the Plotting Convention

Development of the gradiometer design concept will be outlined in Chapter 2, this will include preliminary characterisation measurements on prototype systems that validated each step of the process. The final design of the completed sensor package will be covered in Chapter 3 with a brief overview of the compact laser and control systems of the control package provided in Chapter 4.

System optimisation, characterisation and procedure will be outlined in Chapter 5 and results from all field trials with the completed sensor, along with an explanation of their analysis, will be detailed in Chapter 6.

Note on the plotting convention Much of the characterisation will need to be considered concurrently for both chambers and large datasets, meaning it is useful to be able to plot them together in the same plotting region. For some of the datasets considered it is also not practical to differentiate them with different line or marker styles. Therefore the most effective way to present this data is to use different colours for the different chambers; throughout this thesis all red data sets will represent the Top sensor and any blue data sets will represent the Bottom sensor (as shown in Fig 1.7).

1.5 Statement of Contributions

This work has been undertaken within the scope of a larger project at the University of Birmingham and has been a collaboration between the author and a larger team. Some of the work presented in this thesis is the work of others given to provide a full understanding of the experiment, K. Bongs and M. Holynski have provided general input and advice throughout. Substantial direct contributions have come from: A. Kaushik who built and updated the laser systems, A. Niggebaum who made the computer control system and intensity stabilisation circuits and A. Stabrawa who designed the bias coils (3.2.1) and optical delivery system (3.3). Additional contributions have been supplied by: A. Rodgers who developed the ellipse analysis (6.2), B. Stray who helped in the assembly of the flight case and vacuum systems, D.

Gilbert who helped in the development of the intensity stabilisation and J. Vovrosh who helped in day-to-day system operation.

This contributed work is mostly confined to Chapter 4. General optimisation and day-to-day operation of the system has been a team effort, however all results presented in this thesis are the work of the author.

CHAPTER 2

BUILDING A PORTABLE QUANTUM DEVICE

In order to validate and characterise the potential of a prism MOT gradiometer the design went through two initial prototyping stages before implementing into the final design. The initial stage tested the capabilities of the prism MOT[36] as a compact cold atom source, to ensure that the atom number used in Eq. 1.11 is achievable. The second stage involved placing this prism MOT into a gravimeter with a more developed vacuum system to test that this atom source is sufficiently cold to maintain good contrast and the system pressure was sufficiently low as to make the impact of background collisions negligible (see 2.3.3). The latter sections of this chapter will present the development of the gradiometer system. This chapter started with some underlying theory required to understand the results that will be presented.

2.1 Magneto-Optical Trapping

The MOT is the workhorse of quantum sensors acting as the starting point for everything from rotation sensors[37][38] and inertial navigation systems[39][40] to high precision clocks[41] and atomic gravimeters[35] for more than 30 years[42]. Forming a MOT is a matter of implementing an optical molasses scheme on a magnetically sensitive atomic transition in the presence of a position-dependent magnetic quadrupole field.

To understand this let us consider the case of an atom with two energy levels $|g, l = 0\rangle$ and $|e, l = 1\rangle$ separated by an energy $\Delta E = h\nu$ with a transition linewidth of Γ. Figure 2.1 presents the case of placing this atom in a quadrupole magnetic field with the condition that $B(z = 0) = 0$. Introducing a pair of counterpropagating laser beams (L_1 and L_2) at frequency ν_{laser} makes it possible to drive the $|g, l = 0\rangle \rightarrow$

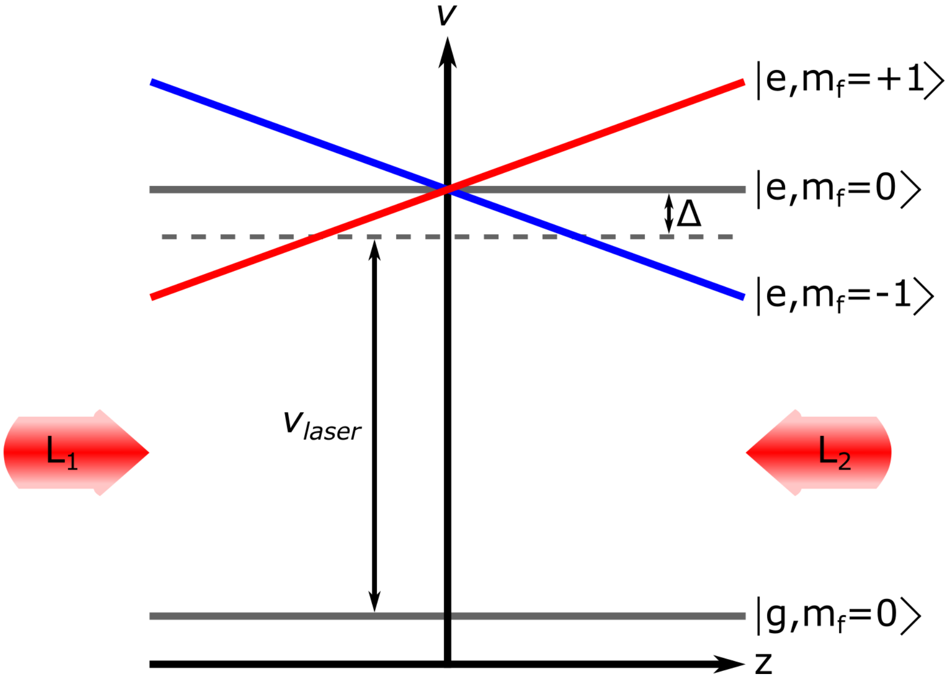

Figure 2.1: Simple energy level diagram for a two level MOT system. In an active MOT system the lasers L_1 and L_2 would have σ_+ and σ_- polarisations respectively. The frequency of the laser is ν_{laser} and Δ is the detuning from the atomic resonance.

$|e, l = 1\rangle$ transition. Consider the case that L_1 and L_2 contain only light polarised to drive π-transitions, in this case conservation of angular momentum means only the $|g, m_f = 0\rangle \rightarrow |e, m_f = 0\rangle$ transition can be driven. For an atom moving towards L_2 with a velocity $\mathbf{v}$ the force experienced will be the difference between the scattering forces from the two light sources,

$$F_1^{\text{scatt}} = \frac{h\nu_{\text{laser}}}{c} R^{\text{scatt}} = \frac{h}{\lambda_{\text{laser}}} \frac{\Gamma}{2} \frac{I_1 / I^{\text{sat}}}{1 + I_1 / I^{\text{sat}} + 4\,(\Delta - \mathbf{k}_1.\mathbf{v})^2 / \Gamma^2} \tag{2.1a}$$

$$F_2^{\text{scatt}} = \frac{h\nu_{\text{laser}}}{c} R^{\text{scatt}} = \frac{h}{\lambda_{\text{laser}}} \frac{\Gamma}{2} \frac{I_2 / I^{\text{sat}}}{1 + I_2 / I^{\text{sat}} + 4\,(\Delta - \mathbf{k}_2.\mathbf{v})^2 / \Gamma^2} \tag{2.1b}$$

$$F^{molasses} = F_2^{\text{scatt}} - F_1^{\text{scatt}}. \tag{2.1c}$$

Here the light intensities are given by I_1 and I_2, the transition saturation intensity is I^{sat} and the wavevectors of the two lasers are $\mathbf{k}_1$ and $\mathbf{k}_2$. Assuming that $I_1 = I_2 = I$ it is clear that for an atom moving towards L_2

that $F_2^{\mathrm{scatt}} > F_1^{\mathrm{scatt}}$. This will produce a force in the direction of propagation of L_2, oppposing the motion of the atoms, that scales with velocity. This is the basis of an optical molasses and acts to slow the atoms in a given direction, cooling them via radiation pressure[43].

Considering now the case where L_2 contains light that can drive a σ_- transition only and L_1 can only drive σ_+ transitions, therefore only the $|g, m_f = 0\rangle \rightarrow |e, m_f = \pm 1\rangle$ transitions can be driven depending on the light source. In this modified case the optical molasses process described by Eq. 2.1c will still occur but there is an additional frequency shift arising from the Zeeman shift on the upper energy levels. This position dependent force will also have a bias towards one of the two lasers. In the region $B > 0$ the scattering rate for the $|g, m_f = 0\rangle \rightarrow |e, m_f = -1\rangle$ will be greater than that for $|g, m_f = 0\rangle \rightarrow |e, m_f = +1\rangle$ as the transition is closer to resonance, with the inverse being true for $B < 0$. The frequency shift due to the Zeeman effect (ν_{ZE}) is

$$\nu_{\mathrm{ZE}} = \frac{g_F \mu_B m_f B(z)}{h} \tag{2.2}$$

$$\Rightarrow \nu_{\mathrm{ZE}}(z) = \frac{g_F \mu_B m_f}{h} \frac{dB}{dz} z = m_f \zeta(z) \tag{2.3}$$

$$\zeta(z) = \frac{g_F \mu_B}{h} \frac{dB}{dz} z \tag{2.4}$$

Where g_F is the Landé g-factor for the fine structure, μ_B is the Bohr magneton and the condition for $B(z = 0)$ has been imposed. This then modifies Eq. 2.1 with a position dependent component

$$F_1^{\mathrm{MOT}} = \frac{h\nu_{\mathrm{laser}}}{c} \frac{\Gamma}{2} \frac{I/I^{\mathrm{sat}}}{1 + I/I^{\mathrm{sat}} + 4(\Delta - \mathbf{k}_1.\mathbf{v} - \zeta(z))^2/\Gamma^2} \tag{2.5a}$$

$$F_2^{\mathrm{MOT}} = \frac{h\nu_{\mathrm{laser}}}{c} \frac{\Gamma}{2} \frac{I/I^{\mathrm{sat}}}{1 + I/I^{\mathrm{sat}} + 4(\Delta - \mathbf{k}_2.\mathbf{v} + \zeta(z))^2/\Gamma^2}. \tag{2.5b}$$

Where the fact that L_2 can only address the $m_f = -1$ state and L_1 the $m_f = +1$ state has been used to set the sign on the $\zeta(z)$ term. It can be seen that for $z > 0$, $F_2^{\mathrm{MOT}} > F_1^{\mathrm{MOT}}$ and for $z < 0$ the converse is true and for $z = 0$ then $F_2^{\mathrm{MOT}} = F_1^{\mathrm{MOT}}$. A position dependent force now exists on the atom which vanishes at the centre, producing the trapping component of the MOT.

From this it can be seen how a one-dimensional MOT could be implemented with the appropriate choice of magnetic field profile, which can easily be extended to a 3D case with a quadrupole field and a six-beam delivery[44] (Fig. 2.2a). Equations 2.5 can also be used to easily show why a MOT with a single beam delivery, like that in Fig. 2.2b, can be used to improve cloud stability. Taking now that $I_1 \neq I_2$ then the forces in Eq. 5.9 will have an intensity dependence, it is easy to see that if $I_1 > I_2$ the condition $F_1^{\mathrm{MOT}}(z=0) = F_2^{\mathrm{MOT}}(z=0)$ will no longer be satisfied. This will cause the position of the MOT to change depending the relative intensities of the two beams.

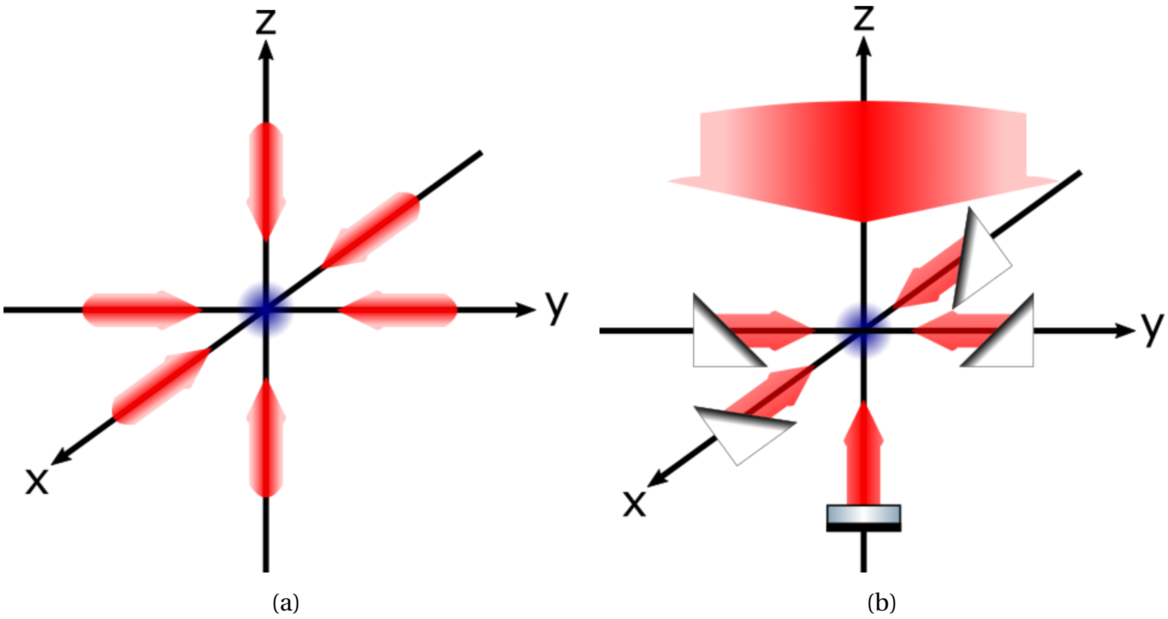

(a) (b)

Figure 2.2: (a) A traditional six-beam MOT with are individual delivery beams for all six co-ordinate axes from six independent delivery telescopes. (b) In a prism MOT a single delivery beam provides the light for all six capture directions.

For a system with a common input beam the condition that $I_1 = I_2$ is inherently true. As these transitions require circularly polarised light it is reasonable to treat any polarisation variation of the light as a change in the useful intensity in the beam. This means that whereas the traditional six-beam MOT is sensitive to both relative polarisation and intensity drift on the input light, in a single beam MOT these are common mode on all six cooling beams providing a more stable system to work from (see Fig. 2.6).

2.1.1 Capture Velocity and Volume

In a MOT system if an atom is moving sufficiently fast it will be able to completely cross the extent of the MOT beam before it is cooled to the the point at which it can be trapped. The maximum velocity an atom can have when entering the cooling beam that will result in capture is known as the capture velocity, v_c. Practically, the capture velocity of a given system is determined by a combination of the intensity of the laser light, its frequency and the size of the beams. To determine the maximum capture velocity of the system the easiest way is to consider Eq. 2.5a for $I \to \infty$; this assumption will be used in 2.4 to compare beam sizes at finite intensities. For $I \to \infty$ the scattering force from any given beam is

$$F^{\infty} \approx -\frac{h\Gamma}{2\lambda_{\text{laser}}}. \tag{2.6}$$

The minus has been introduced as this is a retarding force and the ∞ label is used to signify variables where infinite light intensity has been assumed. For an atom of mass m this results in an acceleration that can be considered velocity independent

$$a^{\infty} = -\frac{h\Gamma}{2m\lambda_{\text{laser}}}. \tag{2.7}$$

With this constant acceleration it is possible to consider the atoms as classical ballistic objects. If the final velocity of the atoms is assumed to be 0 m/s relative to the background atoms then for a beam of width L the capture velocity is

$$v_c^{\infty} = \sqrt{\frac{h\Gamma L}{m\lambda_{\text{laser}}}}. \tag{2.8}$$

In the case of square input beams with side length L a good approximation of the capture volume (V_c) will be L^3 as the overlap region for the 3 beam directions would be a cube with side L. In a six-beam system with high optical power (i.e. $I \to \infty$) we have a relation for v_c. This can be linked to the atom number by considering Equation 3.10 in [45] where it is shown that the saturated atom number (N_0) of a spherical MOT obeys the relation

$$N_0 \propto A v_c^4. \tag{2.9}$$

Where A is the area of a circular input beam. In the case of the prism MOT $A = L^2$ and Eq. 2.7 shows that $v_c \propto \sqrt{L}$. Therefore it is clear from Eq. 2.9 that the atom number scales as $N_0 \propto L^4 = V_c^{4/3}$. This means that in a system of finite intensity the exact prism geometry will have a severe impact on the achievable atom number. The importance of this will be explored in 2.4 and used to validate the specific prism MOT implementation used in our gradiometer.

2.2 Prism MOT Prototype

In order to test design choices for streamlining a single beam delivery system and allow for the characterisation of the MOT, a simple prototype vacuum chamber was produced. This also allowed for evaluation of the in-house manufacturing capabilities and robustness of the design with results from these tests used to inform the design of the final system. During the design of the prism MOT, thought had to be given to how it may be developed long-term into a complete sensor, in particular how the size and weight of the system could be minimised. As discussed in 1.2 the magnetic shield is the leading source of weight in a portable system and by using single-beam delivery the shield can be shrunk considerably. With the removal of the additional telescopes that drive up the volume of six-beam systems the size will instead be dominated by the prisms and the peripheral vacuum components. As prism size will have a direct impact on the capture volume, and thus atom number of the MOT, this size will be set by the sensitivity requirements of the system.

In 1.3 an example case was considered using realistic values for the free parameters in Eq. 1.9 to determine the maximum possible sensitivity of the sensor. For this example a value of the order of 10^8 was used for the MOT atom number, where it was assumed that only 1% of these atoms would be interrogated during the interferometry sequence due to the velocity selective nature of Raman beams (see Fig. 2.18). This atom number was representative of MOTs with a ~25.4 mm beam diameter (capture volume of ~8600 mm^3). As 1.3 showed this would provide appropriate sensitivity for anomalies of interest it was decided that the prism MOT should have a similar capture volume, which would require a

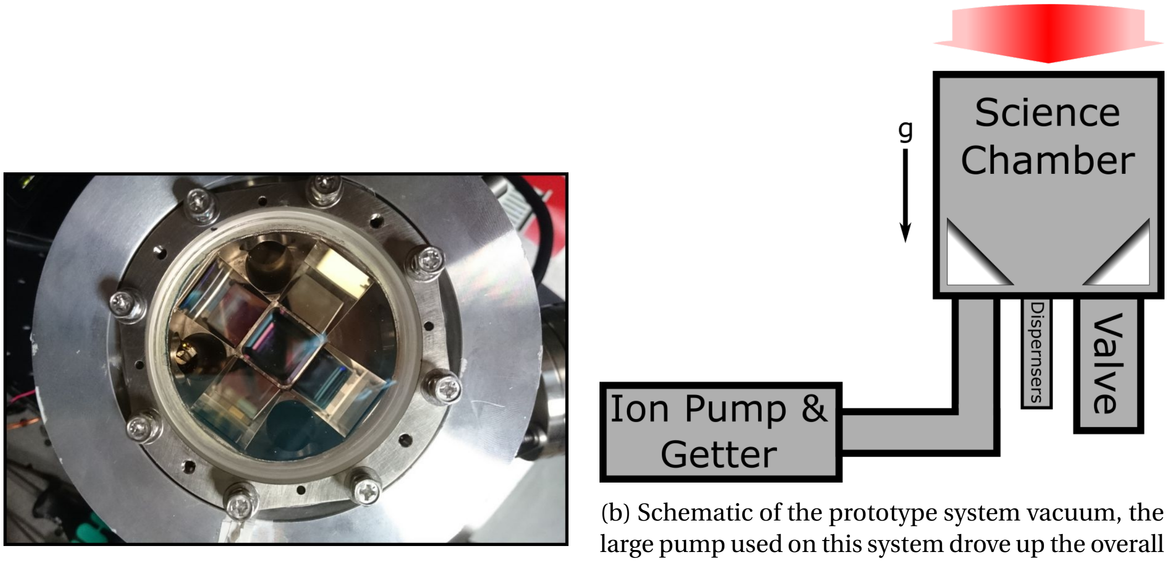

(a) Prism MOT prototype with mounted coils.

(b) Schematic of the prototype system vacuum, the large pump used on this system drove up the overall footprint but would be replaced in later designs.

Figure 2.3: Diagram of the prism MOT prototype.

prism with side length ~20.5 mm. As this was not available 20 mm prisms were used, giving a capture volume of 8000 mm^3. This was to be the limiting factor in the size of the sensor; the system needed to be large enough to hold 20 mm prisms whilst also having the facility to attach vacuum components such as pumps and atom sources.

This was achieved by mounting all vacuum components on a single back plate, opposite but in line with the telescope delivery, to drive down the overall system footprint. Additionally, indium sealed windows provided good optical access whilst not requiring bulky mounting flanges. A dispenser[1] was used as the rubidium atom source, this is a small metal capsule that acts like a filament; when current is passed through it the metal heats up causing rubidium to be released from the surface. This option provides two benefits over others atom sources such as ampoules; by mounting directly onto a CF16 power feedthrough it provides no impact on the footprint and maintains the cylindrical form factor of the vacuum system. Additionally it provides a degree of control of the rubidium pressure allowing the optimisation of loading times to give a desired repetition rate of the sensor.

The vacuum is maintained with a combination passive getter and active ion pump[2]; this supplies a 100 l/s pumping speed split 95 : 5 between the passive and active pumps respectively. More compact

---

[1] *Saes 5G0125*

[2] *Saes NEXTorr 100-5D*

pumping solutions are used in later designs once calculations had been performed for the more complex gravimeter vacuum system proving that more compact components with a lower pumping speed were sufficient (see 2.3.1). The result of these design choices can be seen in Fig. 2.3a, where with the coils used for generating the MOT field mounted there is no visible protrusion, aside from the large attachment needed for the pump.

This prototype system is currently on loan to Gooch & Housego and they have kindly supplied a loading curve produced with their home built laser (Fig. 2.4). This laser currently does not feature a frequency lock and therefore is left free running, leading to the noise on the curve. This measurement was performed with a detuning of ~13 MHz and a total optical power into the telescope of ~90 mW, split between the main cooling light and the repump sideband at a roughly 1 : 10 ratio.

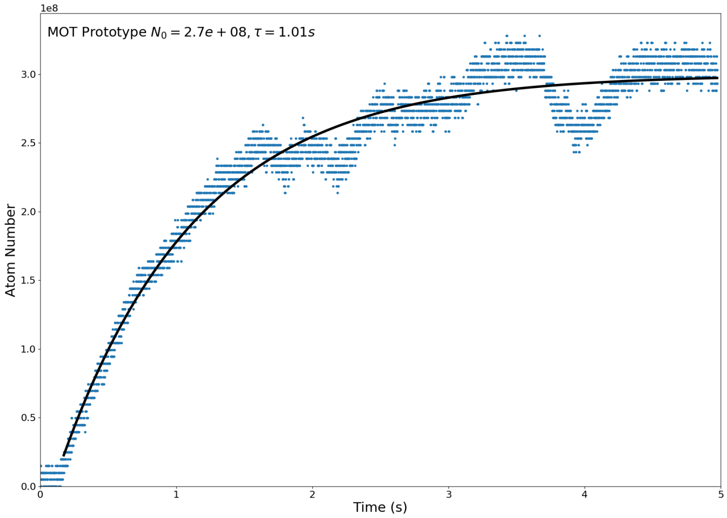

Figure 2.4: Loading curve for the prism MOT prototype demonstrating 2.7×10^8 atoms required for sufficient sensitivity (Eq. 1.11).

From this graph two conclusions can be drawn. The first is that the 20 mm prisms produce a satisfactory atom number at saturation of 2.7×10^8, 10% lower than that assumed in the sensitivity calculation (1.3). Using a higher optical power can potentially increase this value as it would lead to an increased effective capture volume (see the discussion in 2.4). The loading time is below the optimal of 700 ms, which would allow for a 1 Hz measurement rate, but can be optimised by increasing the rubidium pressure in

Figure 2.5: Left: The prototype system housed in its transport cage, below the vacuum chamber a rack mounted fibre based laser (similar to that in 4.1) is mounted. Right: The first picture of the MOT produced in the prism system with a very satisfying red glow.

the chamber (achieved by increasing the dispenser current).

A useful feature of this design is that small alignment posts were machined into the body of the chamber and by gluing the prisms against these the high machining precision of the CNC machines is imposed onto the mechanical alignment of the optics. As a result, if the system is robust at the point of assembly then there should be no alignment creep over the lifetime of the sensor. If sufficiently stable then alignment should be maintained even if the sensor is in transit, requiring minimal adjustment of the input telescope upon arrival to regain the MOT. This rigid, one time, alignment of the reflection optics however removes any potential for adjustment in optimising the sub-Doppler cooling process therefore affecting the final temperature of the cloud. As will be seen in 2.3.2 this was not a problem as temperatures below 3 μK were achieved in the second stage test.

In order to create a portable module the cooling laser system was mounted into a 19 " rack box and integrated with the vacuum system into a single $55\times55\times85$ cm cage structure for ease of transport (see Fig 2.5). This system has travelled a combined distance of over 2000 miles, in planes, trains and automobiles, for presentations at conferences and public engagement events. It was eventually retired in favour a more compact second generation vacuum chamber[46], the design of which was one of many inspired by this system and constructed as part of an undergraduate Masters project.

2.2.1 Prism MOT Stability

In 2.1 the assertion was made that due to the intensity sensitivity of the trapping force in Eq. 2.5 the prism MOT was inherently more stable. Previous work at the University of Birmingham[26] measured the “walk" of the MOT centre for a six-beam MOT over the course of ~45 minutes. Here these results have been reproduced alongside the same measurements performed on prism MOT prototype system.

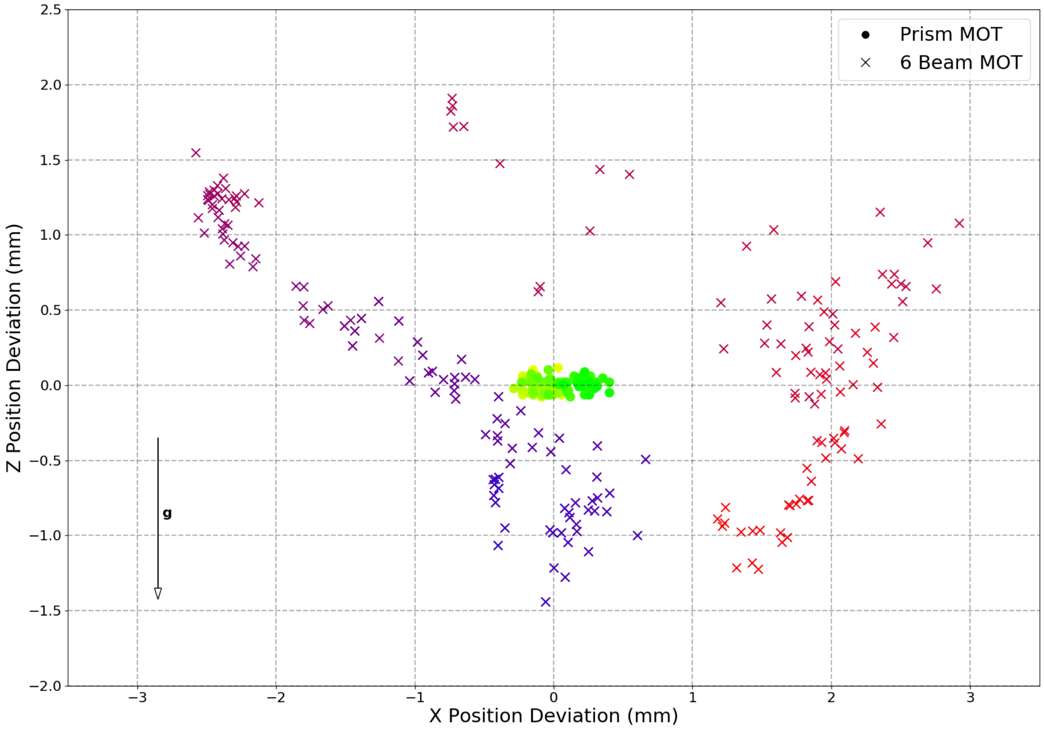

Figure 2.6: Position drift comparison between prism MOT and six-beam MOT. The data sets are both over approximately one hour going from blue to red in the six beam case and yellow to green for the prism MOT. The data for the six-beam MOT is reproduced from [26].

Figure 2.6 shows these two datasets overlaid with one another; the data for the prisms spanned a one hour period with a picture taken every 3 seconds taking the MOT centre as the point of highest intensity. The MOT centre for the six-beam system varies by ~6 mm in the *x* direction and ~3.5 mm in the *z*-direction over the course of ~45 minutes, whereas the prism MOT has a range of motion less than 1 mm in the x and 0.2 mm in the z over a longer measurement period. There is probably less variation in the z direction for the single beam system as this direction is cooled by the centre of the Gaussian intensity profile and as such has substantially higher intensity than the side beams, making it less sensitive to drifts over time. The residual motion present here is mostly due to alignment drifts in the delivery telescope and retro-reflection mirror as this system had no custom designed mounting, an issue which was fixed for the more subsequent designs by securing the telescopes directly to the chamber.

2.3 Prototype Gravimeter

Having validated the design of the prism MOT as a source with sufficient atom number the vacuum system required modification to make it a viable test bed as an atom interferometer. Much of the design is identical to that used for the gradiometer, which will be covered in depth in Chapter 3, here there will be a brief overview of the changes made. Figure 2.7 shows a schematic of the vacuum system which differs from that presented in Fig. 2.3b: the chamber has been extended to allow for a 15.5 cm drop (equivalent of $2T$ =178 ms) and pumping is now performed by a separate ion pump and getter attached directly to the chamber.

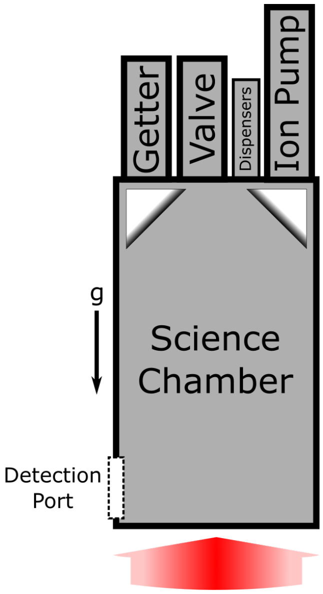

Figure 2.7: Schematic of the gravimeter vacuum system. The ion pump and getter are now two separate components mounted directly onto the chamber which has been lengthened to allow for the drop.

This introduced a potential problem, with a significant increase in the vacuum chamber volume and a reduction in pumping speed, with the new components providing 20:2 l/s passive:active pumping speeds, it may no longer be possible to achieve the same loading times as shown in Fig. 2.4 whilst maintaining a sufficiently low pressure in the drop region. The first step was to perform a MOT loading measurement using a system similar to that used for Fig. 2.4 with a laser similar to that described in 4.1.

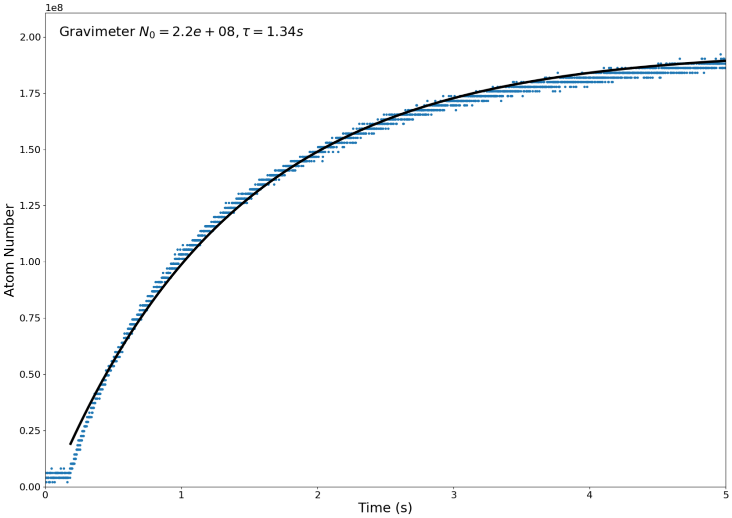

Figure 2.8: Loading curve data (blue) for the prototype gravimeter. By fitting to this data (black) a longer loading time of 1.34 s and lower saturated atom number of 2.2×10^8 is measured when compared to the MOT prototype.

This measurement used the same setup as that used in Fig. 2.4, a detuning of 13 MHz and a total optical power of ~90 mW. From this it is clear that the prism MOT integrated well into the gravimeter design, giving a similar atom number and loading time. The decrease of roughly 20% in atom number and increase of roughly 30% in loading time suggests that there is a difference in the rubidium pressure between the two systems. This can be optimised with a higher dispenser current but will result in a reduced dispenser lifetime and higher overall system pressure. With multiple dispensers it should take a number of years before it becomes an issue. The concern of the pressure during the drop, requires a more involved consideration of the potential gas load on the system given the new pumping speeds.

2.3.1 Pressure Calculation

To reduce the overall system footprint the bulky combination active/passive pump was replaced with two much smaller, albeit slower, pumps. The 95 l/s of passive pumping was replaced by a *Saes CF16 CapaciTorr* that should provide 20 l/s of pumping in the case of perfect conductance. It wasn't possible to mount directly into the chamber due to space concerns instead a custom nipple was used which quotes to give a pumping speed of ~15 l/s at its entrance. The 5 l/s active pumping was replaced by a 2 l/s active ion pump[1].

In an interferometry system the goal is to ensure that the pressure during the drop is as low as possible as background collisions cause a reduction in detected signal. Whereas in systems that load from background vapour it is desired to have a high rubidium pressure in the MOT region to reduce loading times. These provide competing requirements on the system, with higher rubidium pressure reducing the MOT loading time but resulting in a decreased signal. Let us require that over the course of the drop time $2T = 178$ ms such that only a single collision per atom can occur with background vapour, requiring the mean collision time $\tau \leq 2T$. From kinetic theory it is possible to show that a gas at temperature T_K has a mean collision rate of[33][47]

$$\tau^{-1} = n\sigma_c\sqrt{2}\langle v\rangle = \sigma_c P\sqrt{\frac{16}{\pi m k_B T_K}}. \tag{2.10}$$

Collision cross section $\sigma_c = 3\times10^{-17}\,\mathrm{m}^2$[48] assumes that only rubidium-rubidium collisions need be considered. The second equality has been made using the number density $n = \dfrac{P}{k_B T_K}$ from the ideal gas law and average velocity of a background gas $\langle v\rangle = \sqrt{\dfrac{8k_B T_K}{\pi m}}$ (for a gas particle of mass m) from kinetic theory. Requiring that $\tau \leq 2T = 178$ ms and assuming that the background gas is in thermal equilibrium with the chamber ($T_K = 293\,\mathrm{K}$) results in the pressure requirement

$$P \leq 2\times10^{-8}\,\mathrm{mbar}. \tag{2.11}$$

Where $m = 85.47$ au[49] is the mass of rubidium arising from the relative abundance of its isotopes. For the pressure of the system to be limited by the dispensers and not the background pressure, the background level must meet the condition set out in Eq. 2.11, but would preferably be orders of magnitude

---

[1] *Gamma Vacuum: 3S-DI-1H-5K-N-N*

lower. At times long after initial pump down this pressure should be dominated by the inherent gas load of the system, given by window permeation and material outgassing. This means the ultimate pressure can be determined by the ratio of this gas load, Q, and the pumping speed of the system ,S,[50][51]

$$P_{\text{Ultimate}} = \frac{Q}{S}. \tag{2.12}$$

.

In a system with negligible leaks and sufficiently long pumping time the impact of the initial water in the system is negligible and the gas load will arise from both the outgassing of vapour trapped in the materials used to make the chamber and the permeation of helium through the vacuum windows. Table 2.3.1 features the breakdown of the leading factors contributing to the gas load on the system. These contributions arise from the main body of the vacuum chamber and the prism mount which have been machined from solid titanium, with the commercial components made from low magnetic (316LN) steel and all windows made from NBK7 glass. The feedthroughs feature copper pins with ceramic beads for mounting of the dispensers and the optics have been secured in place with TorrSeal.

Material	**Area** (cm^2)	**Outgassing Flux** ($mbar \cdot l/s \cdot cm^{-2}$)	**Outgassing Rate** ($mbar \cdot l/s$)
Titanium	595	7×10^{-12}	4.2×10^{-9}
Steel (316LN)	165	3×10^{-13}	5×10^{-11}
Ceramic	15	1×10^{-14}	1×10^{-13}
Copper	51	2.9×10^{-14}	1×10^{-12}
NBK7	131	5×10^{-12}	7×10^{-10}
TorrSeal	~ 0.4	2×10^{-8}	8×10^{-9}

Table 2.1: List of materials used in the vacuum system along with the surface areas as determined by the CAD model of the system and the outgassing rates from literature [52][53][54][55]. The numbers shown in this table are only a rough approximation as the exact value will depend on heat treatment and pre-baking procedure. The area of TorrSeal has been approximated with a likely value as it is difficult to determine the exact surface area covered.

The total outgassing rate is the sum of contributions from all individual elements of the system to give a net rate of 1.3×10^{-8} mbar·l/s. Another gas load on the system is the helium permeation through the glass windows[51]. The mechanism behind this is that no solid is completely impermeable to all substances so after a sufficiently long time the pressure will rise due to this passive leakage effect [56]. This rate of pressure change depends on the window area (A), thickness (d), the difference between internal and external pressure (ΔP) and the glass permeation rate (K)

$$Q_{\text{He}} = \frac{KA\Delta P}{d}. \tag{2.13}$$

The value of $K_{\text{NBK7}} = 3.5 \times 10^{-11}\ \text{cm}^2/\text{s}$ is the helium permeation rate for $NBK7$ borosilicate[57] glass at 20 °C, approximately room temperature. Placing this value into Eq. 2.13 with the window parameters a permeation rate can be determined. Contributing to this is one 3 " top window used to deliver the light into the chamber, one 1 " bottom window used for the retro-reflected light and eight 1 " side windows used for MOT monitoring and detection. Using the helium partial pressure in air for with a relative abundance of $6 \times 10^{-}6\%$ this gives the value $\Delta P \approx 5 \times 10^{-5}$ atmosphere $= 5 \times 10^{-2}$ mbar taking the internal pressure as 1×10^{-10} mbar.

Window	**Area** (mm^2)	**Thickness** (mm)	**Q$_{\text{He}}$** ($mbar \cdot l/s$)
Top	3320	10	5×10^{-21}
Side	300	5	1×10^{-20}
Bottom	320	10	6×10^{-20}

Table 2.2: Rate of Helium permeation through the windows in the sensor head.

Giving a total permeation rate of $1.5 \times 10^{-19}\ \text{mbar}{\cdot}\text{l/s}$, which is orders of magnitude below the outgassing contributions given in Table 2.3.1. Getter pumps do not effectively pump noble gasses so when considering the ultimate pressure this can only be removed by the 2 l/s active pump. The ultimate pressure of the system can now be found by combining these two contributions with Eq. 2.12

$$P_{\text{Ultimate}} = \frac{Q_{\text{Outgassing}}}{S_{Total}} + \frac{Q_{\text{He}}}{S_{Ion}} = \frac{1.3 \times 10^{-8}}{15+2} + \frac{1.5 \times 10^{-20}}{2} \approx 8 \times 10^{-10} mbar. \tag{2.14}$$

Equation 2.14 only considers the passive background gas load on the system and does not account for the impact of the dispensers. The result from Eq. 2.14 is two orders of magnitude lower than the condition placed on the system by Eq. 2.11. A gas load from the dispensers far in excess of the background contributions will be required to reach a collision rate during the drop sufficient to disturb the interferometer.

2.3.2 Cloud Temperature

Finite-length counterpropagating Raman pulses have a velocity selective effect on atomic clouds[58][59], arising from the fact that a finite pulse in the time domain produces a $sinc$ function in the Fourier domain. Counterpropagating Raman beams are used as this causes the absorption and emission of a photon in opposing directions meaning that the momentum transferred at each of the pulses is given by the sum of the two photon momenta[22], whereas copropagating beams induce a net momentum transfer given by the difference between the two

$$\hbar k_{eff}^{\text{counter}} = \hbar\left(k_1 + k_2\right) \tag{2.15}$$

$$\hbar k_{eff}^{\text{co}} = \hbar\left(k_1 - k_2\right). \tag{2.16}$$

A thermal cloud has a distribution of velocities, with a range of Doppler shifts across the cloud, meaning the atoms can can interact with the range of frequencies present in the finite square pulse. However only a subset will be within the resonance condition for both beams. A shorter pulse will have more frequency components in the Fourier domain and as such will be resonant with a broader velocity distribution of the atoms; longer pulses will have a more monochromatic representation in the Fourier domain and select a narrower distribution.

An interferometer uses three Raman pulses of differing lengths meaning they will interact with different velocity classes of atoms. To ensure that all three pulses will interact with all of the atoms present in the first pulse it is desirable to perform an initial pulse with a more stringent velocity selective condition and blow away the hotter atoms before the first interferometer pulse. A lower temperature cloud is desirable as it will lead to an increase in the percentage of atoms which will be transferred by this selection pulse, increasing the value of N in Eq. 1.9. For a square pulse of duration τ we take the cut-off frequency in the Fourier domain as the full-width half-maximum (FWHM) of the associated $sinc$ function, given by

$\omega_{\mathrm{FWHM}} = 1/\tau$ [60]. For a laser of frequency ω_{atoms} then the Doppler shifted frequency seen by the atoms moving with velocity, v, is

$$\omega_{\mathrm{atoms}} = \omega_{\mathrm{laser}} \left(1 - \frac{2v}{c}\right). \tag{2.17}$$

The factor of 2 in the Doppler term arises from the two counter-propagating Raman beams and the velocity condition must be matched for both. The Raman laser can only drive the transition for atoms whose Doppler shift is less than this FWHM value. Therefore by taking $\omega_{\mathrm{laser}} - \omega_{\mathrm{atoms}} = \omega_{\mathrm{FHWM}} = 1/\tau$ it is possible to use Eq. 2.17 to determine the value of the cut-off velocity,

$$v_{\mathrm{cutoff}} = \frac{c}{2\omega_{\mathrm{laser}}\tau}. \tag{2.18}$$

Using the Boltzmann velocity distribution for the post-molasses atoms and v_{cutoff} for τ =50 μs[61] it is possible to determine the percentage of atoms that will be transferred. To do this for the gravimeter system it requires characterisation of the cloud temperature, which can be obtained via light-sheet detection[62]. This is a type of fluorescence imaging where a thin sheet of light is used to illuminate a subset of the cloud as it falls through the thin sheet and the emitted light is detected with a photodiode perpendicular to the beam. The signal on the photodiode will be proportional to the number of atoms illuminated, so if this is placed in the path of a moving atomic cloud it can be used to measure the spatial extent of the cloud in the direction of motion. After sufficient free expansion times, the spatial distribution of the cloud is dominated by its thermal nature allowing a temperature to be determined.

On the prototype gravimeter using a 7 mm wide and 0.3 mm thick light sheet it has been demonstrated that by modifying the light frequency and intensity the temperature of the atom cloud can be lowered via a process known as optical molasses[63]. With the introduction of a more developed control package it was possible to implement this cooling step in a sequence scheme similar to that in[26]. Figure 2.9 shows the result of a typical light sheet measurement performed on the gravimeter. The fit in this data takes the Gaussian form

$$y(x) = Ae^{-\frac{(x-x_0)^2}{2\sigma^2}} + y_0. \tag{2.19}$$

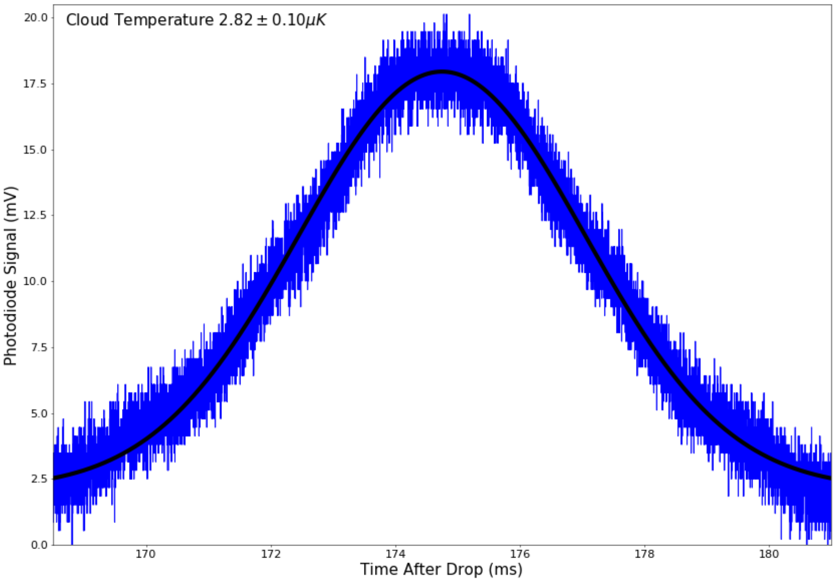

Figure 2.9: Gravimeter light sheet measurement.

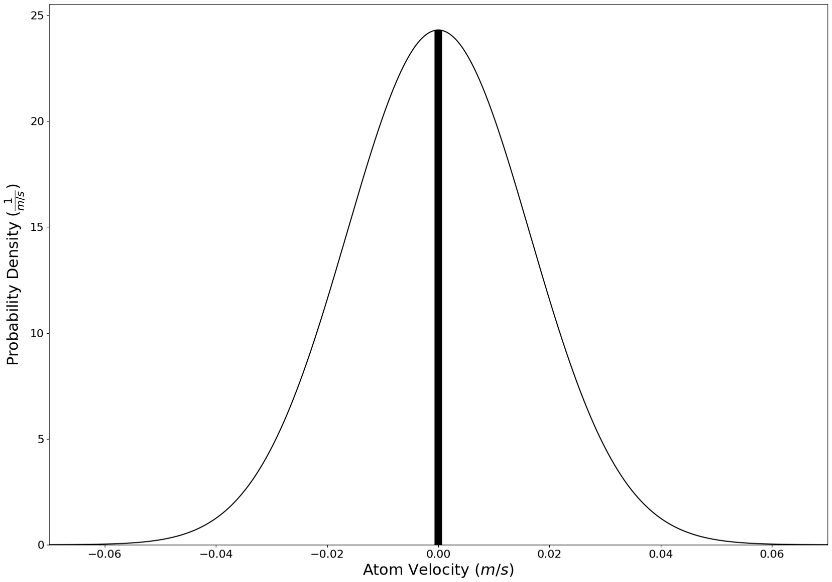

Figure 2.10: Boltzmann distribution of a 2.82 μK cloud. The shaded region shows the atoms that would be transferred by a 50 μs Raman pulse.

The width of the cloud is taken as the FWHM of the Gaussian (2σ), giving the length of time it takes the cloud to cross the light sheet. This can be converted into a cloud diameter by determining how quickly the cloud is moving after accelerating under gravity for the duration of the drop (t_d). Giving a final cloud diameter of

$$d_\mathrm{f} = 2g\,t_\mathrm{d}\sigma. \tag{2.20}$$

Using the initial cloud diameter, d_0 =2 mm[1], to find the change in size over the course of the drop for atoms with mass m and some drop time t_d gives a relationship for the temperature[64]

$$T = \frac{\left(d_f^2 - d_0^2\right) m}{k_\mathrm{B}\, t_d^2}. \tag{2.21}$$

This gives a temperature of 2.82 μK for the data in Fig. 2.9. Allowing the creation of a Boltzmann velocity distribution that will be representative of the post-molasses atom cloud. Considering a velocity selection pulse of duration 50 μs[61] (corresponding to a temperature of ~1 nK) it is possible to determine the percentage of atoms that will take part in the interferometry sequence using Eq. 2.18. Figure 2.10 presents this graphically where the shaded region shows the atoms that would be transferred by the Raman pulse. From this it can be seen that, for an atom cloud of 2.82 μK, a 50 μs velocity selection pulse would transfer ~ 3% of the total atoms. The loading curve for the gravimeter (Fig. 2.8) gives an atom number after a 3 seconds loading time of ~ 1.75×10^8, resulting in ~ 5.25×10^6 atoms present in the interferometry sequence. This is almost double the 3×10^6 atoms assumed in 1.3 for a good single shot sensitivity, suggesting that the temperature possible in a prism MOT gravimeter is sufficient to meet the sensitivity requirement.

2.3.3 Atom Interferometry

Having validated the design of the gravimeter for atom number, pressure and cloud temperature the final step was to ensure that it was possible to perform atom interferometry in a representative environment. The atom used in this system is ^{87}Rb, chosen for the good synergy between its fortunate level structure

---

[1] An approximate value based on the relative size of the MOT relative to the prisms when viewed on the monitoring camera

(as shown in Fig. 2.11) and the easily available telecoms technology. The telecoms standard 1560 nm lasers are easily frequency doubled to address the 780 nm $D2$ line of ^{87}Rb. Additionally the presence of the $|F = 2\rangle \longrightarrow |F^{'} = 3\rangle$ transition means that only a single low power repump laser is required to keep atoms out of the $|F = 1\rangle$ dark state. The 6.834 GHz splitting between the two ground states is a frequency that is easily generated by off-the-shelf components.

Atom interferometry manipulates atomic states and trajectories through the use of $\pi/2$ and π duration Raman pulses; these pulses induce a Rabi oscillation between the two hyperfine ground states of ^{87}Rb with frequency Ω_R. If all atoms begin in the lower ground state $|F = 1\rangle$ then the probability, $|c_{F=2}(\mathrm{t})|^2$, of being in the excited state $|F = 2\rangle$ after some time, t, is[66]

$$|c_{F=2}(\mathrm{t})|^2 = \sin^2\left(\frac{\Omega_R t}{2}\right). \tag{2.22}$$

The π and $\pi/2$ definitions arise from the argument of the sin term. If a population inversion is desired then after some time $t^{'}$ the probability of being in the excited state would be $|c_{F=2}\left(\mathrm{t}'\right)|^2 = 1$ meaning that $\Omega_R t^{'} = \pi$ and for a population split it would require that $\Omega_R t^{'} = \pi/2$. To characterise this in an experiment a multi-shot scan is performed where a Raman pulse of variable duration is performed, the relative population of the two states $|F = 1\rangle$ and $|F = 2\rangle$ is taken. Plotting this ratio as a function of the applied pulse duration a Rabi oscillation can be scanned out as in Fig. 2.12 for the given setup.

A measurement of this type requires sub-microsecond timing control as typical pulse durations are of the order of 10-20 μs. Fast control is provided by the AOM on the output of the laser box driven by a thermally stable DDS source and controlled from a home made FPGA control system.

In Fig. 2.12 the values of the $\pi/2$ and π times have been found by fitting a damped sin curve to the data of the form[67]

$$y(t) = Ae^{-\tau/t}\cos\left(\omega t + \phi_0\right) + y_0, \tag{2.23}$$

Where A is the oscillation amplitude, τ the decay constant and oscillation frequency ω. The two constants y_0 and ϕ_0 are included to account for the offset from zero at the start and would ideally both be zero. Figure 2.12 shows the result of a calibration measurement for the prototype gravimeter, where rather

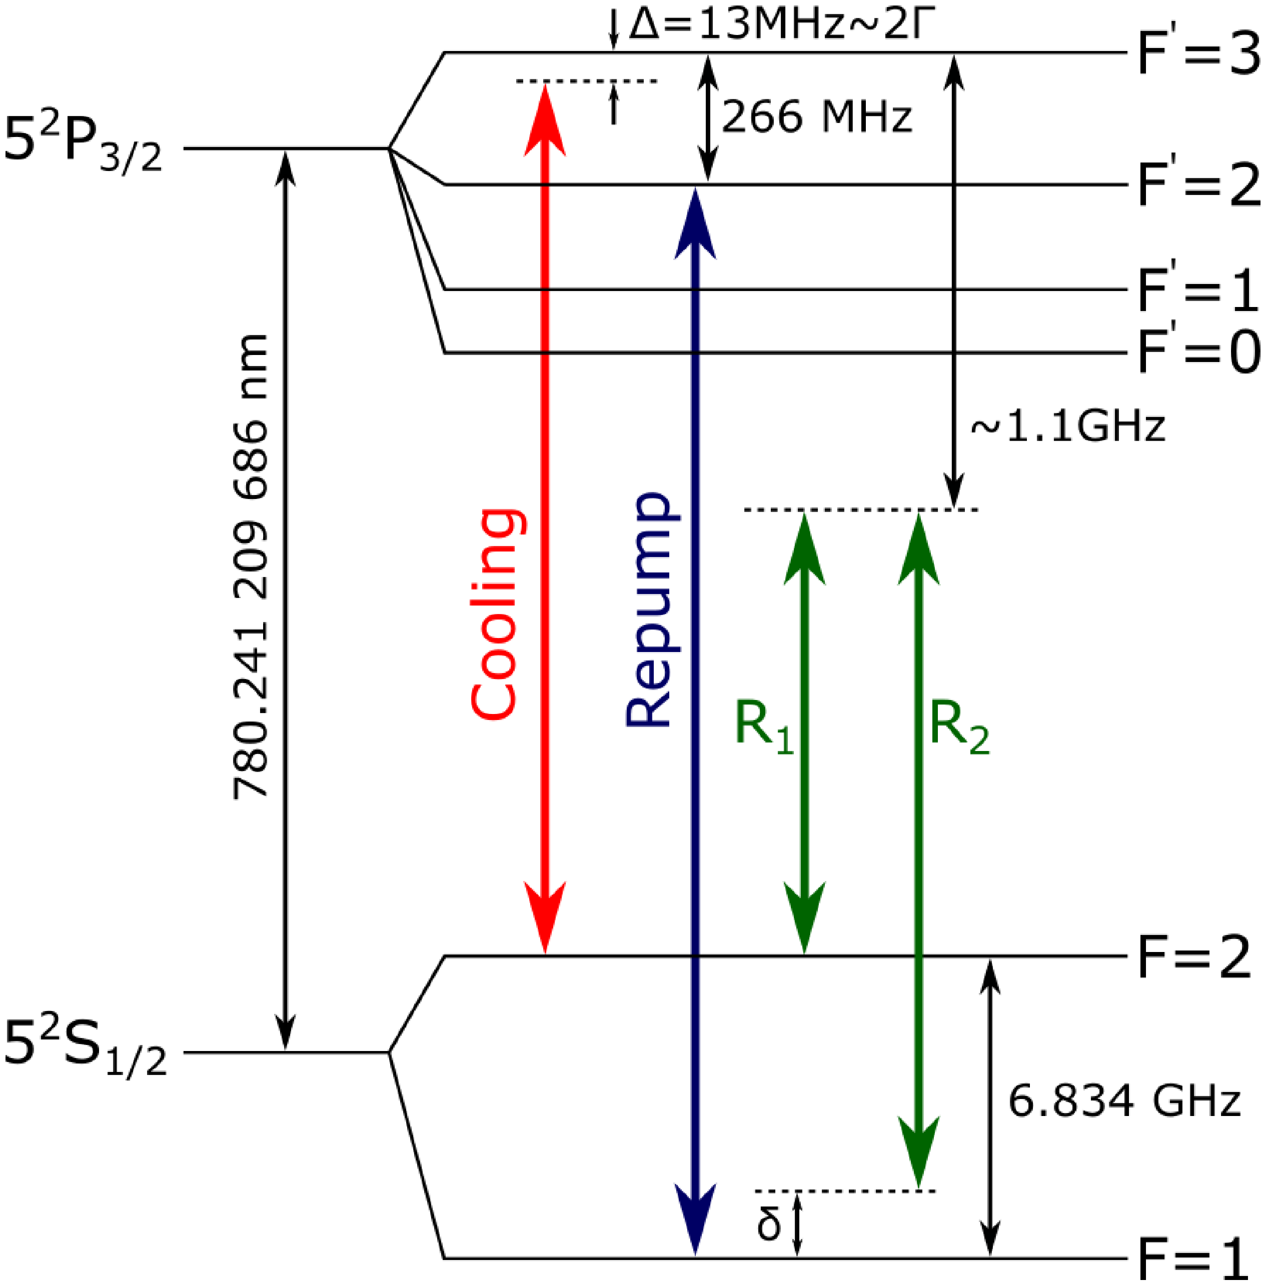

Figure 2.11: Level structure of the D2 transition of ^{87}Rb (adapted from [65]). The cooling and repump beams are used in the MOT and optical molasses stages to cool the atoms and keep them out of the dark state. R_1 and R_2 are the two frequencies of light required to drive the two photon Raman transition.

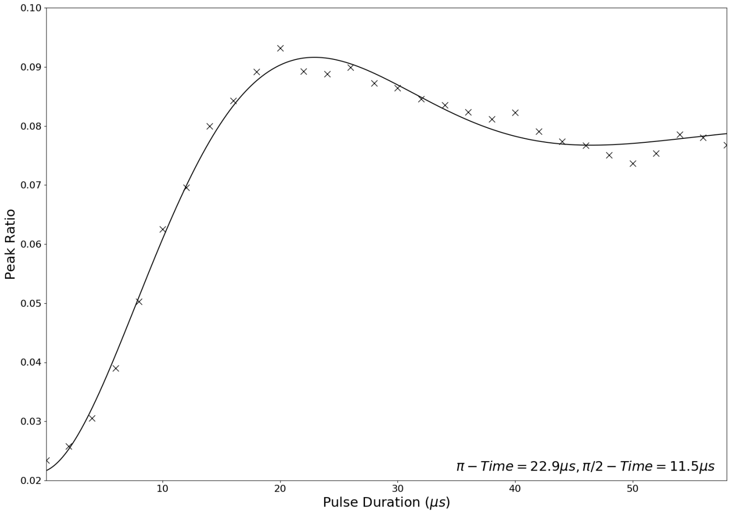

Figure 2.12: Rabi oscillation for the prototype gravimeter during system tests outside of the lab.

than 100% transfer the maximum transfer is more like 9%. This damped nature arises from other loss mechanisms, such as single photon interactions with the excited state and the thermal nature of the cloud. This sets a a maximum possible contrast on the gravimeter of 9%.

The internal energy states $|g\rangle$ and $|e\rangle$ are directly linked with their external momenta as the higher energy state will only be occupied after interaction with the light field[68]. This allows the two states to be labelled with both their internal energy and external momenta. Thus these states become $|g, p\rangle$ and $|e, p + \hbar k_{\text{eff}}\rangle$ where $\hbar k_{\text{eff}}$ is the momentum given to the atoms by the light.

An atomic interferometer utilises a series of $\pi/2$ and π pulses that act analogously to beam splitters and mirrors in an optical Mach-Zehnder interferometer. The initially clean cloud of atoms in the $|g, p\rangle$ state is placed in a superposition of two states with a $\pi/2$ pulse. As this creates a momentum difference between the atoms in each of the two internal states it causes them to spatially separate (see Fig. 2.13). After some free propagation time, T, a π-pulse is applied that inverts the two states; leading one to lose momentum and the other to gain an equal amount of momentum causing them to be brought back together until they overlap after the same free propagation time. Once the states are overlapped a final $\pi/2$ pulse is applied to close the interferometer, if this measurement is performed in the absence of anything which can effect the atomic phase (e.g. gravity) it will recreate the original clean $|g, p\rangle$ state.

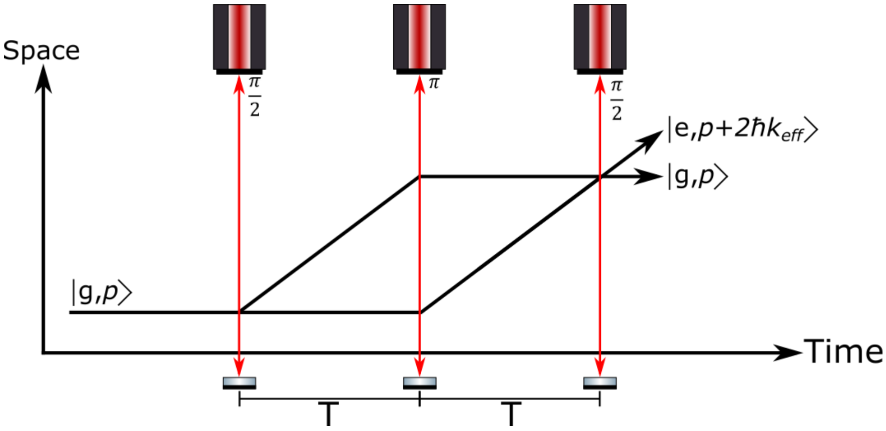

Figure 2.13: The motion of the two atomic clouds in the absence of an external phase shift, in the ideal the final $|2, p+2\hbar k_{eff}\rangle$ state will be empty.

This is analogous to the optical case of a Mach-Zehnder interferometer (Fig. 2.14) where in the absence of a phase shift, i.e. $\Delta\phi = 0$, the signal of one output port will be maximised and in the other will be minimised. If the value of $\Delta\phi$ is scanned this additional phase shift will cause the signal seen on both diodes to oscillate.

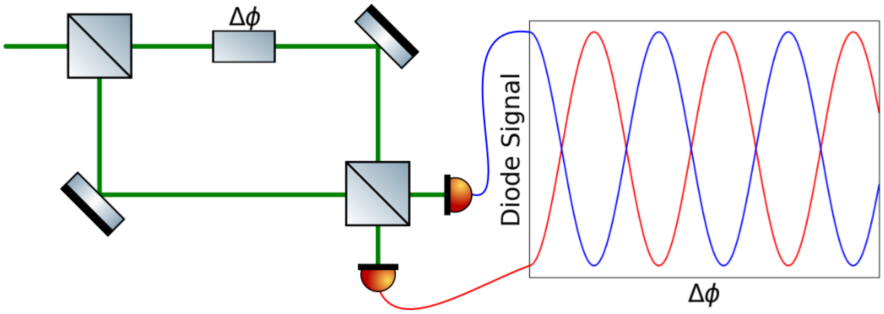

Figure 2.14: Optical Mach-Zehnder Interferometer, $\Delta\phi$ can be seen as analogous to gravity in the atomic case.

For an atomic interferometer this is the same as introducing a phase difference between the two interferometer arms, such as gravitational acceleration. This acceleration will induce a Doppler shift on the light seen by the atoms, moving it away from the resonant case as δ increases. In this off resonant scenario there will be a change in the Rabi frequency of the atoms. The shift to the Rabi frequency means that a π-pulse in the resonant case will no longer be a complete π-pulse in the off resonant case, leading to an imperfect state inversion. This will result in leaving some atoms in the $|e, p + \hbar k_{\text{eff}}\rangle$ output state of the interferometer, it is the total number of atoms in this excited state that is of interest as this state can only be populated due to phase shifts over the drop. Phase shifts will arise from both the gravitational acceleration and any phase change of the laser between Raman pulses, as the phase of the light is imprinted onto the atoms at each pulse. It can be shown[26] [32] that the total phase difference between the two paths can be given by

$$\Phi_{\text{total}} = \Phi_{\text{upper}} - \Phi_{\text{lower}} = k_{\text{eff}} g T^2 + (\phi_1 - 2\phi_2 + \phi_3). \tag{2.24}$$

Where ϕ_n $(n = 1, 2, 3)$ is the phase of the laser at the instant of each of the three pulses. In a similar way to Eq. 2.22 it can be shown that the population of the excited is given by[32]

$$P_{|e,p+\hbar k_{\text{eff}}\rangle} = |C_{|E,p+\hbar k_{\text{eff}}\rangle}|^2 = \frac{1}{2}\left[1 - \cos\left(\Phi_{\text{total}}\right)\right]. \tag{2.25}$$

As the atoms fall they accelerate under gravity, this will lead to a frequency difference between the Raman lasers that is time sensitive in the atomic frame, after time t_{drop} the frequency shift seen by the atoms $\omega_D(t) = -k_{\text{eff}} g t_{\text{drop}}$. This will increase δ leading to a reduction in the transition probability and a loss of fringe contrast. The frequency difference therefore needs to be chirped with rate α giving a frequency difference in the lab frame of

$$\omega(t) = \omega(t_0) + \alpha(t - t_0). \tag{2.26}$$

If $t_0 = 0$ is taken as the time of the first Raman pulse then the phase shift due to the chirped frequency between two pulses is given by $\Phi_{\text{chirp}} = \alpha T^2$. Using this to modify Eq. 2.24 and then substitute into Eq. 2.25 the population of the readout state of the interferometer is

$$P_{|e,p+\hbar k_{\text{eff}}\rangle} = \frac{1}{2}\left[1 - \cos\left(\left(\alpha - k_{\text{eff}} g\right) T^2 + \phi_{\text{laser}}\right)\right], \tag{2.27}$$

Where the bracketed term in Eq. 2.24 has been replaced by the net laser phase ϕ_{laser}. Performing an atomic measurement of gravity in sensitive systems such as those in [32] and [18], it is the value of α that is scanned for multiple T-times (see 5.3). The chirp rate is also used to keep the Raman lasers on the central fringe and if ϕ_{laser} is scanned a sinusoid is traced, this is the method used to generate a gravity fringe in the prototype gravimeter. Keeping $\phi_{1,2} = 0$ the total phase can be set by scanning ϕ_3.

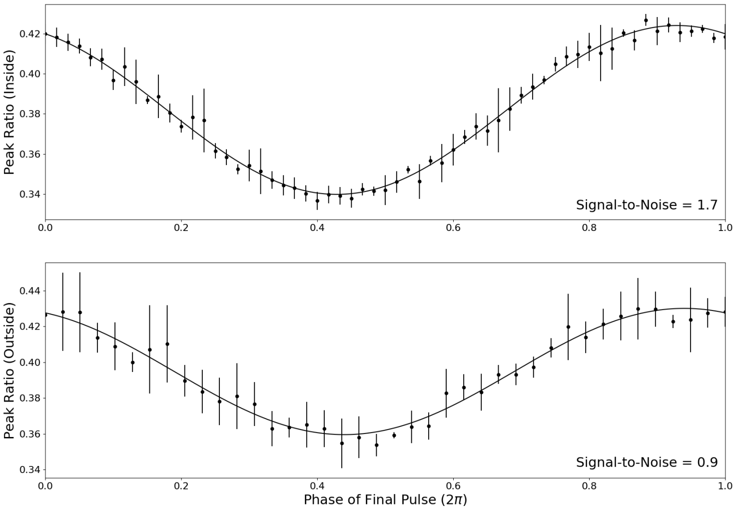

Figure 2.15: Comparison of fringes generated both inside (upper plot) and outside (lower plot) of the lab environment for a T-time of 1 ms.

Figure 2.15 shows a comparison between two sets of data, each of which is the average of 5 complete measurement cycles. The signal-to-noise values listed are the ratio of the RMS of the errorbars to the fit amplitude. The lab-based measurements were those taken the day before the external measurements and the second set were taken in the middle of a warm, sunny courtyard (Fig. 2.16), as can be seen the amplitude and the phase of the fringe are essentially unchanged. However the noise on the data is fairly substantially increased when outside, arising from increased thermal fluctuations and large variation in background light levels. Over the course of these measurements it was found that heat dissipation in the flight case provided a serious problem, as highly efficient packing does not leave much space for air flow, introducing thermal noise on some component such as the VCOs causing them to drift over MHz as the ambient temperature changed by degrees. Neither measurement used a vibration isolation stage meaning that without the common mode suppression a gradiometer benefits from the acceleration of the reference mirror will couple into the measurement.

These problem areas were targeted for improvement in the second generation of the system. The fact that with little adjustment to the system a fringe with a comparative signal-to-noise was generated in both laboratory and external environments confirmed the robustness of the design. This final validation step gave the confidence to move forward with these designs in a gradiometer configuration and provided a direction in which to progress to improve the system.

Figure 2.16: The gravimeter prototype in its first foray outdoors to what has become the preliminary testing ground for all portable systems.

2.4 Implementation in Gradiometry

Having proven the viability of the prism MOT gravimeter the design needed modification in order to function as a gradiometer and benefit from the common-mode noise suppression. A number of potential implementations were considered and ultimately discarded in favour of opposite co-axial dual beam delivery. This section will outline these different options and briefly compare them to explain the reasoning behind the final choice. To make this comparison we will refer back to the discussion of capture velocity and effective capture volume presented in 2.1.1.

2.4.1 Prototype Gravimeter Capture Volume

In all real MOT systems the intensity of the cooling light will clearly not be infinite. However it is possible to compare particular MOT geometries by modelling the effective capture volume with a finite and non-uniform intensity distribution and then compare this to the infinite intensity case. To better explain this let us consider the case of the prototype gravimeter where the MOT intensity is represented by the Gaussian intensity distribution

$$I(x) = I_0 e^{-\frac{(x-x_0)^2}{2\sigma^2}}. \tag{2.28}$$

Figure 2.17 shows an overlay of this distribution with the position of the prisms in the chamber. The central intensity I_0 =2.55 mW/cm^2 has been set according to that measured in the cooling light for the real system and the Gaussian width, 2σ =21 mm, has been set by taking the required diameter of the telescope to be 6σ, ensuring that 99.7% of the power in the beam is accounted for. The value of x_0 can be used to account for position offset between the input telescope and the prisms but here is set to zero. From this it is clear that the capture velocity for the central 2 cm will be much higher than that for the two reflected beams, the symmetry of the Gaussian about 0 cm means the two side beams can be treated as identical.

From this two values of average intensity are found by integrating over the regions of interest; one for the vertical ($I^{\mathrm{V}}_{\mathrm{avg}}$ =2.37 mW/cm^2 $\approx 1.4 I^{\mathrm{sat}}$) and one for the horizontal ($I^{\mathrm{H}}_{\mathrm{avg}}$ =1.08 mW/$cm^2$$\approx 0.65 I^{\mathrm{sat}}$). It is immediately clear that the relationship given in Eq. 2.6 will not hold as these intensities are far below I^{sat}. Calculating the scattering force directly from 2.1 causes Eq. 2.8 to take on the form

$$v_{\mathrm{c}} = \sqrt{\frac{\hbar k \Gamma L}{m} \frac{I/I^{\mathrm{sat}}}{1 + I/I^{\mathrm{sat}} + 4\delta^2/\Gamma^2}}. \tag{2.29}$$

Where the velocity dependence of the force has been neglected for simplicity and all terms are as defined in 2.1.1. Using Eq. 2.29 it is possible to determine the horizontal and vertical capture velocities. Plotting Eq. 2.8 for beam sizes between 0-2 cm it is possible to compare these velocities and find the equivalent beam size for the infinite intensity case. From Fig. 2.18 we can see that the capture velocity in this low intensity case is the same as using substantially smaller beams with a much higher intensity. Using these values the effective capture volume is only 340 mm^3, about 4.3% of the maximum 8000 mm^3 for the infi-

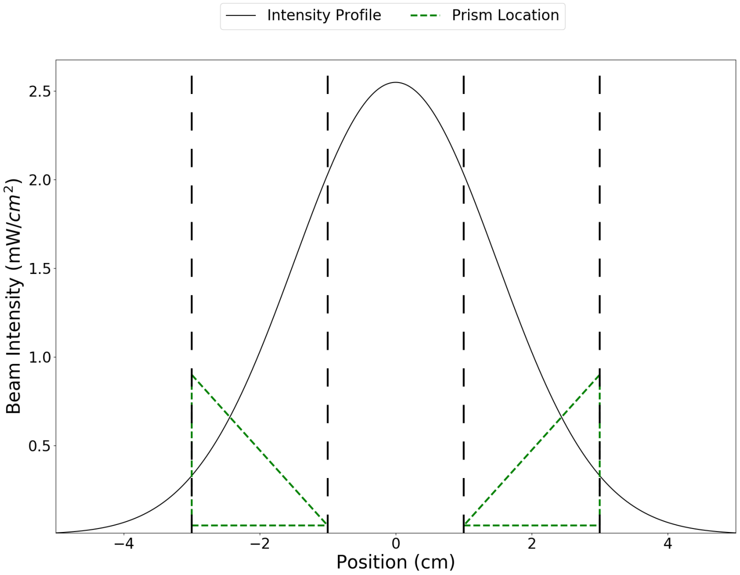

Figure 2.17: Gaussian intensity profile overlaid with the prisms in the relative locations used for the prism MOT.

nite intensity case.

In the infinite intensity case the detuning term is no longer important as it becomes negligible next to the intensity term. Thus in the non-infinite case part of the drop in effective capture volume is due to the impact of this term. The simulation was repeated with $\Delta = 0$ producing an effective capture volume of 8.2% of the maximum (~655 mm^3) compared to the 4.3% with the 13 MHz detuning used in the MOT. This shows that the detuning term will have an impact on the result but the intensity ratio is by far the leading order contribution. The cases considered in 2.4.2 use a non-zero detuning to be more representative of the real world case.

This method can be used to optimise the size of prisms with respect to the predicted input optical power. It is possible to see how this may lead to a situation where a higher atom number can be achieved with smaller cooling beams. This could be done by imposing a top hat intensity profile onto a given geometry leading to a more uniform capture velocity or by reducing the overall system and telescope size leading to a better intensity distribution over the prism area.

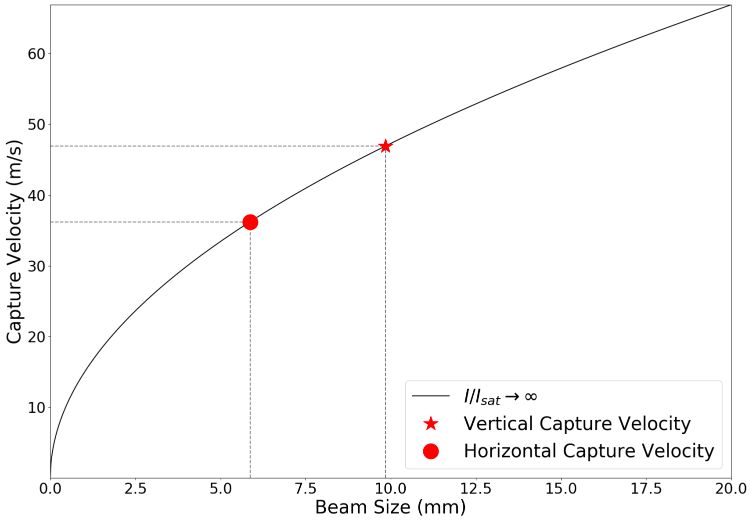

Figure 2.18: Capture velocity with beam size comparing the real gravimeter prototype system to an idealised infinite intensity case.

2.4.2 Gradiometer Geometry

In order to create a prism MOT gradiometer the core requirement is that there be two prism regions separated by the desired baseline. A number of potential solutions to this problem are presented in Fig. 2.19, however a decision had to be made as to which would be implemented. Initially the most attractive design was one with two entirely independent MOT regions (Fig. 2.19a), where the cooling telescopes were perpendicular to the Raman beams. This implementation was low risk as each MOT system could be fully independent whilst enabling the Raman beam to couple the two systems together. However this also presented some major drawbacks. As the Raman beam would need clear line of sight between the two clouds it would have to fit between the prisms in the two regions therefore meaning it would need to be small. This would lead to a greater intensity gradient over the cloud and a greater reduction in interferometer contrast. Additionally implementing this design would move away from the compact cylindrical form factor developed in the gravimeter leading to a much larger and heavier sensor. This idea was therefore discarded as too heavy and not sufficiently sensitive.

The other four designs are different implementations of that used in the gravimeter, to distinguish

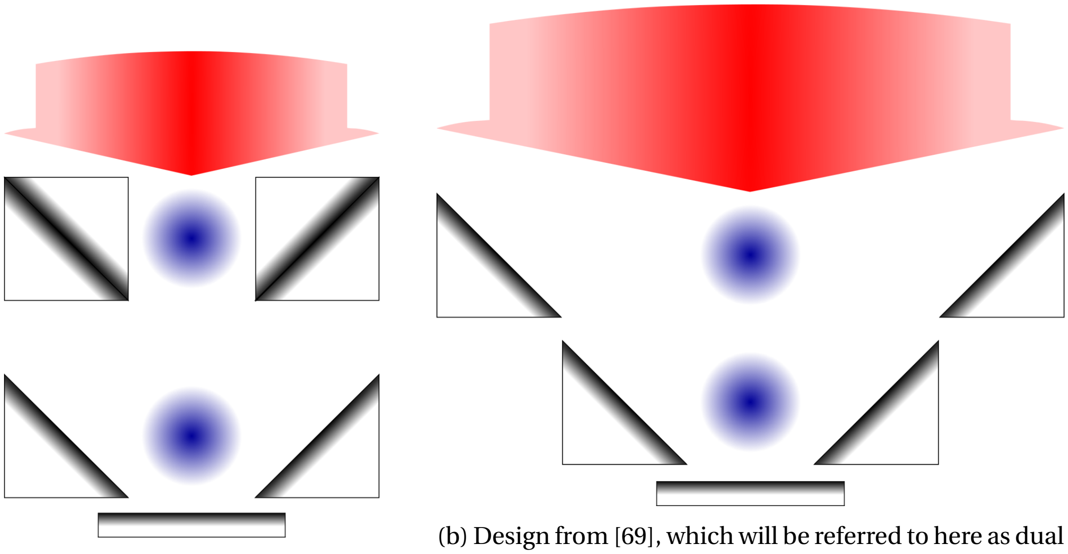

(a) Beam splitter (BS) gradiometer.

(b) Design from [69], which will be referred to here as dual radius (DR).

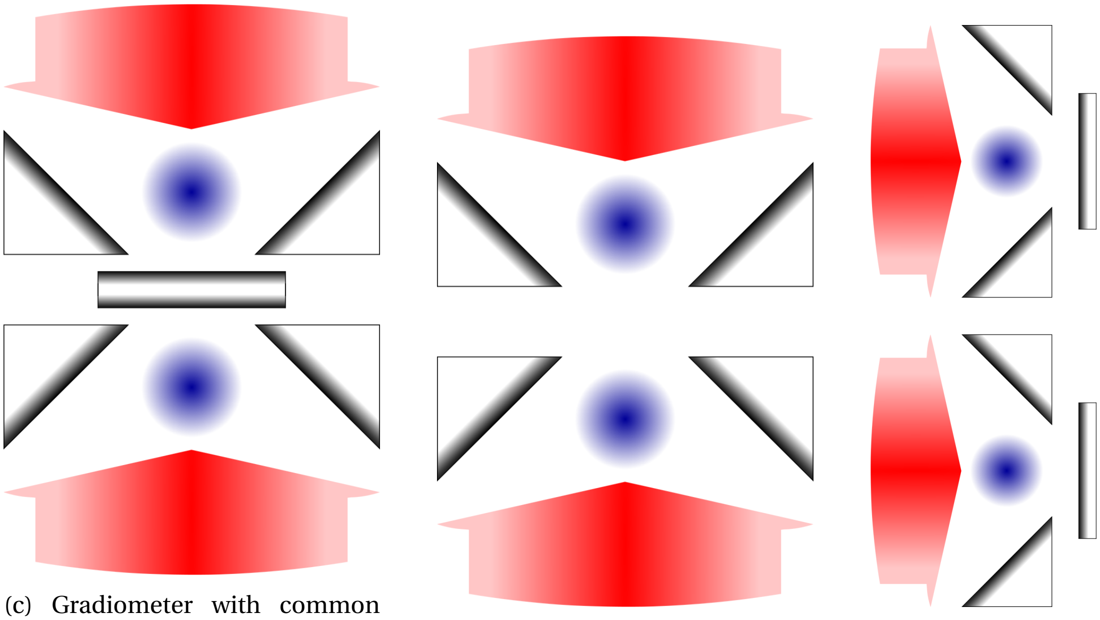

(c) Gradiometer with common central mirror and two beam delivery.

(d) Gradiometer without central mirror but with two beam delivery.

(e) Off axis beam delivery.

Figure 2.19: Potential Gradiometer implementations.

between them let us consider the model for capture volume developed in 2.1.1. For this discussion the designs which take advantage of two counterpropagating delivery telescopes (Fig. 2.19c and Fig. 2.19d) shall be considered as the same dual telescope (DT) approach. That using two sets of prisms in different parts of the much larger input beam (Fig. 2.19b) will need to be considered as two separate regions as they will lie in very different parts of the intensity profile, these are the dual capture regions (DR). For the beamsplitter (BS) implementation in Fig. 2.19a it is assumed that the beamsplitters are perfect 50 : 50 beam splitters.

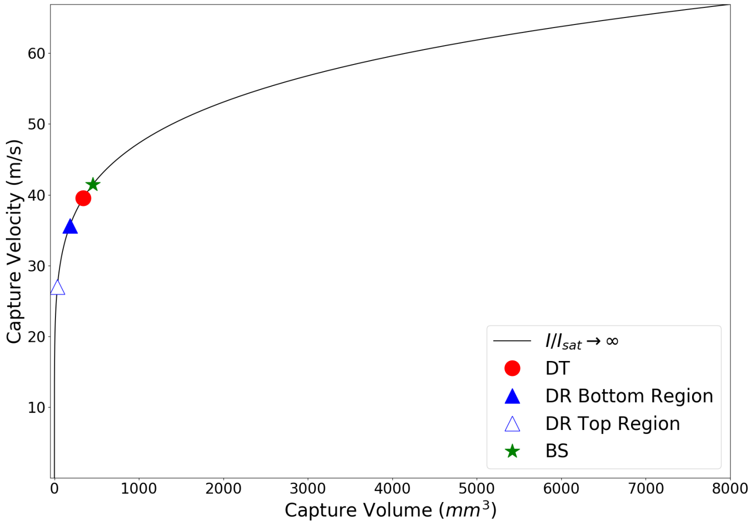

Figure 2.20: Capture velocity versus capture volume for the designs shown in 2.19a, 2.19b and 2.19c

The values used in the model are the same as those used in the example in 2.4.1, $I_0 = 2.55\,\mathrm{mW/cm^2}$ and σ =21 cm for prisms of length 20 cm. For the DR case the value of σ has been scaled up to account for the extra 40 cm region arising from the additional prisms and the value of I_0 has been scaled down to ensure that there is the same total optical power in the prism region. Treating the capture region as a cuboid the capture volume is then

$$V_c = d_\mathrm{H} \times d_\mathrm{H} \times d_\mathrm{V}. \tag{2.30}$$

Where d_H and d_V are the effective horizontal and vertical beam sizes respectively. The ideal volume curve is taken by cubing the infinite intensity beam size, d^{∞}, and referencing it to the same capture velocity, that is

$$V_c^{\infty} = d^{\infty 3}. \tag{2.31}$$

The same comparison between v_c and the ideal measurement is made to give the effective capture volume. The important comparison here is between the volumes of the different implementations; as discussed in 2.1.1 the atom number, N, is proportional to $V_c^{4/3}$, and Eq. 1.9 shows that the single shot sensitivity scales with $1/\sqrt{N}$. This means that the sensitivity of the sensor will scale as $1/V_\mathrm{c}^{2/3}$, meaning that the larger the capture volume the more sensitive the gradiometer.

The results in Fig. 2.20 show that the most effective use of optical power is the BS approach giving an effective capture volume for one MOT region of 5.6% of the infinite intensity case, with the DT approach at 4.3% of maximum and the two DR regions at 2.2% and 0.4% of the maximum. From this it is clear that the DR approach is much less efficient than the other two approaches. The BS design has a higher capture volume than the DT design however the outputs of the 50 : 50 beam splitters would not be perfectly stable and would vary with polarisation of the input light, leading to an uncommon drift in both atom number and cloud temperatures. These would both produce an uncommon phase noise on the ellipse (see Chapters 5 and 6 for details), thus this approach was discarded.

This left the option of the DT approach, as the two implementations are sufficiently similar it was possible to attempt both with a single system. Using the modular design of the gravimeter a gradiometer was implemented by placing two of these designs facing each other and in the middle a double-sided mirror was mounted with the option to remove it and attempt the alternate implementation. Both designs have their own benefits; by keeping the modules separate using a central mirror, the MOT intensity within a chamber is entirely common for a given system and will provide the stability seen in Fig. 2.6, however it means that the Raman beams are also decoupled so that intensity or polarisation drifts would introduce uncommon phase noise onto the measurement. On the other hand removing the mirror would have the inverse effect; the vertical MOT beams would now be common between the chambers but uncommon in a single MOT system and could thus suffer from the instabilities seen in the six beam MOT data from Fig. 2.6, but the Raman beams would be perfectly common and thus any phase noise would be common between the two clouds. As some of the results in Chapter 5 will show, the approach with the mirror was initially implemented and provided a more stable MOT source, but the mirror was ultimately discarded as the uncommon noise on the Raman beams was seen to be a greater concern.

CHAPTER 3

SENSOR PACKAGE

Manipulating atoms to sub-μK temperatures requires isolating them from the external environment: vacuum prevents collisions between the cooled, trapped atoms and the hot atmospheric gas and magnetic shielding is required to allow precise manipulation of the magnetically sensitive atomic substates (see 2.1 and 5.4.1). The sensor package supplies this controlled environment for the atoms and the light is generated by the control package as in Chapter 4.

The discussion in this chapter will continue from the end of Chapter 2, providing details of the specific gradiometer implementation. Utilising two independent gravimeter modules (modified from 2.3) and placing one above the other in a gradiometer configuration (Fig. 3.1). The optical input for the chamber delivers both the Raman and cooling light along the chamber's common central axis, by retro-reflecting the input Raman light off a common reference mirror it is possible to maintain complete independence in the two gradiometer arms whilst benefiting from the common mode suppression.

Removing the central reference mirror allows the light to directly couple the two gravimeter systems. The sensor has adaptive capabilities beyond this impactful change to the optical system. As will be shown in 3.4 each of the sensors was separately mounted into its own rigid structure providing the facility to make changes to the sensor package geometry with minimal changes to the complete system. It is possible to change the baseline, total sensor height and alter the system into a launch configuration taking advantage of the full chamber length, with only a few minor changes to the sensor infrastructure. This also allowed for easier manufacture and assembly whilst producing a less risky platform for transport; if there was a leak in a single sensor then the other would not be contaminated.

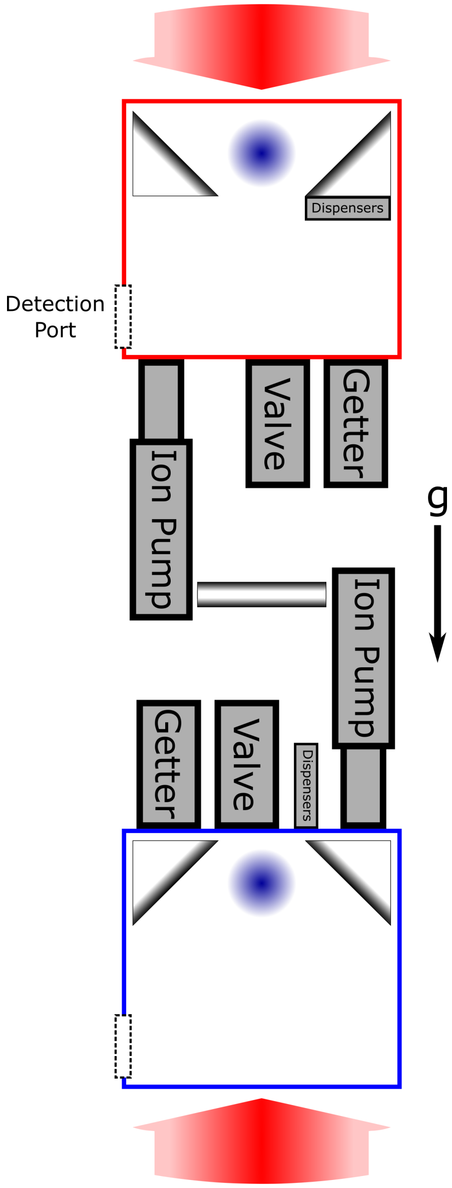

Figure 3.1: Final implementation of the gradiometer, two separate sensor heads (colour coded as 1.4) coupled together by a common mechanical structure, common light input and a central double sided mirror.

Feature	Value	Comments	Related Constraint
Drop Distance/Time	155 mm/$2T = 178$ ms	$2T = 366$ ms with launch	Sensitivity from Eq. 1.11
Cylindrical form Factor	Sensor Outer Diameter = 220 mm	Facilitates compact geometries as per Ch. 2	N/A
Weight of Sensor	~50 kg	Can be positioned by a single person	Survey Viability
Chamber Diameter	In Drop: 86 mm Maximum: 97 mm	Set by prism size of 20 mm	Sensitivity from 1.3
Independent Sensors	N/A	Eases manufacture/assembly and facilitates adaptability	N/A
Pressure	$\sim 5\times10^{-9}$ mbar/ Low 10^{-10} mbar	With/without dispenser load, measured with ion pumps	Eq. 2.11
Adaptability	N/A	Can be used to launch or drop atoms and test both DT geometries from 2.4.2	N/A
Coaxial Dual Beam Delivery	N/A	Required for cylindrical form factor	N/A
Umbilical Connection	5 m	Provides a single input connection	N/A

Table 3.1: Outline of key features of the experimental system.

Two way communication is provided between the sensor and control package (Ch. 4) with a 5 m umbilical connected to a distribution box at the base of the sensor. The umbilical delivers light and power to the sensor whilst providing transmission of diode signals and fibre delivery to the control package to monitor the sideband ratio on an optical cavity. This flexible connection enables surveys to be performed over localised regions, moving only the sensor whilst leaving the control package stationary. All the top level features outlined here and in earlier chapters are summarised in Table 3.1.

Each sensor head is a fully integrated atom interferometer in its own right with a design almost identical to that of the original prototype gravimeter; to better understand the way in which either sensor head fits together a cross section of the CAD model is presented in Fig. 3.2. Each sensor head features its own vacuum system to isolate the atoms, a double layered magnetic shield over the full length of the system and a second double-layered shield encompassing the central interferometry region. This provides passive magnetic shielding in the MOT and an increased shielding factor in the interferometry region, the high shielding in the central region and the bias coils provide a well defined field direction for the inter-

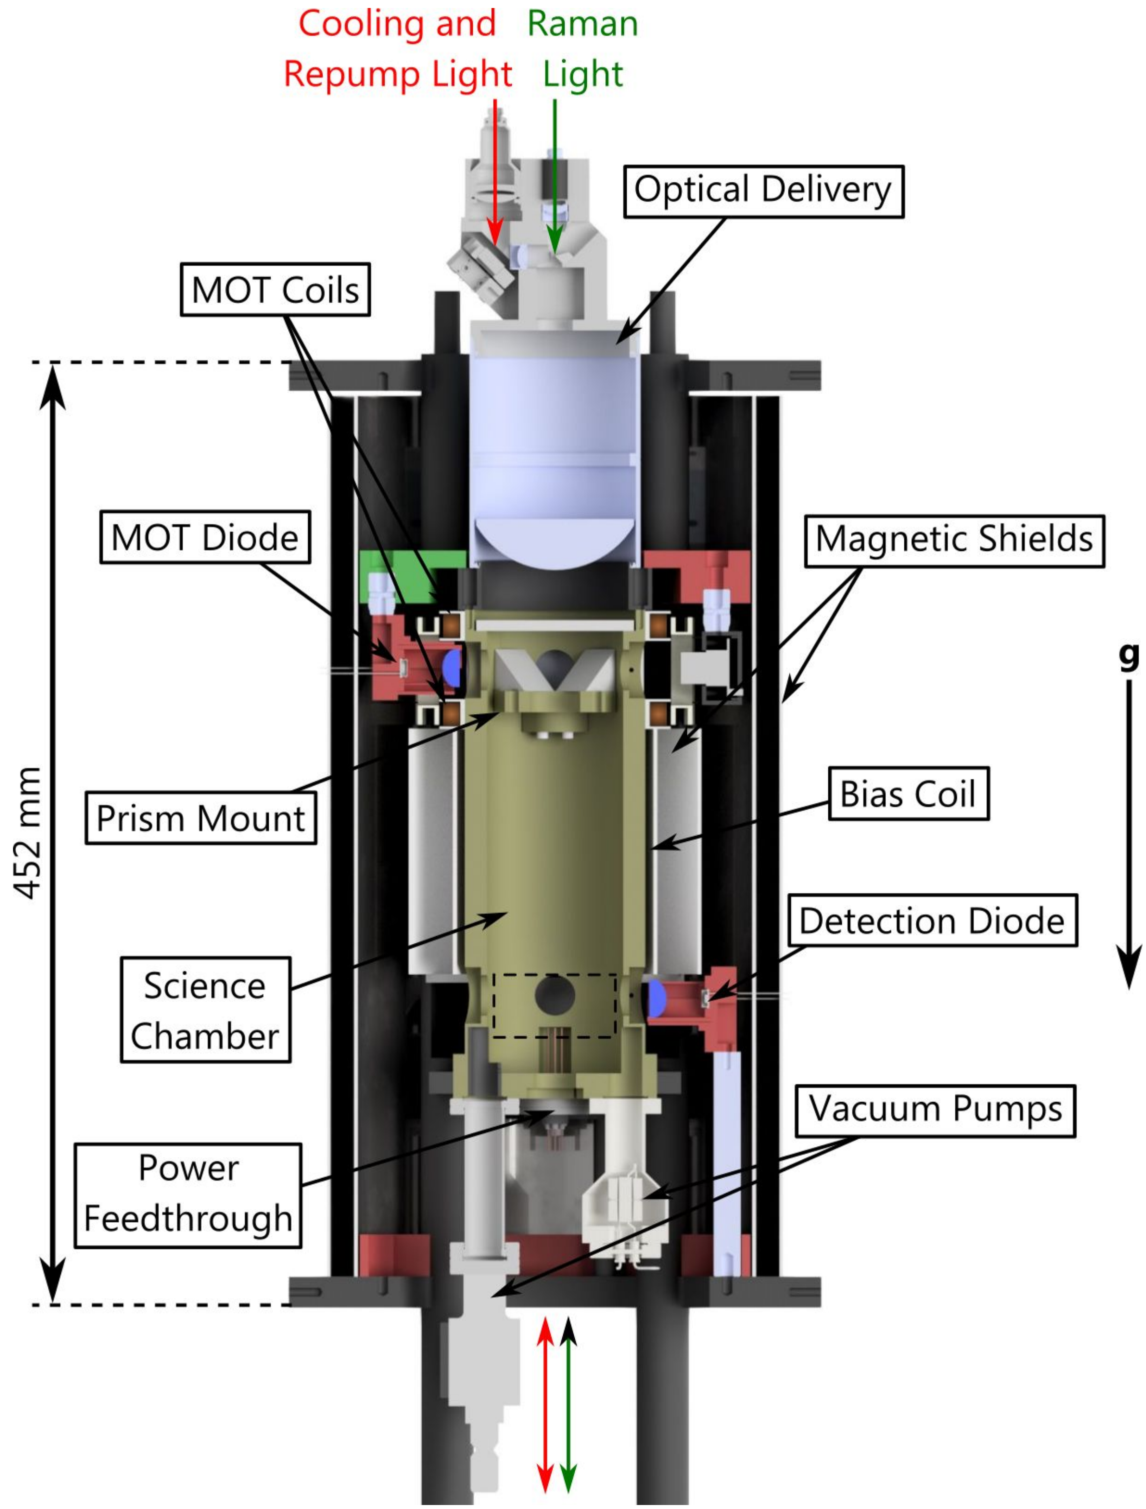

Figure 3.2: Cross section of the Top sensor head, the region enclosed by the black dashed region represents the location of the prisms in the Bottom sensor head.

ferometry sequence. A photodiode in the MOT region monitors the cloud during the atom preparation and a photodiode at the bottom of the drop provides the facility to read out the atomic states (see 5.1.3).

3.1 The Vacuum System

The vacuum system is little more complex than the diagram in Fig. 3.1 presents. The central science chamber is pumped with the same passive/active pumping components as the gravimeter (see 2.3.1), however the active pump has been placed on the end of a 76 mm long CF16 nipple to reduce the magnetic field introduced in the MOT region by the pump's permanent magnet. A CF16 valve is used to seal the system off from the atmosphere and a four-pin power feedthrough supplies current to the dispensers, all of which is mounted to the back of the chamber (see Fig. 3.4). Internally the prisms are secured to a separate mount that is lowered to the correct depth and glued to the wall of the chamber (see 3.1.1) and the ten windows are secured with indium seals (3.1.2).

Placing the active pump at the end of a nipple will act to reduce its effective pumping speed in the chamber due to the limited capacitance. Effective pumping speed at the entrance to the chamber (S_{eff}) is related to the true pumping speed, S, and the nipple conductance, C, by the relation[70]

$$\frac{1}{S_{\text{eff}}} = \frac{1}{S} + \frac{1}{C}. \tag{3.1}$$

As the length of the nipple, L =76 mm, is much greater than its diameter, d =16 mm, it is reasonable to treat it as a long round tube. Assuming that the gas inside the vacuum system is at the same temperature as the atmospheric gas outside the system ($T_{emp} \approx$293 K) the conductance of the nipple is given by[71]

$$C = 121\frac{d^3}{L} = 121 \times \frac{\left(1.6 \times 10^{-3}\right)^3}{7.6 \times 10^{-3}} = 6.5 \times 10^{-3}\,\mathrm{m^3/s} = 6.5\,\mathrm{l/s}. \tag{3.2}$$

The effective pumping speed of the active pump is therefore 1.5 l/s. A drop of 3% in total pumping speed for non-inert gases but a drop of 25% for noble gases, using this modified value for active pumping in Eq. 2.14 leads to an increase in the ultimate pressure of the system of $\sim 1.5 \times 10^{-11}$ mbar. The impact on the system pressure is a change of less than 2% whilst almost tripling the distance from the ion pump magnet to the MOT region, giving an order of magnitude reduction in the field seen by the atoms.

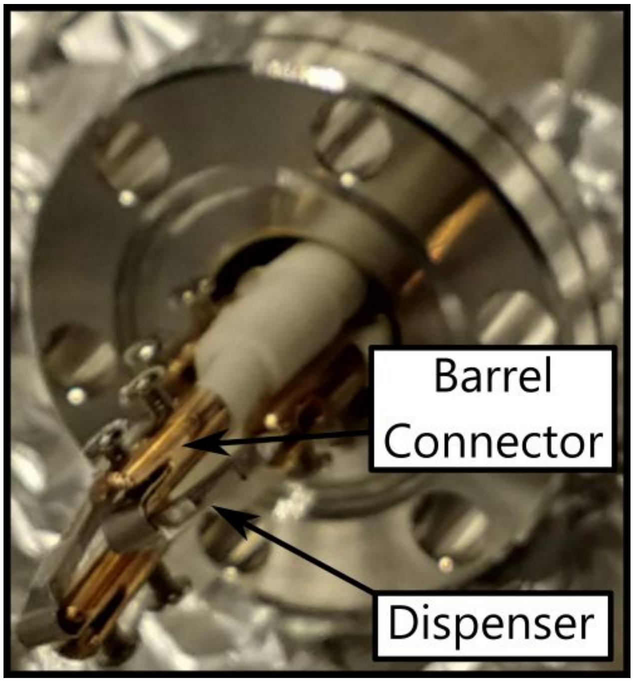

Figure 3.3: Two dispensers mounted onto a single four-pin power feedthrough with barrel connectors that have had a slit cut into them. Ceramic beads are used to ensure that the barrel connectors do not come into contact with the chamber wall, causing a short.

The four-pin copper electrical feedthrough[1] supplies the connection to the in vacuum rubidium dispensers. The Bottom chamber has the dispensers mounted directly to the feedthrough as in Fig. 3.3 while the Top chamber the prism mount (see 3.1.1) and connected to the feedthrough with four lengths of 1 mm Kapton insulated wire[2]. Placing the dispenser on the prism mount moved the atom source closer to the MOT region and was hoped to provide a high rubidium pressure in the MOT region whilst limiting the pressure during the drop. This was deemed unnecessary for the Bottom chamber as the gravimeter prototype proved the effectiveness of the feedthrough mounted design.

The science chamber is the region in which the atom cloud is prepared and manipulated as it provides space for the drop after the initial atom cooling stage. The science chamber consists of a tube of titanium 233 mm long; with an internal pocket 221 mm deep and an inner/outer diameter of 65/85.9 mm within which the atom manipulation takes place. At one end the chamber diameter increases to 97 mm to provide room for four CF16 knife edges to interface with the vacuum components, this presents the same minimal footprint as the MOT prototype (Fig. 2.3a).

To position the prism mounts at correct height and orientation, four shallow grooves (~0.5 mm deep with a radius of 32 mm) are machined into the internal edge of the vacuum chamber (see Fig. 3.4). These match protrusions on the prism mount (see 3.1.1) that were machined to produce a tight fit between the two components, by changing the length of the groove they can be used to set the depth of the prisms within the chamber. In the Top chamber the base of these grooves are 38.7 mm from the front window and 183.8 mm in the Bottom chamber. This allows an almost identical design to be used for both the Top and Bottom chambers, changing the length of the groove to fix the prism at the desired height.

3.1.1 Prism Mount

Increasing the chamber length to allow the cloud to drop meant it was no longer possible to reach the base of the system by hand to secure the prisms in place, to solve this the prisms were first aligned on a separate mount before lowering into the chamber. Alignment was achieved with the same technique as the MOT prototype, posts machined into the mount with a precision CNC machine. The depth of the prisms in the chamber was set by small grooves in the wall of the chamber (see Fig. 3.4), the radius of

---

[1] *Kurt J. Lesker EFT0043032*

[2] *Kurt J. Lesker FTAK10010*

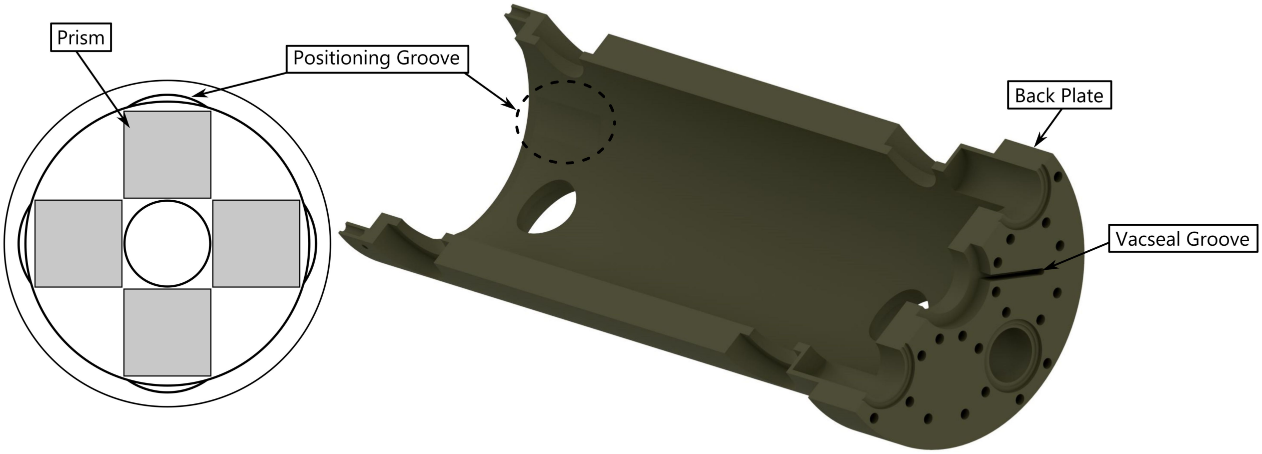

Figure 3.4: Left: Birds eye view of the chamber front showing the location of the positioning grooves relative to the prisms. Right: A cross sectional view of the Top science chamber where the depth of the positioning groove has been circled.

curvature of these grooves matches the curved protrusions in the prism mount. The close fit of the two parts meant it was difficult to lower the mount into position but aids in keeping it in place, between this precise fit and the use of *TorrSeal* at the base of the groove the prism mount was firmly secured.

The approximate quantity of *TorrSeal* in Table 2.3.1 is the estimated total surface area, arising from gluing the prisms to the mount and the mount to the chamber. To minimise the impact of this high gas load the mount was prebaked at 160 ° for roughly 12 hours, this is a necessary step due to the limits placed on the bakeout temperature by the indium (see 3.1.2). The relatively short drop times used in this sensor mean the gas load from *TorrSeal* does not limit the performance of the sensor (Eq. 2.11). In systems with longer drops this would be problematic so finding a replacement for *TorrSeal* is desirable for future development, with a mechanical solution being the ideal.

After the pre-bake the mount was lowered into the chamber with a specially designed tool (Fig. 3.5) using the M6 threads in the mount. Once the glue was given time to set a vigorous shake test was performed by hand. This was to emulate placing the system in transit as if it could not survive a manual test the vibrations in a moving vehicle would likely dislodge the optics.

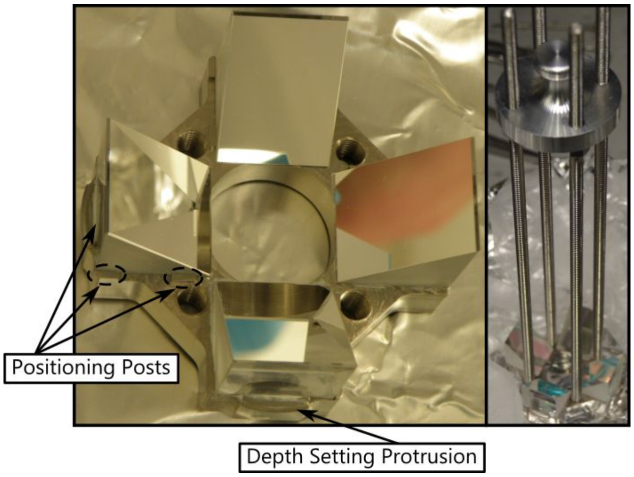

Figure 3.5: Left: Prism mount for the gravimeter system showing the parts used for mechanical alignment. Right: The tool used to lower the mount into position.

3.1.2 Indium Sealing

In the design of the sensor ten windows were required on each chamber for monitoring the atoms and to allow for delivery of the light and monitoring. Of these ten windows eight are 1 " diameter and 5 mm thick windows[1] used to provide side ports for monitoring the MOT or detecting the atoms, one is a 72 mm diameter and 5 mm thick window[2] at the front of the system to allow full illumination of the prism region and one is a 1 " diameter and 10 mm thick $\lambda/20$ flatness window[3] on the back of the system. These are affixed to the chamber with a soft metal indium seal. Indium grants an advantage over other soft metal seals, such as lead[72]; the compression flanges used to push the window onto the surface can be removed once the weld has formed. This allows the windows to be recessed into the chamber without protruding from the surface and maintaining the cylindrical form factor.

1 mm diameter indium wire is used to form the cold weld between the window and the chamber[73]. This bond can be formed without baking by preparing the indium with an acid bath before applying pressure[74]. This method was not an option due to a lack of appropriate health and safety equipment requiring an alternate solution. The compression on the windows is maintained whilst the system is pumped down and baked, the resulting pressure differential is sufficient to keep the window affixed to

---

[1] *Thorlabs WG11050-B*

[2] *Knight Optical WQK7501-AI1*

[3] *Newport 10BW40-30AR.16*

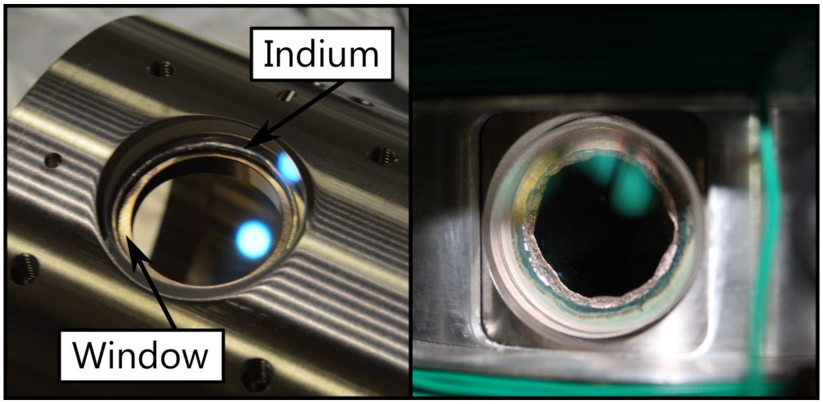

Figure 3.6: Left: The uncompressed indium ring under the window with a clear gap around its edge. Right: The compressed indium seal post bake, under compression it has spread into the chamber opening reducing the clear aperture.

the chamber once the compression flange is removed.

The compression flanges for this system attached to the threaded holes around the window recesses of the science chamber. The process for creating an indium seal was as follows; the indium was cleaned with acetone and optical tissue then cut to approximately the correct length[1]. This was placed around a jig of the correct diameter and the wire was cut to length at an angle to provide a larger bonding surface, these ends were then pressed together carefully to form a uniform ring of indium[2]. This ring was then placed between the window and the chamber, as evenly about the window aperture as possible (see Fig. 3.6), a similar ring of Teflon was placed above the window to protect it and even out the pressure on the indium.

For the side windows these flanges were in two parts as the curved surface only provided ~5 mm of thread, making precise application of pressure difficult. The first part attached directly to the chamber and the other screwed into this whilst applying pressure to the window (see Fig. 3.7a). To ensure even distribution of pressure the side windows had six securing threads present for their compression flange and the more problematic windows had eight. Initially each screw was tightened by hand until they came into contact with the top of the compression flange. Using a torque wrench each screw was then tightened in order (see Fig. 3.7b) until it reached 0.5 Nm of torque, this order was then repeated each time

---

[1] All tools went through the same cleaning process as the chamber.

[2] A rule of thumb is that it should not be possible to tell where the bond was made

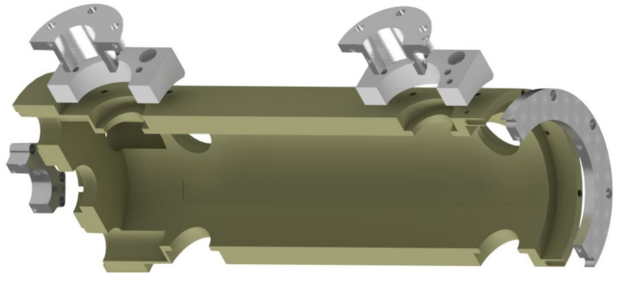

(a) The compression flanges for the different window types.

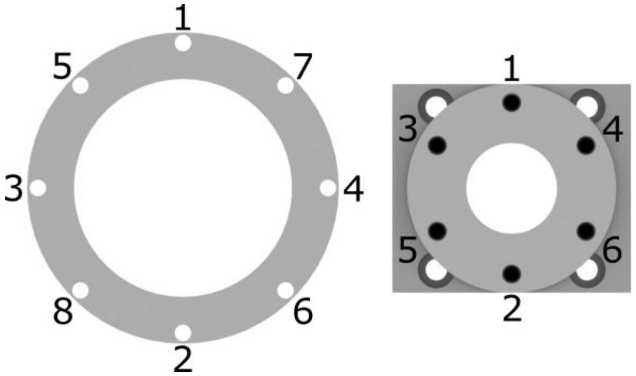

(b) The order in which screws need to be tightened to evenly compress indium under a window.

Figure 3.7: Design and use of compression flanges for indium sealing windows to the chamber.

upping the torque applied by 0.2 Nm. Once at 1.5 Nm the final step was repeated, i.e. every screw was tightened twice at the 1.5 Nm value.

Once compressed the indium should present a smooth and shiny surface (Fig. 3.6) with no distinctive features. If it is possible to identify a seam or join in the layer then it will be present a leak risk. To the side of each window recess there is a 1 mm diameter hole that is parallel with the surface of the recession; in the eventuality that the window seal leaked these would be used to apply sealant directly to the indium rather than needing to flood the area, this functionality was never needed. Indium sealing windows places a limit on the maximum baking temperature of the system, the melting point of indium is 156 °C[73] however local history with indium sealing suggests that heating much above 120 °C can create problems. This lead to the baking limit on the system of ~110 °C, far from either of these limits.

3.1.3 Cleaning and Bakeout Procedure

Preparing an ultra-high vacuum (UHV) system requires a high degree of cleaning to remove material that will introduce a substantial gas load to the system. Once cleaned the system is heated whilst attached to an external roughing pump (known as baking) to remove any residual gas from the system; with sufficient cleaning the baking procedure will mostly remove water. The components manufactured by the in-house workshop (science chambers and prism mounts) required intense cleaning. All those purchased from external companies arrived clean and sealed so were not reprocessed. The first step of cleaning uses paper towels to remove substantial grease deposits from the surface of the material before, a series of solvents and an ultrasonic bath[1] were then used to clean all remaining residue.

The cleaning procedure for the gradiometer system is the same as that for the prism MOT and gravimeter prototypes. An initial 1 hour cleaning cycle with the ultrasonic bath and 1% solution of detergent[2] and tap water was used to remove any remaining machine grease. These components were then rinsed in tap water before a 1 hour ultrasonic cycle in distilled water. The same 1 hour cycle was then performed for an isopropanol bath. The isopropanol was then switched to a clean batch and an ultrasonic cycle was performed for 30 minutes. After this cleaning procedure the components were left to dry before being placed in a 200 °C vacuum oven and left to prebake for 55 hours. This pre-bake shortens the bake required once the system is fully assembled due to the temperature limits imposed by the indium.

For the gravimeter this process resulted in a pressure as measured by the ion pumps of $\sim 10^{-10}$ mbar after a 2 week bake at ∼i110 °. The discrepancy between this value and that given in Eq. 2.14 is expected as the outgassing rates of materials vary substantially with the exact bakeout procedure and the values used for Eq. 2.14 were approximate based on the literature. The gradiometer followed the same bakeout procedure and achieved similar levels in pressure without dispenser load; with dispenser load the pressure as given by the ion pumps $10^{-10} - 10^{-7}$ mbar depending on operating current. The high end of this range was only reached as a dispenser began to run out and was heated in excess of 700 °, the typical operating pressure is at the low $\sim 2 \times 10^{-9}$ mbar (corresponding to a mean collision time of 1.8 s).

3.2 Magnetic Field Control

Controlled manipulation of the magnetic substructure of ^{87}Rb is necessary to create a cold atom gravimeter. The use of a gradient field creates the position-dependent force in the MOT (2.1) and a uniform bias

---

[1] *Ultrawave U500H*

[2] *Liquinox*

field provides a well defined direction for the Raman interactions. To provide this controlled environment is a twofold process; the system initially requires isolating from external fields and then the desired field is generated with a set of solenoids.

The magnetic isolation is provided by two sets of double layered μ-metal shields. These shields are open ended cylinders, taking advantage of the cylindrical form factor and so will provide no shielding along the shield's central axis but substantial shielding radially. The attenuation quoted for the gravimeter shields[1] is $> 30,000$ in the centre of all four layers of shielding. This will provide substantial shielding in the interferometry region and has been measured to reduce the Earth's magnetic field to below 1 μT in the MOT region. This generates a Zeeman frequency shift of $<$10 kHz, three orders of magnitude less than the 13.2 MHz detuning of the cooling laser and 21.9 MHz/cm from the MOT coils.

3.2.1 MOT Coils

Having sufficiently attenuated the background Earth field it was necessary to generate the trapping field for the MOT. As shown by Eq. 2.5 and Fig. 2.1 the field required to generate a MOT has a symmetric field gradient of opposite sign about a central point of zero field. To generate this field profile a separate mount was made to allow for the coils to be easily slid on and off the system. The coils for the prism MOT prototype generated a field gradient of ~3.1 G/cm/A, but were found to generate a substantial amount of heat despite the good air flow around them. To reduce the power consumption and heat generation of the coils time was spent modelling the field profile for different coil parameters (separation, number of turns, wire thickness, coil diameter) within the geometric confines of the existing design.

The result was a pair of coils of inner diameter 90 mm, outer diameter of 106 mm and 75 turns per coil. This produced a field gradient of ~3.9 G/cm/A whilst almost halving the power consumption. The model, generated numerically using the Biot-Savart law, is presented in 3.8 beside measured values for the MOT field and the calculated field gradient, showing a good agreement between the two.

---

[1] With an almost identical geometry.

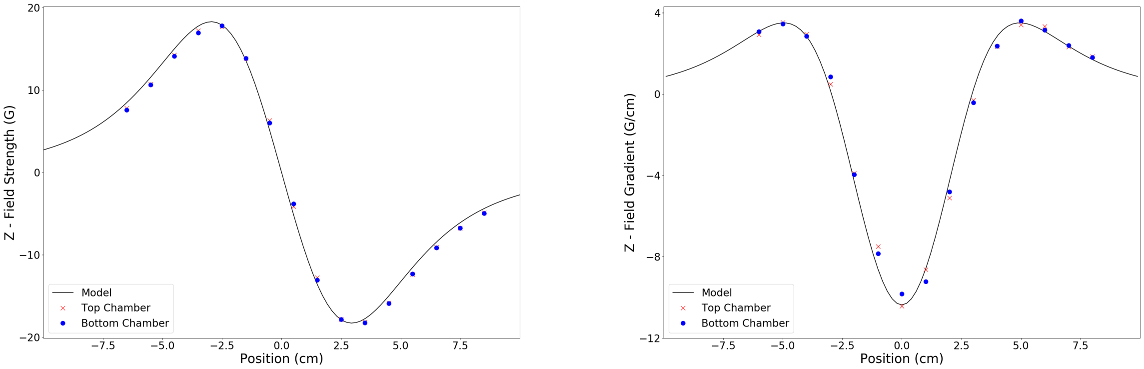

(a)

(b)

Figure 3.8: MOT magnetic field (3.8a) and field gradient (3.8b) at 3 A operating current.

3.2.2 Bias Coils

The interferometry region requires a uniform and well defined field direction. A high level of radial field attenuation is provided by the four layers of shielding in the interferometry region, meaning that any non-negligible field contributions will be axial only. To generate a well-defined quantisation axis and break the degeneracy between the $\Delta m_f = 0$ transitions (see 5.4.1) a single layered, 18 turn, solenoid was placed inside all layers of shielding. The goal of this design was to produce a compact coil that reached the desired field strength in as short a region as possible and produced a good uniformity over the interferometry region. The result utilised a 3D printed mount with grooves into which 0.9 mm diameter wire was set and produced a field uniformity of $< 0.2\%$ over the central shielded region (Fig. 3.9).

Sensitive cold atom gravimeters remove first-order magnetic shifts by preparing in the magnetically insensitive $m_f = 0$ sub-state of ^{87}Rb hyperfine ground state. Second order effects are more persistent and apply a shift to the measured gravitational acceleration that cannot be separated from the gravitational measurement[75]. These second-order effects lead to a stringent requirement on the uniformity of the magnetic field profile over the interferometry region. A gradiometer system such as this does not suffer from these effects if they are temporally stable; these higher order effects will induce this same shift to the value of gravity measured by each cloud[1] that will result in an offset in the gradient measurement. If measurements relative to a reference are taken, as in a survey, these offsets will subtract out of the final

---

[1] It is not required that this shift be the same for both.

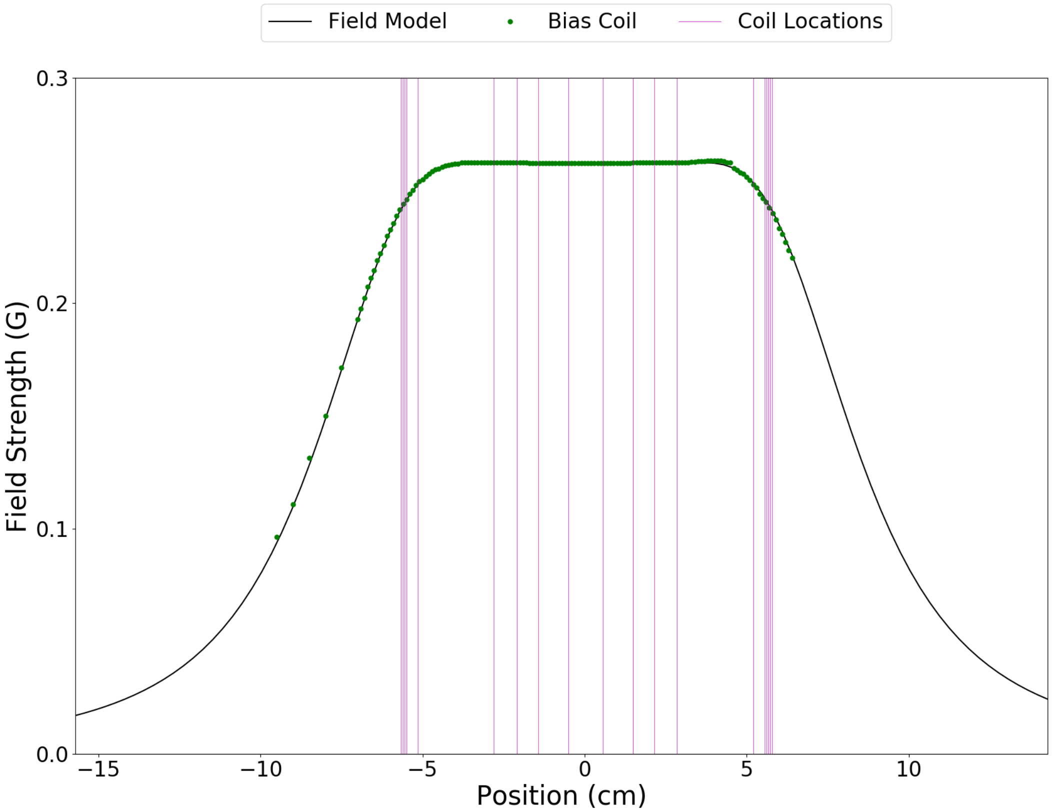

Figure 3.9: Measured and simulated bias field profile of the gravimeter coil at a current of 0.2 A.

result, meaning that the field uniformity requirement is suppressed for a survey-ready gradiometer.

This suppressed magnetic gradient impact is useful given the profile of the bias coil in Fig. 3.9. A high concentration of turns is present at the edge of the coil allowing the field to reach its maximum value within 1 cm of the coil edge, a gradient this sharp would be problematic if the effect were unstable. This rapid gradient is balanced by a reduced concentration of turns near toward the inner edge of the coil resulting in a uniform field. The data presented here was measured on the gravimeter prototype coils and using a current of 200 mA, this will be representative of the gradiometer coils with an identical design. The gradiometer field profiles have been measured using atomic spectroscopy, the results will be outlined later (5.4.1).

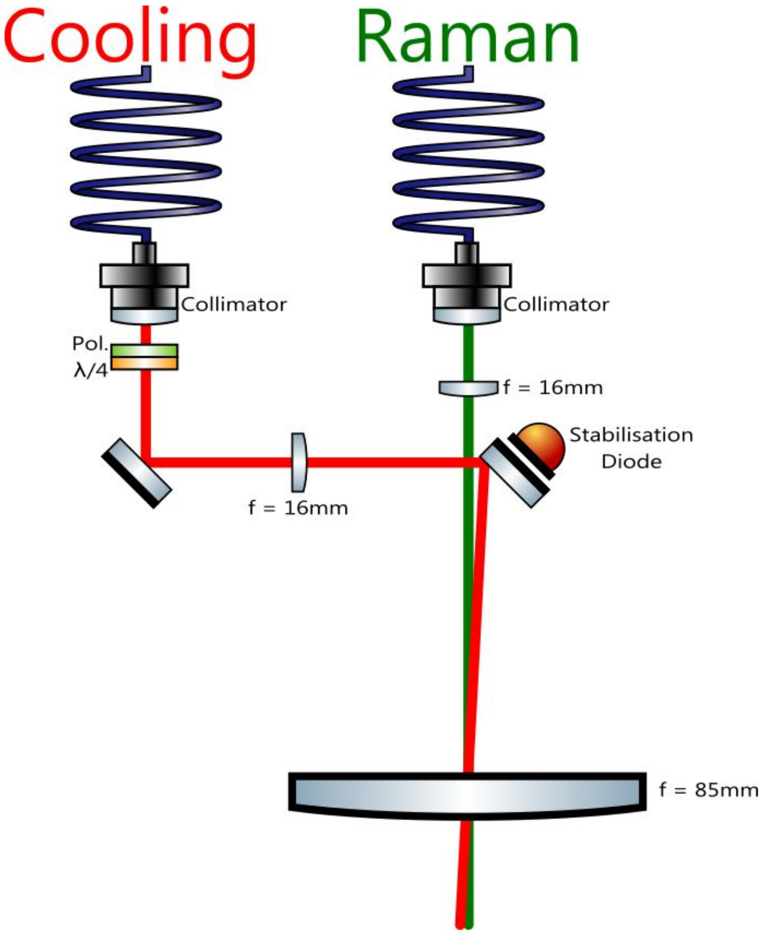

Figure 3.10: Diagram of the main experiment telescope. Both cooling and Raman beams are delivered to the telescope from the same direction and reflected to ensure that delivery into the chamber is near coaxial.

3.3 Optical Delivery

The optical delivery system is the sub-system responsible for delivering the light from the control package to the sensor and the atomic readout signals from the sensor to the flight case. The main delivery telescope had to fit within the system's spatial constraints and maintain the cylindrical form factor. Additionally the limited power of the Raman laser meant that the cooling and interferometry beams could not be the same diameter, requiring a more complicated optical system. These issues were solved with a telescope that can be seen in Fig. 3.2, with a schematic in Fig. 3.10.

Polarising fibres[1] deliver the light directly into one of two collimators[2], with output beam sizes of ~19 mm and ~7 mm for the cooling and Raman beams respectively. The cooling arm features a polariser (an artefact from before the introduction of polarising fibres) and a quarter waveplate. The light from this collimator is reflected off a 0.5 " mirror on a kinematic mount, providing the only alignment capability in the system. This gives the facility to adjust the alignment of the cooling light to place the focus of the beam near the edge of the D-mirror before it is collimated by a 3 " lens producing an collimated output beam with a $1/e^2$ diameter of ~53 mm.

The Raman beam focuses to a point near the edge of the D-mirror placing the foci of the two beams as close together as possible. After passing through the same 3 " collimating lens the Raman beam has a $1/e^2$ diameter of ~20 mm. This results in some small offset angle between the cooling and Raman beams, with the Raman beam mechanically locked coaxial to the science chamber. The stabilisation diode is mounted behind the second-back polished mirror to pick up the light that leaks through and provide a signal to the intensity stabilisation circuit (see 4.3) which is used to remove the intensity drift arising from polarisation fluctuations.

The full optical delivery system is shown in Fig. 3.11 where the prism mount is used to represent the location of the MOT. Considering the system from the top down; the telescope is as previously described and level with the MOT are two pieces of imaging equipment used to monitor the pre-drop cloud. One is a small, untriggerable CCD camera to monitor the MOT visually on a screen, the other is a photodiode and lens system used to measure loading curves. Below this is the detection region with a second diode used for detecting the fluorescence of the cloud during the detection sequence when illuminated with the main MOT beam, the signal is then amplified and transmitted back to the flight case. A 3D printed periscope and opposing mirror generate a balanced light-sheet for time of flight measurements, as discussed in 2.3.2. More details of the detection procedure will be given in Ch. 5.

Interface Box

The sensor and control package are connected with a 5 m umbilical, at the base of the sensor is a box which provides an interface with the electrical connections and two fibre inputs, this box then distributes

---

[1] *IX Blue IXS-POL-780-125-10-Box-FA-NA*

[2] Cooling: *Thorlabs F810APC-780*, Raman: *Thorlabs TC12APC-780*

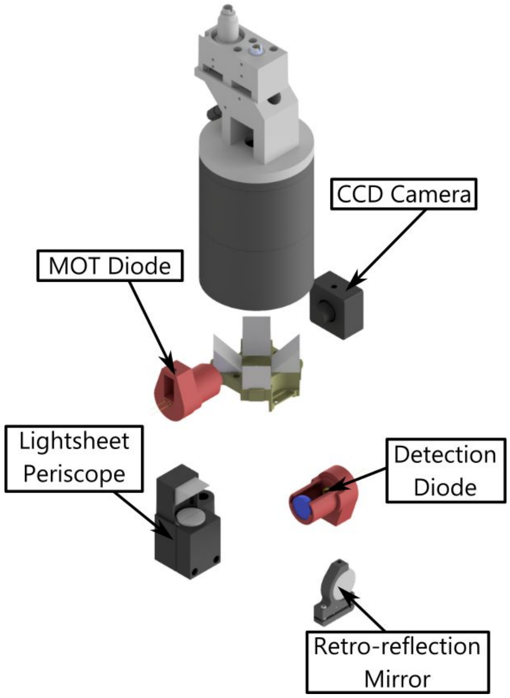

Figure 3.11: The chamber based optical delivery for the Top chamber

the light to the appropriate part of the system. Figure 3.12 shows the details of this distribution - the polarising fibres on the input are to ensure that any polarisation noise induced over the umbilical is removed before distribution. The optical switch delivers the cooling light to either the telescopes or the light-sheet periscope. The light-sheet output also acts as a light dump for the cooling light during the drop due to light leakage through the AOM. The switch induces a ~ 50% loss in cooling power so it would be favourable to remove if the light-sheet and dump facilities were not required.

The monitor outputs were initially used as pickup points for the stabilisation circuits, however it was not sufficiently representative of the light into the chamber so was switched for the telescope based diodes. The Raman monitor is now delivered to a cavity providing information on the sideband ratio (see Ch. 4). The secondary set of polarising fibres ensure that additional polarisation fluctuations are removed before the light enters the chamber, these fluctuations are then projected as intensity nosie and can be removed by the stabilisation system.

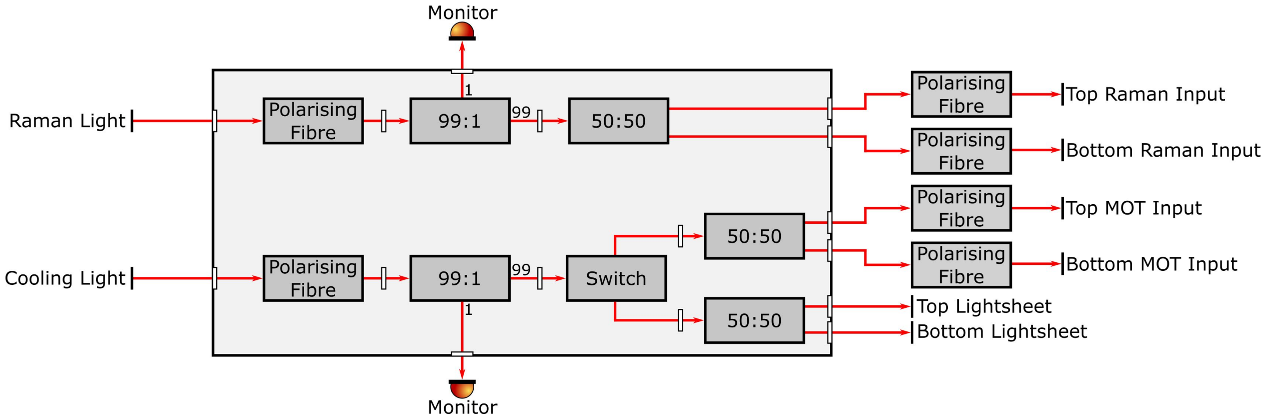

Figure 3.12: The fibre distribution system for the experiment light. The whole system, along with electrical distribution, fits into the base of the system with a volume of ~8 dm^3.

3.4 Integration

To couple together the two sensor heads via an external mechanical mounting structure and common light sources an adaptable mounting structure was developed. The integrated sensor assembly mounts the sensors facing each other, separated by their baseline, as in Fig. 3.1 on a rigid mechanical structure. This structure is custom made forty piece assembly (Fig. 3.13) and provides much of the “Adaptability" as in Table 3.1. Despite the forty individual parts there are only eight unique designs; enabling easier mass manufacture should this sensor be adapted into a product.

The sensor consists of four separate modules as seen in Fig. 3.13 with each module labelled from [1] to [4]. [1] and [2] are the Bottom and Top sensor heads respectively; the vacuum components and coils near the top of [1] and the detection diode visible near the bottom shows this is the Bottom sensor as the MOT is formed at one end and will fall towards the telescope. The two sensors are joined with module [3] on which there is a central cross piece into which the retro-reflection mirror is mounted. This cross piece can be removed without disassembling the rest of the system to allow the chambers to be directly coupled together with the light. Module [3] also allows for a change to the system baseline by exchanging the uprights for those which are longer or shorter, changing the length of this module only and the separation of the two chambers. Module [4] is the system base with similar uprights to the rest of the support structure, the length of these can be easily changed to alter the total height. The base of [4]

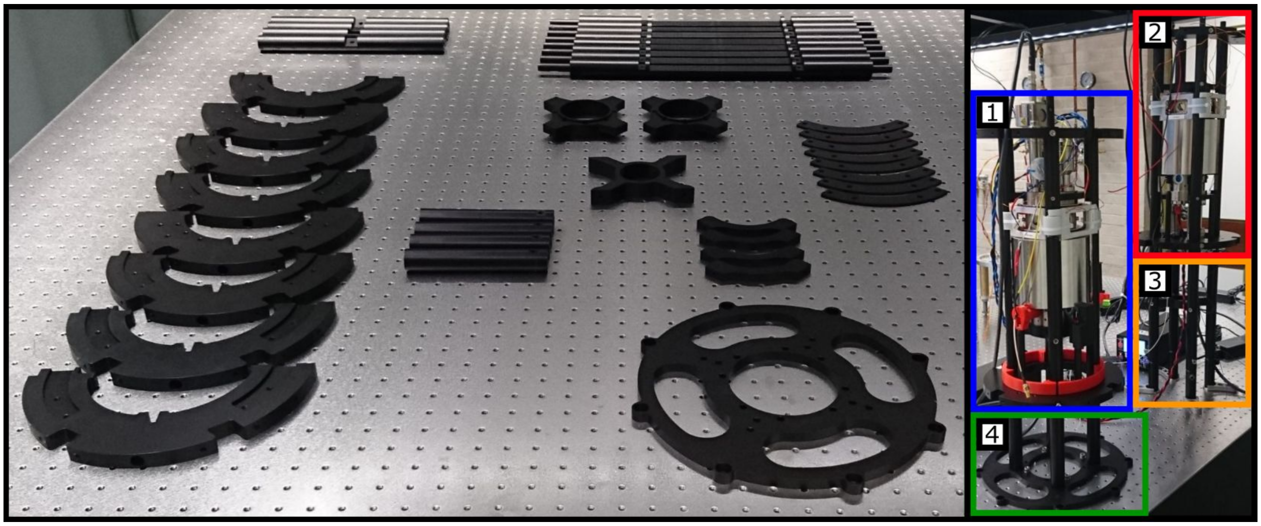

Figure 3.13: Left: The forty piece mounting structure before assembly. Right: The four modules the sensor package can be easily broken down into after assembly.

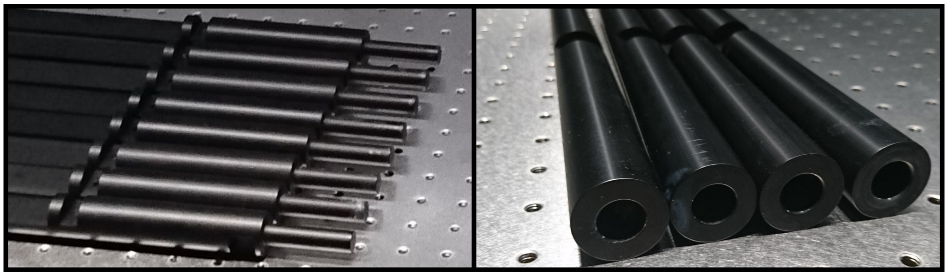

Figure 3.14: Pole and pockets pairings of the mounting structure, the end of each set of uprights features one of these features allowing for easy modification of the system modules.

connects to a separate tripod base with fine machine threads on its legs, this allows precise adjustment of the system tilt. The interface box (Fig. 3.12) is built into this tripod base.

To connect the pillars end-to-end they feature a set of 0.5 " protrusions and pockets as in Fig. 3.14. The 0.5 " diameter allows direct interfacing with standard commercially available optomechanics, a particularly useful feature during assembly allowing the sensors to be secured to the table before they were mounted onto the rest of the system (as can be seen in module [3] in Fig. 3.13). These pockets are present on all the pillars of modules [3] and [4], whereas the protrusions are on the edges of the pillars in modules [1] and [2], this produces complete symmetry in the assembly for modules [1], [2] and [3] about the central mirror mount in module [3]. This means that to switch from the drop configuration to a launch configuration the sensor just needs to be lifted out of module [4], inverted and then slid back into module [4]. A functionality that has never been utilised as the laser scheme needed for the launch is still under development. The final assembly introduces a set of plastic covers (covered internally with black tape) to block out background light and make the system splash proof. Standing at roughly 175 cm tall, with an outer diameter of ~30 cm and weighing just under 50 kg the system is both smaller and lighter than most members of the team.

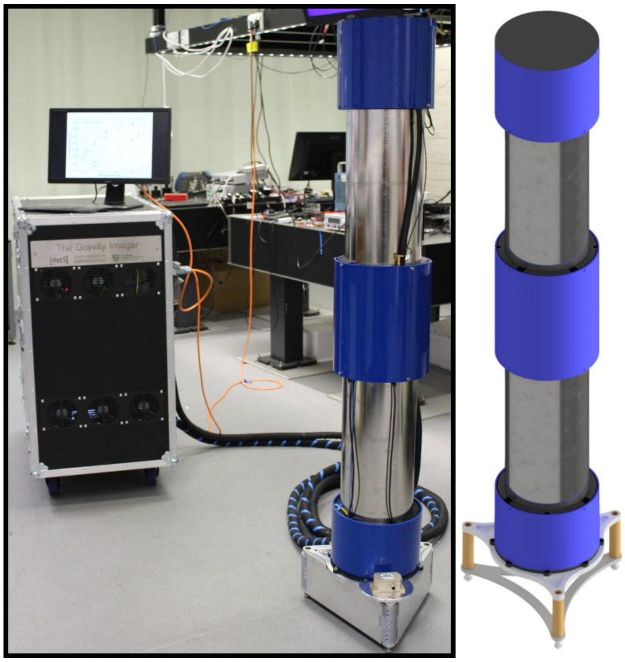

Figure 3.15: The detailed CAD model of the system presents a fairly accurate representation of the final product, as can be seen here.

CHAPTER 4

CONTROL PACKAGE

There are two sides to creating a portable atom interferometer. The sensor package as described in the previous chapter is responsible for isolating the atoms from external influences whilst allowing their controlled manipulation with light. The other side is dedicated to the control of the experiment; generation of all light responsible for interacting with the atoms and the core control of the system, timing and sequence structure. In this system these two sides have a distinct physical split; the sensor head as described in the previous chapter and the control package which is contained within a 90 cm deep 20 U high 19 " rack unit mounted into a portable flight case. With a 5 m long umbilical providing two way communication for light, power and detection signals between them. This umbilical allows for the sensor to be freely moved in the field, with small scale surveys to be performed by moving the sensor head only and leaving the flight case in a fixed location. Figure 4.1 shows a top level breakdown of the distribution of subsystems within these two packages.

The choice was made to construct all control elements into 19 " rack units as they provide a standardised interface in a compact form factor. These are mounted within a flight case that has sufficient room for 20 U[2] on a single side and is deep enough to fit two full sets of 40 cm deep 19 " boxes back-to-back. Mounting the control into this rugged, portable platform allows it to be protected and easily moved in the field.

Much of the space in the flight case is taken up by power electronics: five 2 U programmable benchtop power supplies and a single uninterruptable power supply (UPS) occupying the full bottom 2 U of the flight case. A new custom-made power supply is being characterised that merges the 10 U of vari-

---

[2] The height of any 19 " rack box is given in terms of U, where 1 U≈4.45 cm.

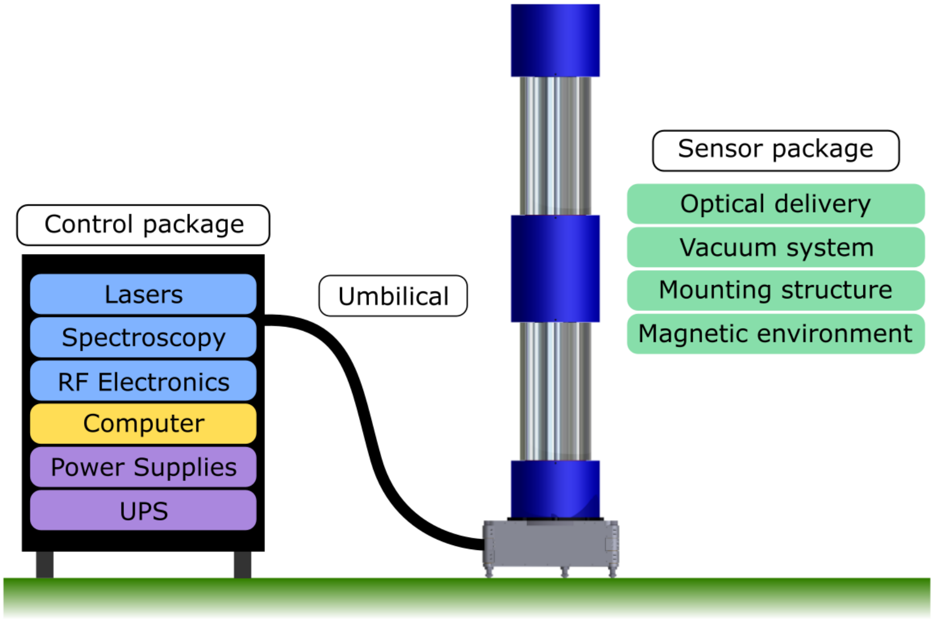

Figure 4.1: To use a biological analogy, the control package is the “brain" and the sensor package is the “body" with the umbilical acting as the nervous system.

able supplies and 1 U of IGBT control boards into a single 1 U box, freeing up a quarter of the flight case volume. Of the remaining space 10 U is dedicated to the home-built lasers plus their associated locking systems; one for the cooling light and one for the Raman, more compact lock boxes and careful design of laser components provides potential for this to shrink. The rest of the space is occupied by control electronics, constant voltage supplies, RF electronics, intensity control and the core computer control.

The purpose of the flight case is to generate and control the light required to interrogate the transitions of ^{87}Rb as shown in Fig. 2.11. There are three transitions of interest: the main $|F = 2\rangle \longrightarrow |F^{'} = 3\rangle$ cooling transition, the $|F = 1\rangle \longrightarrow |F^{'} = 2\rangle$ repump transition and the two photon Raman transition on $|F = 1\rangle \longrightarrow |F = 2\rangle$ via the intermediate detuned state. The cooling laser is the common light source for the L_1 and L_2 beams as in Fig. 2.1. It is necessary to provide repump light as it is possible for the cooling light to excite an atom into the $|F = 2\rangle \longrightarrow |F^{'} = 2\rangle$ transition which then has a chance to decay into the $|F = 1\rangle$ state, removing it from the cooling cycle. Repump light is used to keep the atoms out of this dark state. To do this the lasers need to be frequency stabilised to an atomic transition in a way that allows for a few 100 MHz of control in order to produce a molasses sequence or frequency chirp (see 5.1). The system

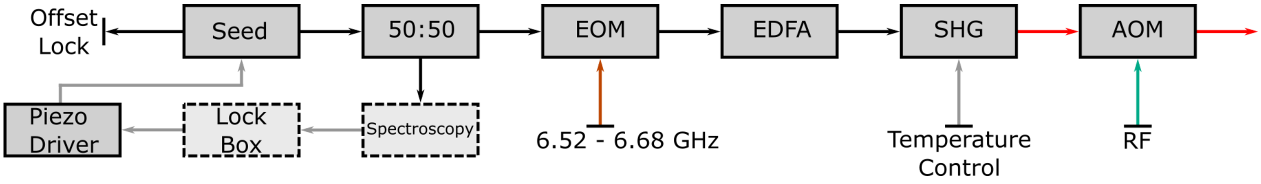

Figure 4.2: Schematic breakdown of the cooling laser, this is all contained within a single 2 U rack unit with the exception of the spectroscopy and lock box.

outlined in this chapter is responsible for generating all the frequencies and timing needed to control the atoms. For a complete overview diagram see **Appendix A**.

4.1 Cooling Laser

A fully fibre-based laser system generates both the cooling and repump light integrated into a single 2 U rack mountable box, the output light is fully stabilised and controllable. The frequency is locked with modulation transfer spectroscopy (MTS) to a transition in ^{85}Rb with a DDS providing frequency control over a ~200 MHz range. The intensity is controlled with an AOM fed with an RF signal with a variable amplitude gain that is set to maintain a constant intensity into the chamber (see Fig. 3.10 and 4.3).

The seed lasers used in light generation are a telecoms standard[1] that generate light at 1560 nm, this is frequency doubled to 780 nm, making it resonant with the ^{87}Rb D2 transition. This allows us to take advantage of these compact, low linewidth, fibre coupled lasers (150 kHz in use[33]) with a low frequency noise over a large operational temperature range (15 °C - 55 °C) [76]. Complete fibre delivery was used as a free space system would be unstable during transport requiring realignment at each new survey site.

Figure 4.2 shows the implementation used in this system generating ~40 mW of 1560 nm light from the main output and a secondary monitor ~1 mW that provides the input for the Raman offset lock. The main output is split 50 : 50 between the output of the light chain and the spectroscopy, used to lock the cooling light to the $|F=3\rangle \rightarrow |F^{'}=4\rangle$ transition of ^{85}Rb (see 4.1).

The electro-optical modulator (EOM) produces a sideband on the cooling light with a controllable modulation frequency generated by a VCO[2] with a frequency range of 6.52-6.68 GHz[3]. Control of the

---

[1] *NKT Photonics : Koheras Basik E15*

[2] *Mini-Circuits ZX95-6640C+*

[3] Measured to have a higher maximum frequency than the quoted 6.64 GHz

modulation frequency allows the sideband to be kept resonant with the $|F=1\rangle \rightarrow |F^{'}=2\rangle$ repump transition whilst the carrier frequency is scanned. With only a single seed laser used for both cooling and repump light there is a substantial reduction in system complexity, but does not allow for generation of light resonant with the repump transition only. This means there is no facility to produce a blow-away pulse on the $|F=1\rangle$ state and complicates two-state detection. An approach under investigation is to overdrive the EOM with a higher-power VCO in order to fully suppress the carrier.

The modulated laser frequency is then amplified with a 2 W erbium doped fibre amplifier[1] (EDFA). This is frequency doubled with a second harmonic generator[2] (SHG) that has an efficiency of 50%. An AOM (transmission efficiency of ~ 50%) on the output provides variable control of the light intensity (see 4.3) in addition to providing fast switching. Due to transmission of the 0^{th} order when the AOM is “off" ~5 μW of light is leaked requiring the optical switch in Fig. 3.12 to ensure that the atoms do not interact with this near-resonant light during the drop.

Spectroscopy

The spectroscopy box is housed in a separate 2 U rack unit and enables the frequency stabilisation of the cooling light to the $|F=3\rangle \longrightarrow |F^{'}=4\rangle$ transition of ^{85}Rb. The MTS module is the only part of the control system that features free space optics, constructed using miniaturised optical mounts (as featured in [77]) which were developed as part of the QUANTUS collaboration[78][79]. These have been proven to withstand repeated 50 g decelerations[80]. Figure 4.3 shows the layout of the spectroscopy box, the input light is modulated at ~1.13 GHz using an EOM. This is the frequency difference between the $|F=2\rangle \longrightarrow |F^{'}=3\rangle$ transition of ^{87}Rb and the $|F=3\rangle \longrightarrow |F^{'}=4\rangle$ transition of ^{85}Rb allowing the probing of the ^{85}Rb atoms in the rubidium cell.

The EDFA amplifies the input light, dumping 90% of this light, this is because if the EDFA does not operate at a sufficiently high amplification it enters an output mode with discrete jumps in intensity to give an output that averages the right power level. MTS produces a sharp, Doppler free error signal with a background that is independent of fluctuations[33][81] making it perfect for a system that will be placed in unstable environments. To provide frequency control of the cooling laser the driving frequency of the

---

[1] *Koheras Boostick OEM*

[2] *WH-0780-000-F-B-C NTT Electronics*

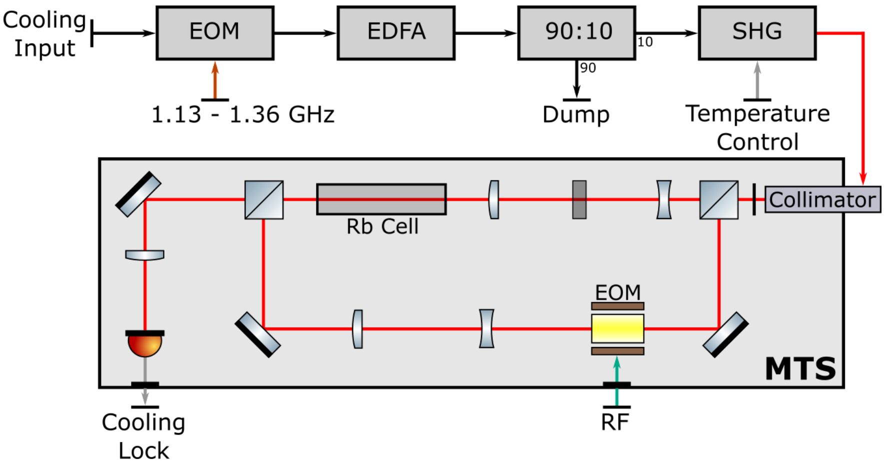

Figure 4.3: Schematic of the spectroscopy box, contained in a single 2 U rack box with the MTS module contained within its own smaller enclosure.

first EOM in the chain is a mixture of a constant 1.057 GHz and a variable DDS[1] frequency. This sets the position of the locked side-band relative to the experimental output of the cooling laser, shifting this modulation frequency thus moves the carrier frequency equivalently.

4.2 Raman Laser

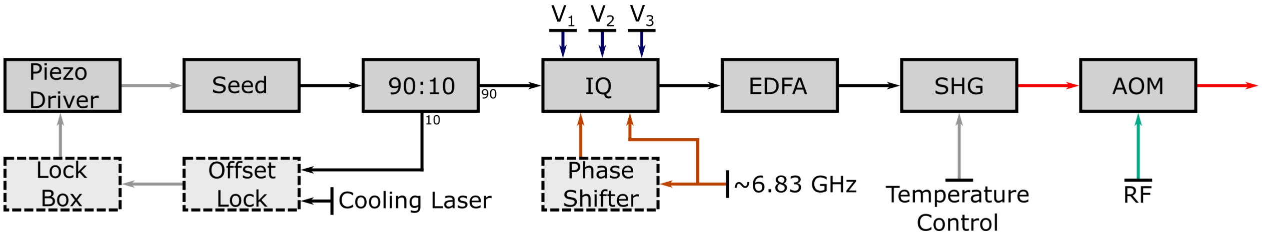

Figure 4.4: Raman laser for the system, it is a modification of the same fibre laser used for the cooling light.

The Raman laser is used to generate the light for the interferometry pulses; requiring two frequencies with a controllable frequency separation equivalent to the hyperfine ground state splitting of ^{87}Rb (~6.834 GHz). The implementation is very similar to that of the cooling laser with some modifications; rather than locking to an atomic transition the Raman is slaved to the cooling laser with an offset lock. In

---

[1] *Novatech 409B*

place of the EOM used to generate the repump sideband an IQ modulator[1] generates a single frequency sideband. For a detailed treatment of the operation of an optical IQ modulator see [82] and [32].

The offset lock generates a beatnote between the Raman and cooling lasers that is filtered and then sent to a comparator board. On this board the input frequency is split between two arms; one with a high-pass filter (HPF) that removes frequencies below a cutoff (f_{HPF}) and then attenuates by -3 dB, in the other there is no filter and a -6 dB attentuation with a reverse bias. Recombining these two signals results in a sharp zero crossing when the input beatnote is greater than f_{HPF}, as the arm with lower attentuation is no longer filtered. The board outputs a voltage proportional to the sum of the two contributions, therefore as the slave frequency is scanned there is a sharp zero crossing from negative to positive voltage at f_{HPF}[83].

In the offset lock the beatnote is mixed with an additional 390 MHz, shifting the effective frequency at which this zero crossing occurs and producing a larger detuning than possible with the comparator board alone. The resulting detuning from the $|F=2\rangle \rightarrow |F^{'}=3\rangle$ transition is ~1.1 GHz, a larger detuning would be desirable but the low Raman power would cause this to have a detrimental effect. This system can be implemented with either an IQ modulator or an EOM, however the unwanted 1^{st} order sideband generated by the EOM introduces a spatial variation in the Rabi frequency over the course of the drop (Eq. 5.12). This spatial variation leads to decreased contrast on the Rabi oscillations, removing it reduces the noise on the ellipse (see 5.4).

4.3 Intensity Stabilisation

Using fibre only delivery removes many of the risks associated with alignment drift during transport that would require readjustment upon arrival at a site. Additionally it removes the need for ultra-clean conditions in the optical distribution as free space systems require consistently high surface quality over the lifetime of the experiment. Fibre delivery comes with a major drawback; applying variable stress, either mechanical or thermal, will lead to fluctuations in the polarisation of the output light[26][84]. The 5 m umbilical fibre means both of these stress sources will be present; to remove the impact of these variations a polarising fibre is placed on all telescope inputs, projecting polarisation noise onto intensity. Variation in intensity is easier to correct for with the introduction of a stabilisation system utilising the

---

[1] *iXBlue MXIQ-LN-40*

AOM to intensity control in the lasers. The control is provided with the circuit in Fig. 4.5 and the light intensity level is monitored using the diode behind the back-polished mirror in the optical delivery telescope (see Fig. 3.10).

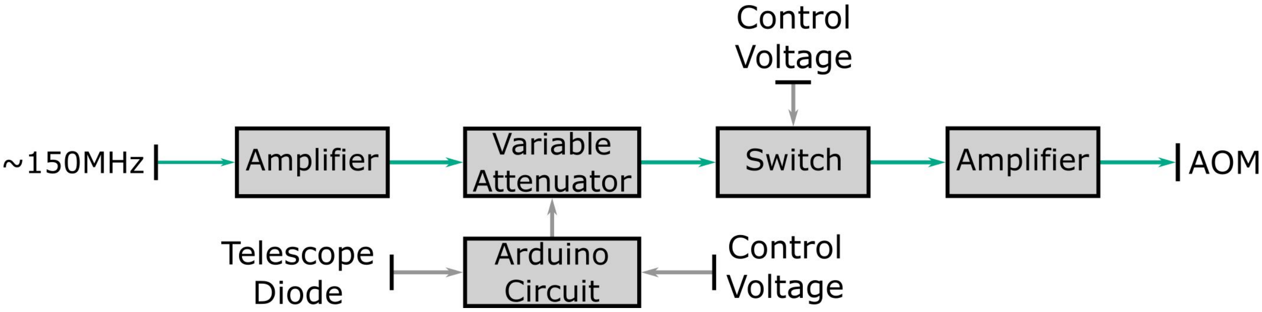

Figure 4.5: Applying a variable gain to an input diode signal allows control of the light level as measured at the telescope whilst also facilitating stabilisation.

Stabilisation is implemented with a variable attenuator on the 150 MHz driving frequency of the AOM, so the power of the driving frequency can be adjusted to maintain a constant light level. All of the optical processes involved in atom interferometry are intensity dependent, however some require variable intensity (see 5.1) therefore stabilising to a rigid set point is not possible. To stabilise but maintain the required control the light level set in the sequence changes an output voltage, the Arduino receives this voltage and uses it to set a variable gain, this feeds a variable RF attenuator which in turn controls the output of the AOM. The power of the AOM driving frequency is then proportional to the desired output level, maintaining a stable light level whilst providing absolute intensity control.

Both Top and Bottom telescopes feature a stabilisation diode as seen in Fig. 3.10, allowing a choice of stabilising to either the Top or Bottom telescope in the current implementation. Should this not provide sufficient stability there is the possibility of introducing an additional AOM to give independent control of both experimental arms. Figure 4.6a shows the level of variation seen in the MOT light over the course of a little over an hour. Using the model presented in 2.4.1 this ~ 8% variation in intensity leads to a ~ 9% variation in effective capture volume and as such should result in a similar level of atom number fluctuation.

Looking at the difference between the two chambers allows the effectiveness of this approach to be quantified. A ~ 8% overall variation is reduced to a ~ 2% variation in the difference between the arms,

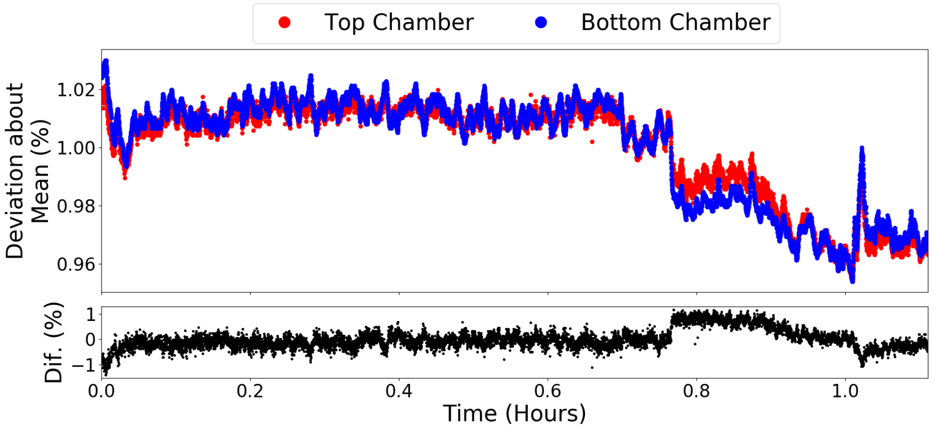

(a) Free running MOT light.

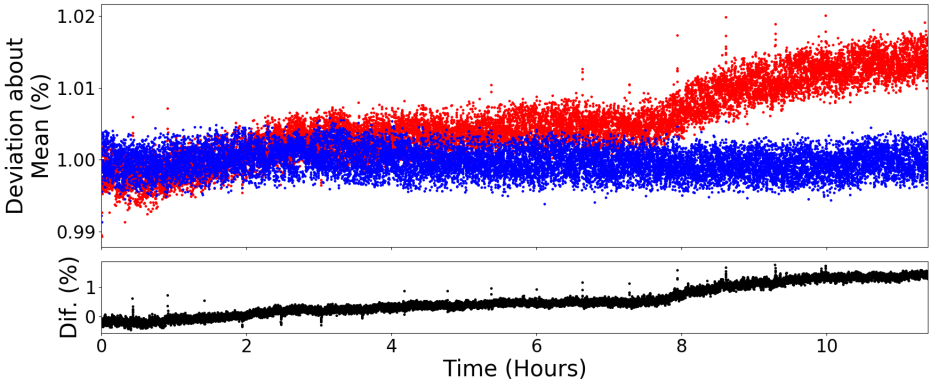

(b) MOT light stabilised to the Bottom telescope input.

Figure 4.6: MOT light level variation for both the case of stabilised and free running. Each of the datasets has been normalised separately to its mean in order to provide a better comparison. The few rogue points visible near the end of the stabilised data set arise due to the duty cycle of the device, a sequence is regularly resubmitted leading to a downtime and thermal changes to the AOMs in the lasers.

meaning the majority of the variation in the light is therefore common to the two chambers and must arise from drifts before the arms are split. The remaining 2% will arise from the separate delivery and will not be possible to remove in the current implementation of the system. Figure 4.6b reinforces this conclusion, with the light stabilised to the Bottom chamber. The stabilisation has clearly reduced the noise on the Bottom input to below 1% and that in the Top to ~ 2% meaning that the Bottom has been stabilised at the cost of the Top. Using the model built in 2.4.1 these 1% and 2% variations in light level become approximately 1.5% and 3% variations in effective capture volume, a better than twofold reduction in the noise.

Section 5.1 will give a detailed outline of the experimental sequence but it briefly needs to be considered here. The total length of a sequence is ~2.3 s consisting of a 0.5 s for Raman stabilisation, 1.5 s MOT loading and ~300 ms for the interferometry sequence and readout. The time given over to Raman intensity stabilisation is before the MOT is loaded, there is therefore a 1.5 s window during which time the Raman intensity is free running and can drift before it interacts with the atoms.

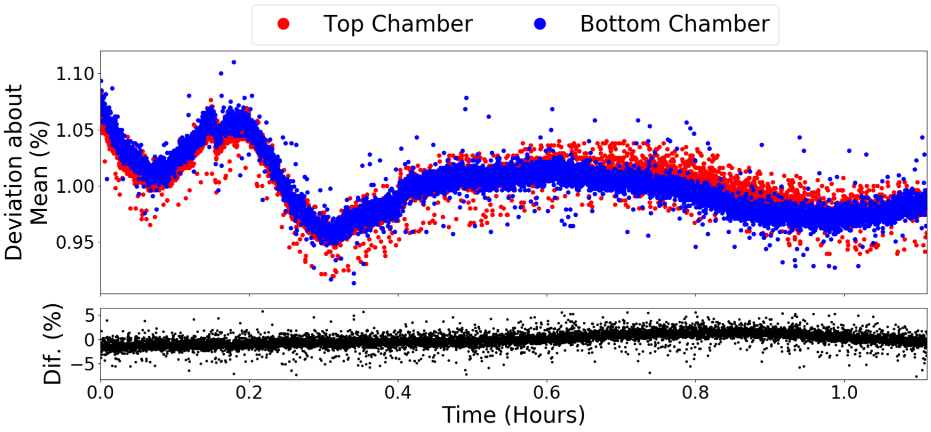

(a) Free running Raman light.

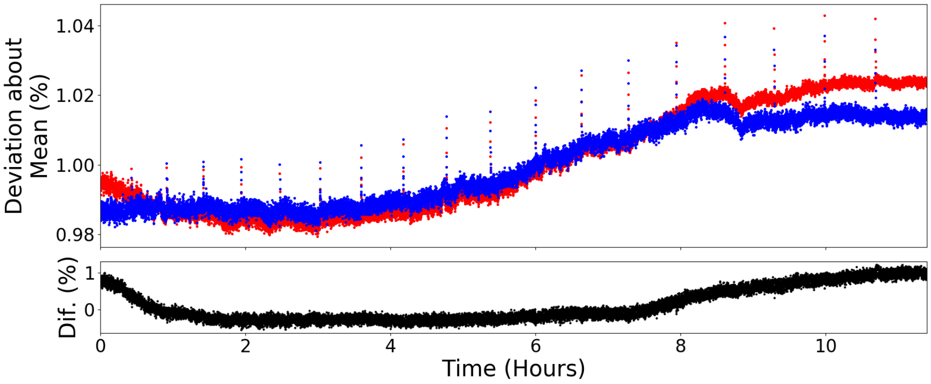

(b) The light level has been stabilised to the Bottom telescope input.

Figure 4.7: Drift of the Raman light level with time for both the stablised and free running cases. The rogue points in the stabilised data set are similar to those in Fig. 4.6b and likely arise due to the same thermal cycling of the AOMs.

Similar data for the Raman stabilisation is presented in Fig. 4.7a, the rogue points around the edge of the data set appear to occur at the start each new new batch of experimental runs, it is unclear precisely why this happens. The results seem to support the data seen in the MOT stabilisation with regards to the ~ 2% uncommon noise between the outputs suggesting that this is the fundamental limit to this style of the intensity stabilisation. There is a common long term drift of ~ 4% over ~12 hours despite the stabilisation, down from ~ 20% in ~1 hour for the free-running case.

4.4 Computer Control

The core of the computer control system is the *National Instruments sbRIO-9627* FPGA board, providing 100 digital outputs and 4 16-bit analogue outputs. These generate a set of 40 digital outputs via a set of 5 buffer boards. The outputs are used to control most of the experimental sequence with a timing resolution of 100 ns, by supplying voltages to VCOs, switch control and generating triggers. Rather than use the inbuilt analogue channels a choice was made to derive these from the digital outputs as they presented a timing jitter on the order of 10 μs, which was not seen on the digital channels. As the digital channels output a 0 - 5 V signal is adapted to provide the 0 - 20 V output range required to drive the VCOs.

Python is used as the experimental interface with the control system. A set of centralised classes produces a database of all control functions and a given sequence calls these into an individual script. This ensures that certain parts of the sequence, i.e. the MOT loading and optical molasses, use consistent parameters in all scripts whilst it was possible to explore other parameters without the risk of impacting any other measurement scripts. This also allowed for easy interfacing with multiple hardware components; the FPGA control, DDS communication and oscilloscope[1] readout to occur in the same base script as the sequence definition. The FPGA core is clocked to its internal 40 MHz reference; a potential future improvement is to exchange this to the 100 MHz source used to lock the DDS to provide a common experimental clock and improve the timing resolution, this is currently not seen as an issue.

---

[1] *Picoscope 5442D*

CHAPTER 5

EXPERIMENTAL PROCEDURE AND CHARACTERISATION

Characterisation of systematics is required to best optimise the experimental sequence and generate the best results for a specifics of a system, such as beam alignment and optical power. The virtue of a gradiometry survey is that offsets which are constant in time will be removed during the subtraction. Measures were introduced to mitigate the impact of time varying systematics, such as the intensity stabilisation in 4.3.

To perform a meaningful gradient measurement with a cold atom source, a set of specific actions must be performed at specific times. These procedures are common amongst atom interferometers but the exact implementation will vary from system-to-system. A sequence begins by preparing the atomic source by trapping and cooling $10^7 - 10^8$ ^{87}Rb atoms to around the Doppler limited temperature of 146 μK[65]. A process known as “optical molasses" can then be implemented on the MOT to further reduce the motion of the atoms below the Doppler limit[85], typically producing single μK temperatures as in the gravimeter measurement in Fig. 2.9. Following the preparation of the atomic source a sequence of coherent Doppler sensitive Raman pulses are used to couple together the two hyperfine states of ^{87}Rb such that the final state readout depends on the gravitationally induced Doppler shift experienced by the clouds. This chapter will present the standard experimental sequence used for the results in this thesis and some characterisation of systematics will then be outlined.

5.1 Experimental Sequence

Over the lifetime of the experiment a number of hardware changes have been made to the system, some of these have required changes to the optimum sequence; the most substantial of these was the removal of the central retro-reflection mirror changing from the implementation shown in Fig. 2.19c to that shown in Fig. 2.19d. This made all light common for the two clouds; all drifts in the Raman light would now be common mode and their effect suppressed in the ellipse. However the vertical MOT beams will be delivered by different input telescopes, in the worst case scenario this could lead to the drift seen in Fig. 2.6 returning along the gravity axis.

Over time changes to the system hardware or in understanding of the system have required changes to the experimental sequence. Some historic data (e.g. 5.3) was obtained with an implementation different to that presented here, the version shown in this section is accurate as of August 2018 and will be a good representation for most data taken as part of this thesis. All sequence based timings will be given with respect to a common $t = 0$ reference, this is taken to be the moment at which the MOT light power ramp finishes (as shown in Fig. 5.1) and the atoms can be described as in free fall.

5.1.1 Atom Preparation

As described in 4.3 the first 0.5 s of any given sequence is used to stabilise the intensity of the Raman light. To do this the Raman light is pulsed on from t =-2.024 s to t =-1.524 s during which time the intensity stabilisation is triggered. The atomic sample is then prepared by loading a MOT from background vapour from t =-1.524 s to t =-24 ms. During this time the prism region is illuminated with light that is red detuned by 13.2 MHz from the $|F = 2\rangle \rightarrow |F^{'} = 3\rangle$ and a repump sideband resonant with the $|F = 1\rangle \rightarrow |F^{'} = 2\rangle$ transition. From t =-1.524 s to t =-24 ms the MOT coils are powered with a current of 3 A providing the magnetic quadrupole field needed for the spatially dependent MOT force in three dimensions.

MOT loading in the gradiometer system results in $\sim 5 \times 10^{7}$ atoms after 1.5 s, which will have a temperature at the 100's μK level. The Doppler-limited temperature (146 μK[65]) of the MOT corresponds to a cloud RMS velocity of 20.5 cm/s, over the course of a 175 ms drop the cloud will expand to a diameter of ~7 cm. This would lead to a substantial drop off in signal size and contrast with many of the

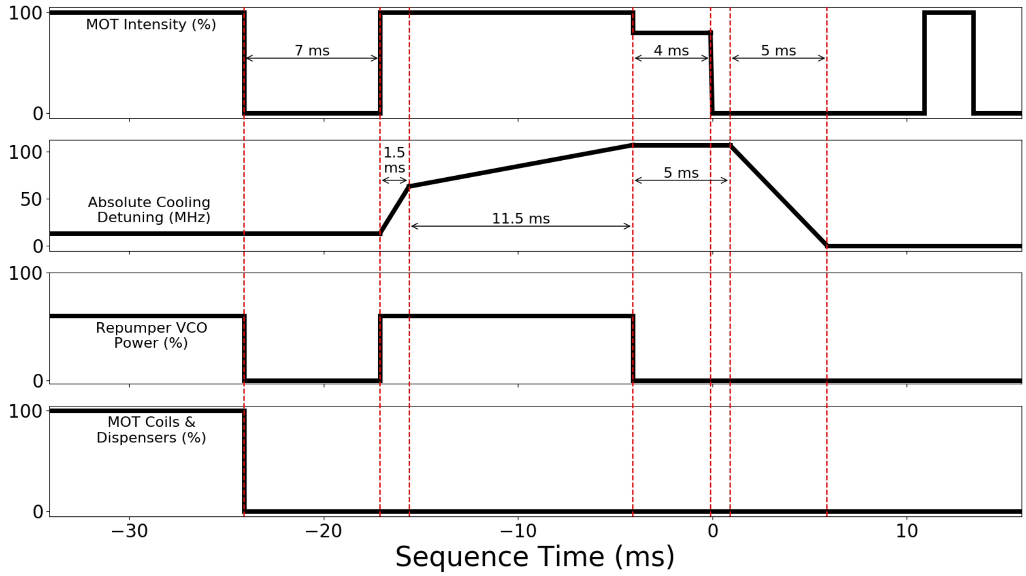

Figure 5.1: Optical molasses cooling sequence and state preparation.

atoms apertured by the detection window or seeing substantially different Raman intensities as the cloud spreads over the Gaussian intensity profile. Cooling the cloud further utilises an optical molasses sequence to provide cooling below the Doppler limit by deactivating the coils, removing the trapping force, and changing the light frequency and/or intensity to gradually cool the atoms whilst reducing the recoil effects. The optical molasses sequence (shown in Fig. 5.1) was developed from those used in [26] and [61] and then optimised for this system.

The molasses sequence has two features of note; one of these is the fact that the MOT light turns off between t =-24.1 ms and t =-17.1 ms. At t =-24.1 ms dispensers and MOT coils are deactivated, given the sub-microsecond response of the AOM and the 1.2 μs turn off time for the IGBTs[1] (with a ~1 ms turn off time measured for the coil) means the light turns off before the current carrying components. The dark time allows the eddy currents generated by this rapid switch off to dissipate without inducing a kick on the atoms, this has been found experimentally (Fig. 5.2) to reduce the temperature of the Top cloud and effect the time of arrival for both clouds. During this time period the cloud is at the MOT temperature and will therefore experience rapid expansion when the confining forces are removed, a shorter dark time will

---

[1] *8650713*

result in a more localised cloud and a better fringe contrast, the result is a trade off between these two effects. The shortest dark time presented here is 4 ms as at shorter times the Top cloud was impossible to distinguish above the background light.

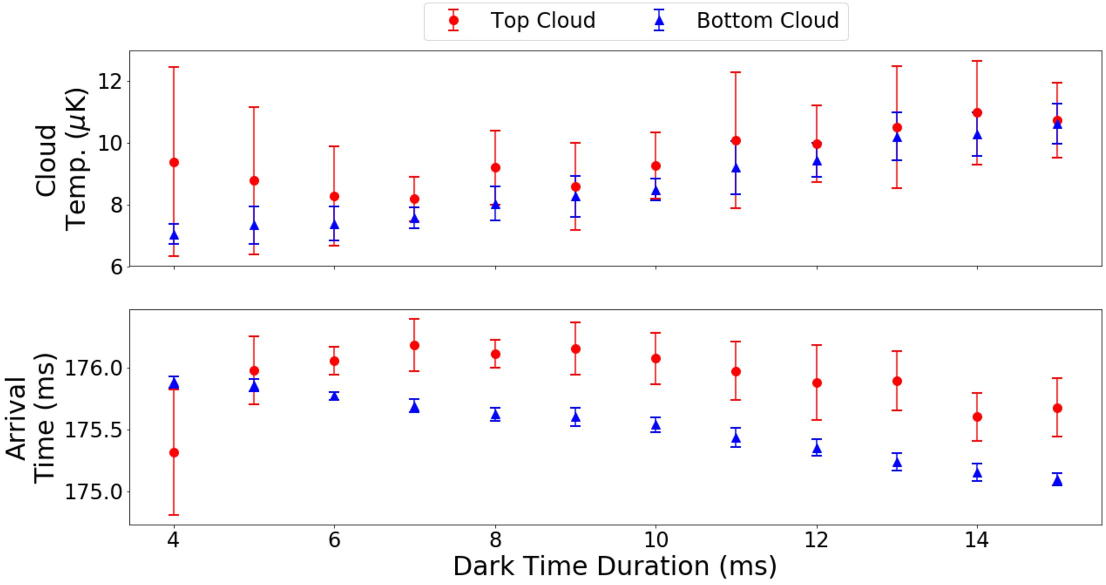

Figure 5.2: Effect of pre-molasses dark time on cloud at detection.

For the Bottom cloud increasing the dark time appears to only have a negative impact, with its temperature increasing in a linear fashion by ~4 μK over the 11 ms scanned and the arrival time trending away from the value of 178 ms as expected from ballistics. Conversely, the temperature of the Top cloud decreases till a dark time of 7 ms before beginning to increase, with a similar convergence towards the expected time of arrival levelling out in the 7-9 ms dark time range before diverging. This is to be expected as during the dark time the atoms are in free fall until the molasses process begins and the optical forces are reintroduced. As a result of this free fall time after 4 ms (or 15 ms) the two clouds will be 0.08 mm (1.10 mm) lower in the system than they were as a MOT. This would result in a difference in time-of-arrival between the 4 ms and the 15 ms dark times of ~0.6 ms. The arrival times measured in the Fig. 5.2 vary by 0.78 ms for the Bottom chamber and 0.87 ms for the Top meaning this initial free fall will be a substantial contribution to this offset. The calculation of the time-or-arrival offset has been made using the assumption that as soon as the optical forces return the centre of mass of the cloud will instantly stop falling, this residual velocity will introduce an additional effect that has not been considered here.

The difference in response for the two clouds potentially arises from the proximity of the dispenser to the MOT region in the Top chamber. A high current carrying wire in close proximity to the cloud will generate a strong source of eddy currents that need to dissipate. At a dark time duration of 7 ms the linear response seen in the Bottom cloud begins to manifest in the Top cloud, given that this corresponded to a similar temperature of ~8 μK and time of arrival in both clouds this value was taken to set the dark time during the sequence.

The other feature in the sequence is the dip in MOT power at t =-4 ms which coincides with the deactivation of the repump light. The power drop is to account for the light from the repump sideband that will return to the carrier when the EOM is turned off (~ 20% of the total optical power). The sideband is switched off to provide state preparation for the system; due to limitations of the cooling laser it is impossible to have a pure repump pulse. The consequence of turning off the sideband is that a clean $|F = 1\rangle$ starting state for the cloud can be prepared collecting the atoms in the dark state then blowing away any remaining $|F = 2\rangle$ atoms with a pulse of $|F = 2\rangle \rightarrow |F^{'} = 3\rangle$ light.

The effectiveness of this state preparation method can be tested by removing the blowaway pulse at t =10.9 ms. Scanning of the duration of the state preparation time will provide information of the minimum off time required to reach maximum population in the clean state. The longer the repump light is left off, the longer the atoms that decay into the $|F = 1\rangle$ state will be removed from the cooling cycle. It thus becomes a trade off between signal and cloud temperature. The results from Fig. 5.3 are gathered with the state readout method that will be described in 5.1.3. In these results it can be seen that after 4 ms almost 100% of the atoms can be transferred in the Bottom chamber and 85% in the Top. At ~5 ms the population of the $|F = 1\rangle$ state begins to drop off suggesting that the effect of removing the atoms from the cooling cycle is becoming dominant. These competing factors placed the optimum duration of the state preparation time at 4 ms, roughly the midpoint of the peak in Fig. 5.3. The final step in the sequence is the aforementioned blowaway pulse at t =10.9 ms, delivered through the MOT port, this ensures that the initial cloud is in a clean $|F = 1\rangle$ state.

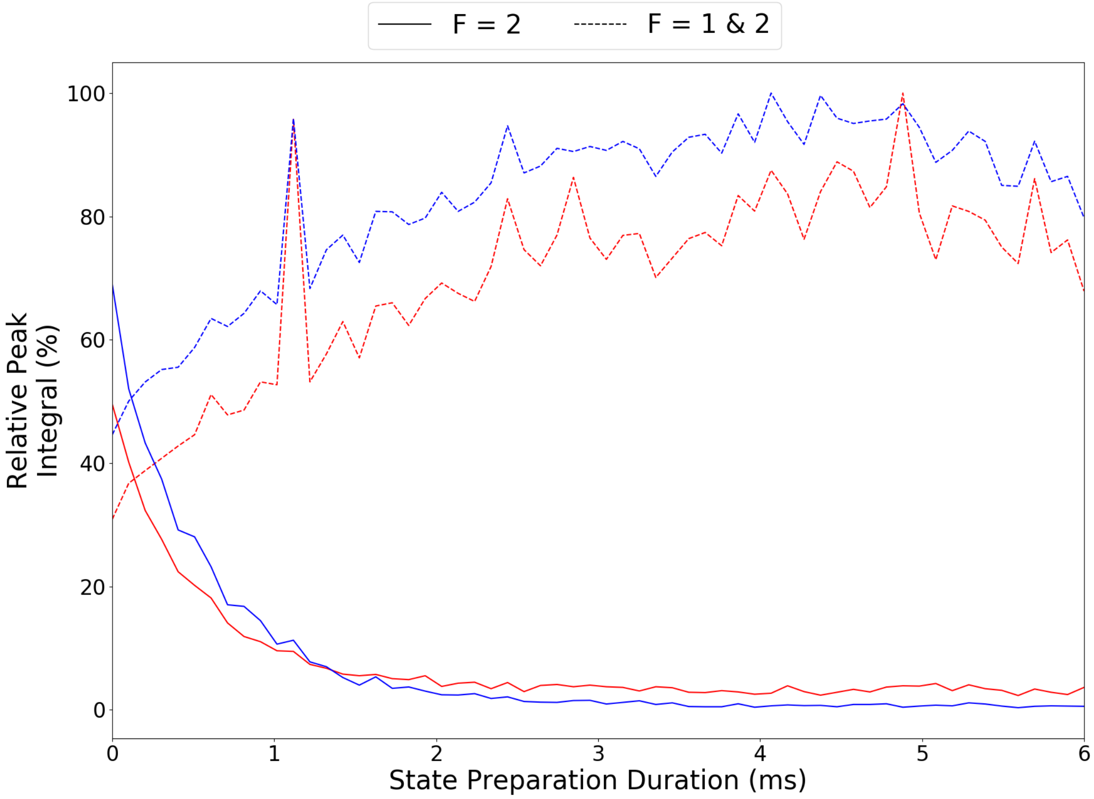

Figure 5.3: Scan of state preparation time, 0% corresponds to no atoms detected above the background pulse but 100% is just the maximum of the dataset. See 5.1.3 for the detection process used to gather these results.

5.1.2 Raman Pulses

As described in 2.3.3 the frequency difference between the Raman beams needs to be chirped with a rate α to ensure that the two photon process stays on resonance with the ground state splitting. Magnetically sensitive substates in the hyperfine ground states of ^{87}Rb means that if the degeneracy is not strongly broken then weak external fields can introduce a varying transfer rate in addition to the gravitational phase shift. In an ideal system the atoms prepared in a clean $|F=1\rangle$ state would be placed in a large magnetic field to break the degeneracy of the states and a microwave frequency pulse used to transfer only the atoms in $|F=1, m_f=0\rangle$ into $|F=2, m_f=0\rangle$ and the atoms in the $|F=1, m_f=\pm1\rangle$ state would then be blown away. This process requires both an independent repump laser and a microwave delivery source in the sensor head, neither were possible to include in this system so a complete state preparation was not available.

The impact of the atoms in these magnetically sensitive states can be mitigated, by placing the falling cloud in a bias field that is strong enough to break the degeneracy between the states. It is possible to address the $|F = 1, m_f = 0\rangle \rightarrow |F = 2, m_f = 0\rangle$ with minimal impact from the other states (see 5.4.1 for a more details). It can be seen from the spectrum scan (Fig. 5.17) that by $t = 65$ ms the atoms are sufficiently far inside the bias coils (3.2.2) that this degeneracy is broken. It is therefore possible to interrogate the cloud inside this region with minimal worry about external magnetic fields.

During an interferometry sequence the separation between the Raman pulses (T-time) is $T = 50$ ms, it was found that the optimum time to start a sequence is t =68 ms, placing the other two Raman pulses at $t = 118$ ms and $t = 168$ ms. The pulses at $t = 68$, 168 ms are the $\pi/2$ pulses and that at $t = 118$ ms is the π pulse, ideally the pulse type would be the only impact on the duration however (as will be discussed in 5.4.2) there is an inhomogeneity in the Rabi frequency with height. The frequency chirp on the Raman modulation frequency is not a true continuous chirp and is performed by making a phase continuous step in the RF signal that drives the Raman IQ modulator 1 ms before the pulse occurs.

Before any given experimental run the Rabi frequency is characterised and due to the spatial inhomogeneities an oscillation is performed at each of the three pulse heights. As discussed in 2.3.3 the definitions of the $\pi/2$ and π pulse durations come from the time taken to induce a Rabi flop between two states and place an atom into either an equal superposition of the two states or completely invert the states respectively. To measure this in a real system an AOM applies a variable length Raman pulse to the cloud, the relative state population is then read out. Plotting this ratio against the length of the applied pulse traces out a Rabi oscillation.

The oscillation in Fig. 5.4 is heavily damped due to the large range of velocities in the cloud perpendicular to the Raman beam. The atoms are spread over a large area of the Gaussian intensity profile of the Raman beam leading to the presence of many different Rabi frequencies overlapped for a given measurement. The same effect will result in increasingly damped oscillations as the cloud falls further. The fit in Fig. 5.4 follows the same form as Eq. 2.23 with $\omega = \Omega_{Rabi}$.

In an undamped system the π-time is half the period of oscillation of the Rabi frequency Ω_{Rabi} and

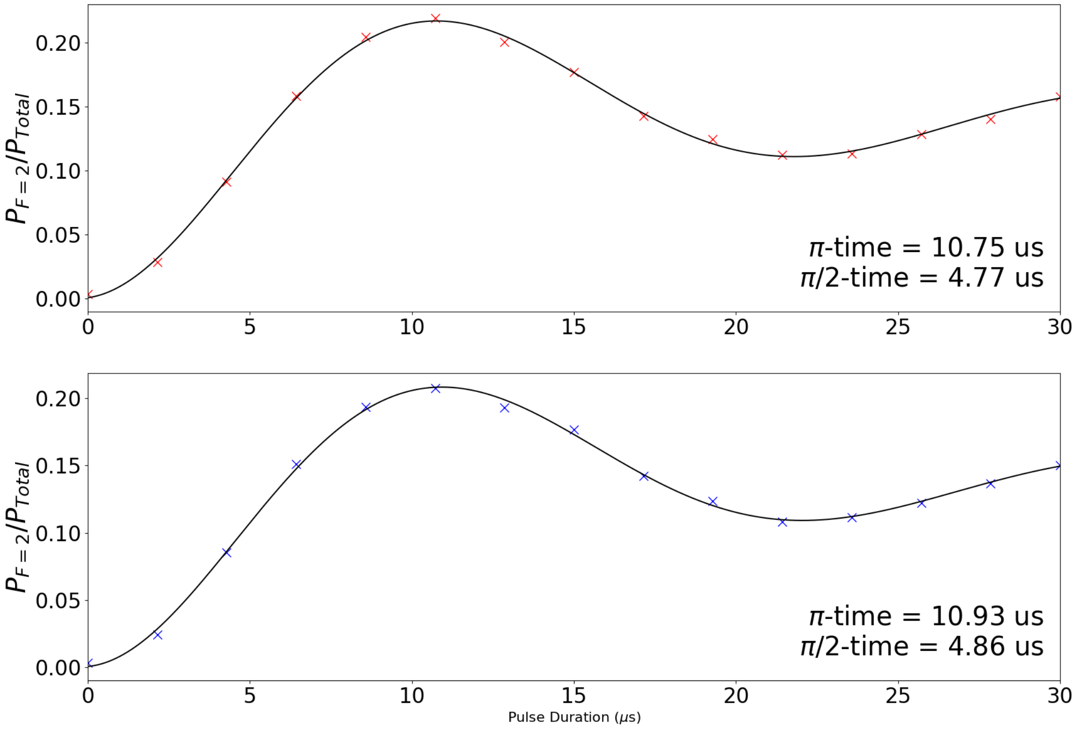

Figure 5.4: Typical Rabi oscillation for a Raman pulse at t =68 ms.

the $\pi/2$-time is half of that. However for a heavily damped oscillation such as that in Fig. 5.4, this definition no longer holds as the the time at which half the population is transferred is not the same as half the time the whole population is transferred. As a result the definition is taken to be the time at which half the population is transferred, these are the $\pi/2$-times as presented in Fig. 5.4. Figure 5.5 shows the Rabi oscillations for the times of the three interferometry pulses, as the cloud falls there is clearly an increased damping effect on the oscillation. As the thermal atoms fall further they are given longer to expand and will be spread over a greater extent of the intensity profile of the Raman beams, leading to more different Rabi frequencies present in the oscillation.

A choice was made here to select the π and $\pi/2$ times based on the average Rabi frequency at a given height. This will produce a fringe with the largest amplitude as the most atoms will be interacted with as the average Rabi frequency is used. This means that hot atoms have an increased impact on the final result, an alternative under investigation is to use the π and $\pi/2$ times as defined by the Rabi oscillation at the highest point of the system, where the cloud is most dense and thermal effects will be minimised. This will allow the Raman pulses to interact with only the coldest atoms, increasing the contrast but leading to a drop in the fringe amplitude.

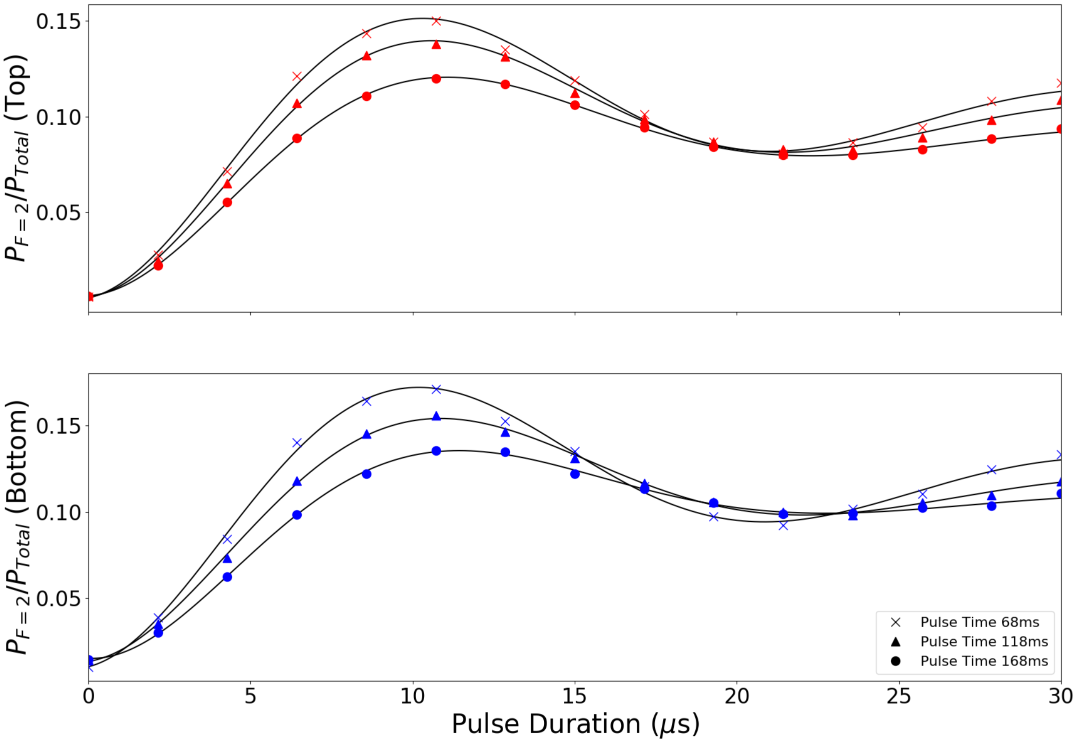

Figure 5.5: The result of performing a Rabi oscillation at each of the three pulse locations. The increased damping effect arises from the thermal cloud spreading out over a larger distribution of Raman intensities.

5.1.3 Detection

To provide phase readout of the interferometer the population of the $|F=2\rangle$ state needs to be determined; if the state preparation produces a completely clean initial state then the only atoms that will end in the excited state must arise from the Raman pulses. There are a number of ways to generate this readout that provide varying levels of information. They all rely on one of two measurement techniques: absorption imaging[86] which measures the shadowing effect of the cloud when backlit by resonant light or fluorescence imaging[87] which measures only the light scattered by the cloud relative to a background level. Variations on these techniques can be employed to gain more information about a cloud, for instance a light sheet measurement as outlined in 2.3.2 is a type of fluorescence imaging that provides both spatial and time-of-arrival information about a cloud.

This system implements a basic form of fluorescence imaging which enables the cylindrical form factor; the cloud is illuminated by a set of pulses from the MOT input port removing the requirement for a separate delivery telescope. These are timed to illuminate the cloud to provide the signal and then only the chamber to provide a background level. Figure 5.6 shows the timing of the detection sequence; first a 100 μs pulse of cooling light only is shone onto the cloud resonant with only the atoms in the $|F=2\rangle$ state, 50 μs later both cooling and repump light illuminate the cloud for 100 μs giving a readout of the atoms in both the $|F=1\rangle$ and $|F=2\rangle$ states. After a 1 ms gap the same cooling and repump light is used to illuminate the region again, this is to clean the detection region by acting as a blowaway and provide a diagnostic tool for system noise; for example noise on the light frequency can be seen as an oscillation in the decay of this peak. Following a further 1 ms break the initial pulse sequence is repeated to give a background light level; first cooling light only then cooling and repump light. Subtracting this from the first two peaks will suppress background light level fluctuations. 50 μs later a pulse of Raman light is shone into the chamber for 300 μs to provide a diagnostic on the Raman intensity.

The power in the repump sideband for detection sequence 5.6 is not the same as during the MOT loading and molasses sequence (Fig. 5.1). As mentioned in 4.1 the EOM in the cooling laser has the option for a high power input, this higher power option is used to generate the sideband in detection. This provides more light resonant with the atoms of interest (those in $|F=1\rangle$) and results in a higher signal-to-noise ratio on this peak.

The pulses are labelled $1-6$ as they appear in the detection sequence and it is the sum over their duration that is used in the data analysis. These sums will be labelled with the peak numbers, i.e. for peaks $1-6$ they have sum $I_1 - I_6$ (see Fig. 5.6). If the first pulse is considered sufficiently short, minimal atoms are transferred from the $|F=2\rangle$ state into the $|F=1\rangle$ and this can be treated as a two state readout. Atom number fluctuations can then be removed with the following normalisation process

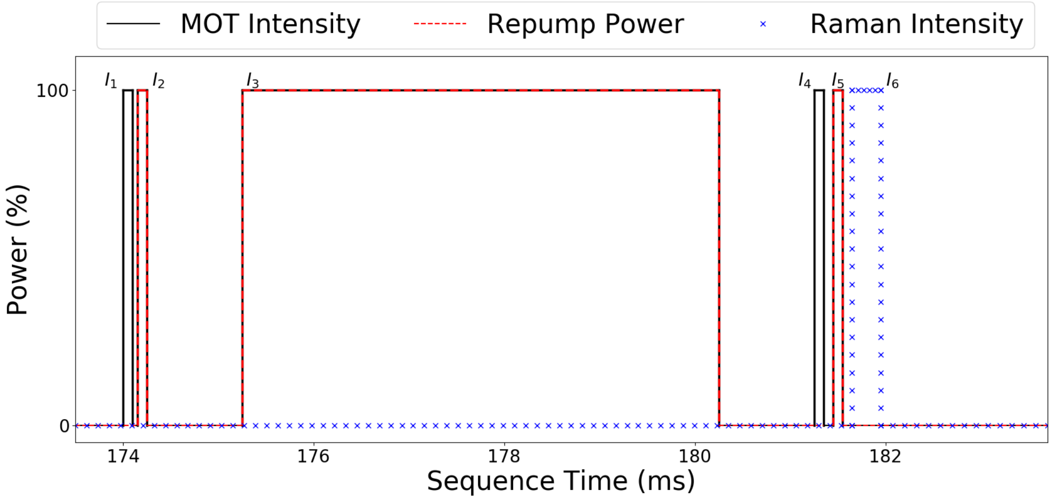

Figure 5.6: Detection sequence for the system. The peaks are labelled by their position in the sequence, i.e. from left to right they take labels 1 – 6.

$$\frac{P_{F=2}}{P_{Total}} \simeq \frac{I_1 - I_4}{I_2 - I_5}. \tag{5.1}$$

Equation 5.1 is referred to as the “Peak Ratio" and provides the readout for all atomic data in this thesis. Some older data in this thesis used longer timings in the detection sequence (pulse 1 and 2 durations of up to 3 ms) in these the normalisation assumes that I_1 is the signal only from atoms in $|F=2\rangle$ and I_2 only from atoms in $|F=1\rangle$. For this data the normalisation is slightly different

$$\frac{P_{F=2}}{P_{Total}} \simeq \frac{I_1 - I_4}{(I_1 - I_4) + (I_2 - I_5)}. \tag{5.2}$$

This is only the case for a small amount of the data and it will be made clear when it is relevant. Unless otherwise stated Eq. 5.1 provides the value of the “Peak Ratio".

5.2 Temporal Variation

In a lab based experiment it is important to understand as much about each component and remove as many noise sources as possible. The same is true in a portable system, however the bounds put on acceptable noise levels will be different as the system only needs to be “sensitive enough" as described in 1.3. In this case some of the most interesting results arise from comparing two of the designs shown in Fig. 2.19. As previously mentioned the original experimental setup utilised a central retro-reflection mirror (Fig. 2.19c) to couple together the phase of the Raman beam whilst keeping each sensor as an independent module, after approximately a year of operation this central mirror was removed (Fig. 2.19d). Similar characterisation was performed in both cases and in some cases it has been seen to reduce noise sources whereas it has increased it in others. This section will make a comparison between these and present some additional system characterisation that has impacted the way the experiment has developed.

5.2.1 MOT Stability

In an atom interferometer a measurement is made with repeated iterations of the experimental sequence as outlined in the previous section whilst scanning the phase relationship between the two Raman lasers in the final pulse. This means that any inconsistencies between shots will manifest as an inconsistent reference and thus introduce noise on the signal. Shot-to-shot variation in MOT atom number will lead to changes in the absolute value of I_1 separate to any transfer induced by the Raman pulses it is possible to remove these changes through normalisation (Eq. 5.1).

Similar to initial atom number fluctuations if there is a variation in the cloud temperature then there will be a variation in the amplitude of I_1. This comes about as a result of the discussion in 2.3.2 where if the temperature of the cloud is doubled then the number of atoms that can be interrogated by the same duration Raman pulse reduces by a third due to the increased velocity distribution. To characterise the variability in this a series of light sheet measurements were performed over the course ~6 hours and put through the same processing as the data in Fig. 2.9. This measurement was performed in the both gradiometer configurations after optimising for each.

The first point of note is that the average temperature is similar for both configurations; 7.2 μK and 8.0 μK for the Top and Bottom chambers respectively with the mirror, 8.3 μK and 7.1 μK without the central mirror. The variation in the measured temperature is substantially worse without the mirror, both the chambers giving a deviation of ~1.6 μK up from ~0.5 μK with the mirror. This is easily understood if one considers Eq. 2.1, the radiation force experienced by an atom is related to intensity of the light. Figure 4.6b shows that, even when stabilised, the intensity from opposing telescopes has a slow differential drift

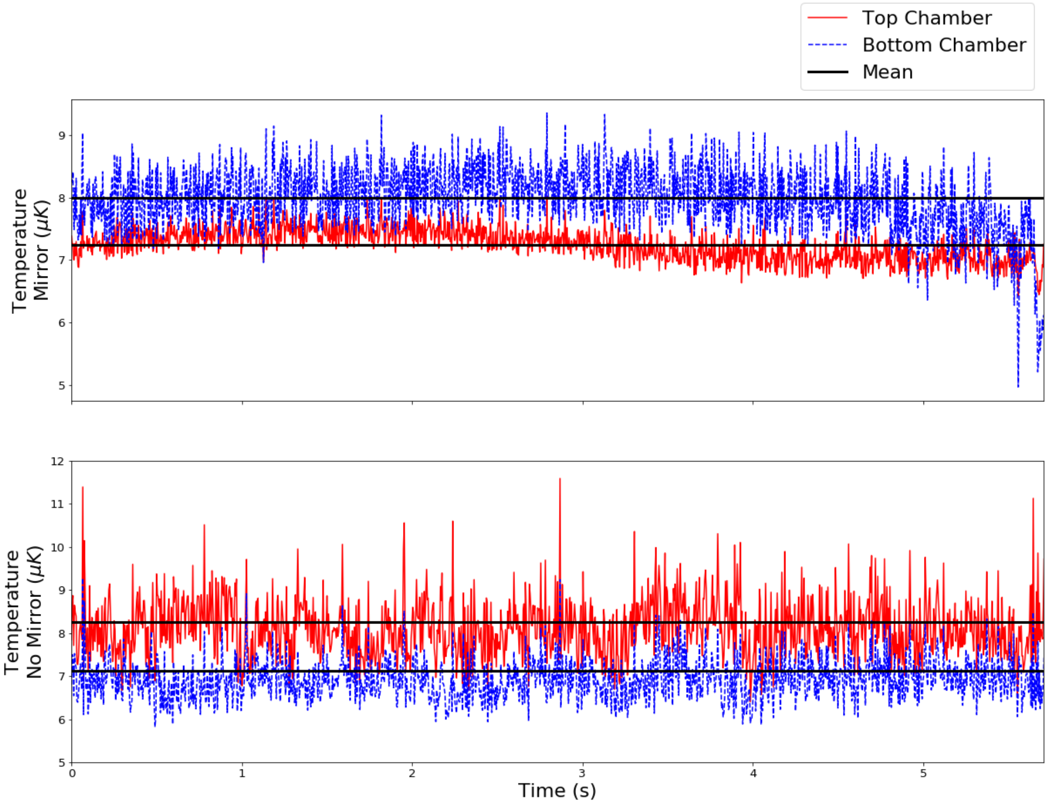

Figure 5.7: Top: The variation in cloud temperature with time with the central retro-reflection mirror. Bottom: The same result after the central mirror has been removed. There is a distinct increase in both the average temperature and the noise on the temperature in the case that the mirror has been removed, this is to be expected as variations in cooling light will no longer be common within a given module.

of ~ 2% and a quick noise of ~ 0.5%. During the late stages of the optical molasses sequence the detuning becomes large meaning that the atoms have an increased sensitivity to light intensity, this differential drift can lead to unbalanced cooling along the vertical direction during the molasses stage.

The same light sheet data can be used to determine the consistency of the time-of-arrival of the two clouds in the detection region as seen in Fig. 5.8. If the time-of-arrival variation arises as the result of the power instability during the molasses stage this means the atoms receive an initial “kick" before the drop which would register as a spurious acceleration in any measurements. This could also result from the MOT “walking" along the vertical direction as in Fig. 2.6, by removing the mirror the clouds have become susceptible to relative drift of the light out of the two experiment telescopes. This potentially means the walk seen in the six-beam case may be present along the vertical direction for the prism MOTs.

From ballistics it can be found that the difference between a time of arrival of 176.2 ms and 175.8 ms corresponds to a difference in starting position of 0.7 mm. This is well within the drift seen for a six-beam MOT in Fig. 2.6 due to the uncommon drifts between the cooling beams, so this conclusion has some merit. Both of these effects are additionally problematic as there are spatial asymmetries between the two clouds, meaning that either a position or velocity offset can introduce phase differences. These inhomogeneities will be explored later in this chapter.

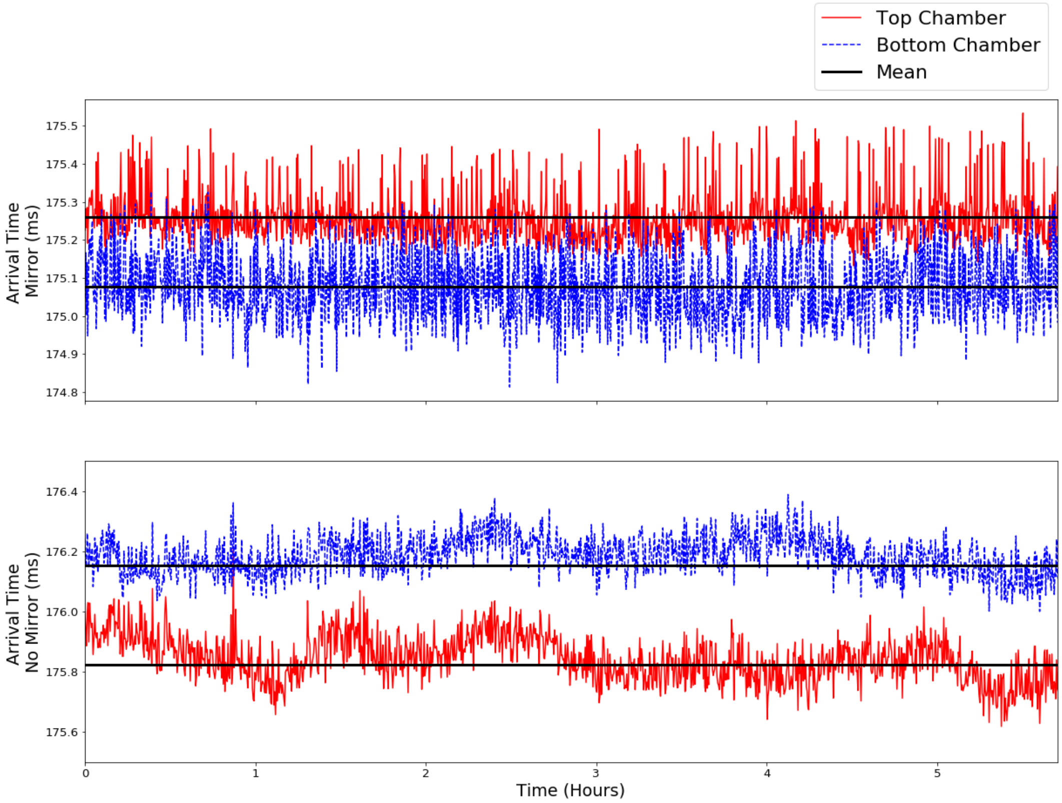

Figure 5.8: Variation in the arrival time.

Figure 5.8 shows an interesting effect of removing the mirror, one which was expected but noteworthy. By removing the isolation between the two sensors the independent variation in light level between the two inputs has created a less stable atom source, however this noise is common between the clouds; with a correlation in the dataset when the mirror is absent. Coupling the noise together in this way gives the potential to take advantage of the phase noise suppression that will arise from the use of common Raman beams, the main motivator behind removing the mirror.

5.2.2 Frequency

During the optical molasses stage the effectiveness of the cooling depends directly on the force of the light, over the course of the sequence; it is important that the way this force changes is consistent from shot-to-shot. The temperature noise of the cloud seen in Fig. 5.7 must arise from a variation in the force experienced by the atoms. The scattering force (Eq. 2.1) has a dependence on both intensity and frequency, it has been shown that when stabilised the intensity of the cooling light has a 0.5% variation over time (Fig. 4.6b). This will have an impact on the temperature variation of the cloud additionally as the light level into the chamber is low, i.e. $I/I_{sat} \approx 1$, the stability of the cooling laser frequency can have a large impact on the force seen by the atoms.

Figure 5.9a shows the result of measuring the frequencies of both the cooling and Raman laser with a *High Finesse WSU2* wavelength metre over a 14 hour period. The internal calibration laser frequency was also recorded and the noise level presented here with the dashed lines is the standard deviation of the frequency of this laser. Every half an hour the wavelength metre recalibrates to this calibration frequency causing the discrete jumps in cooling frequency. The data shows that the noise on the cooling laser is well within the noise of the calibration laser so can be considered stable to within 2 MHz over the full 14 hour period. The Raman presents similar discrete jumps but features an additional variation above this 2 MHz noise floor which is not present on the cooling laser.

The Raman laser is slaved to the cooling meaning that any frequency noise on the cooling would be passed directly onto the Raman. The initial 6 MHz spike on the Raman is not present on the cooling and thus it must arise as part of the locking sequence for the Raman. Figure 5.9b shows the effect this frequency variation would mean as a force on the atoms, found by adapting Eq. 2.6

$$F^{\text{scatt}} = \frac{\hbar\omega}{c}\frac{\Gamma}{2}\frac{I/I^{\text{sat}}}{1+I/I^{\text{sat}}+4\left(\Delta-\omega_N\right)^2/\Gamma^2} \tag{5.3}$$

The frequency variation as on the y-axis of Fig. 5.9a is given by ω_N, such that $\omega_N = 0$ when the laser is at the desired detuning. The values plotted in Fig. 5.9b are the force applied to the atoms by light of the given frequency as a fraction of the force at the intended frequency ($\Delta =$13.2 MHz). It is not possible to

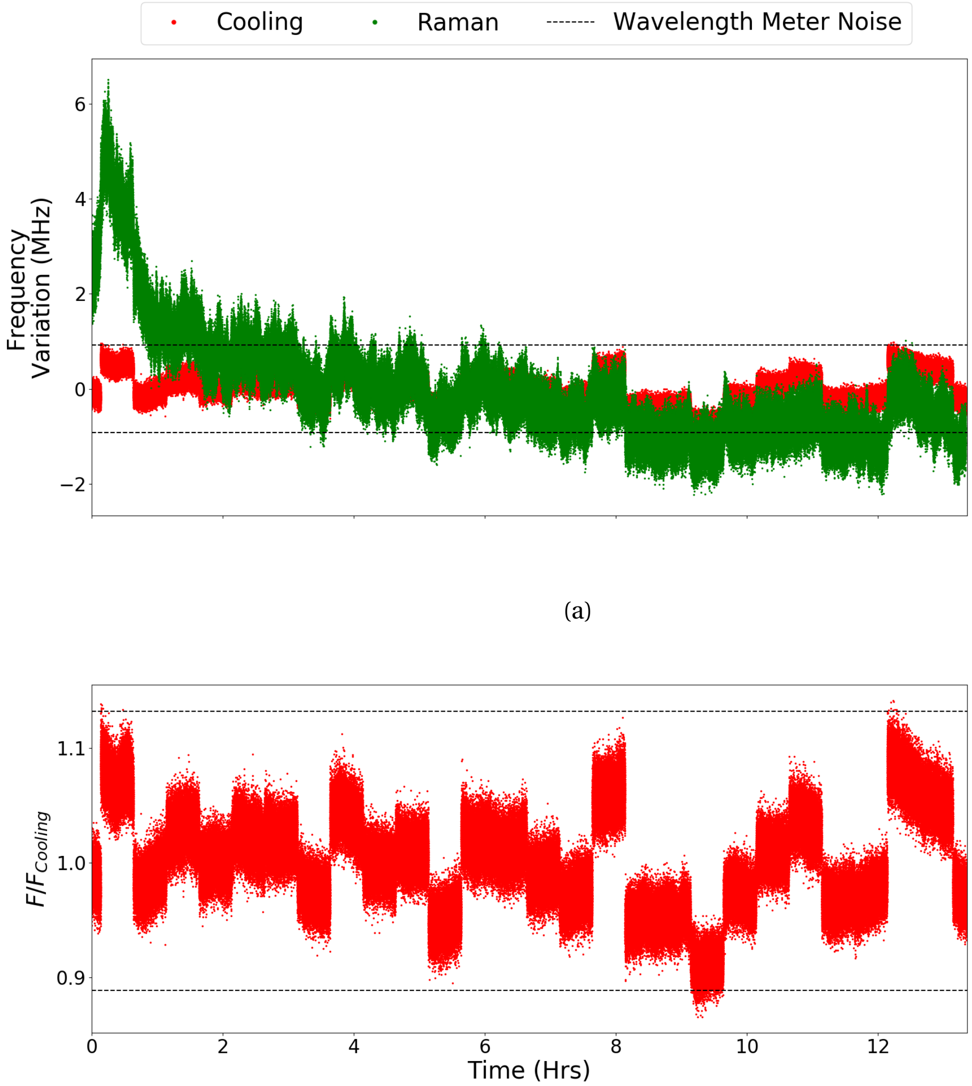

(a)

(b)

Figure 5.9: (a) Measurement of the long term stability of the experimental lasers as performed with a wavelength metre that has a noise level of ~2 MHz (represented by the dashed black line). (b) The impact this frequency noise would have on the scattering force of the cooling laser as a fraction of the force at the standard experimental detuning of 13.2 MHz.

know for certain due to the frequency noise on the wavelength metre but the noise on the cooling light as presented here could lead to as much as a 20% variation in the cooling force experienced by the atoms. Frequency variation of the Raman laser would change the effective Rabi frequency of the atoms (see Eq. 5.4) but given the ~834 MHz detuning a 6 MHz noise will produce at most a 0.5% change. It will be shown later in this chapter that larger variations occur due to spatial inhomogeneities and intensity variation so this noise can be considered negligible.

5.2.3 Raman Interactions

As outlined in 2.3.3 the principle of operation behind an atom interferometer is to use a set of two-photon Rabi oscillations to couple internal energy levels to the external momentum state of the atom. A two-photon Rabi oscillation between states $|g\rangle$ and $|e\rangle$ has a generalised Rabi frequency, Ω_{ge}, which is proportional to the Rabi frequencies of the two component transitions, Ω_{gi} and Ω_{ie}, via some intermediate state $|i\rangle$. This is given by the relation[88]

$$\Omega_{ge} = \frac{\Omega_{gi}\Omega_{ei}}{4\pi\Delta_R}. \tag{5.4}$$

Where Δ_R is the detuning from a resonant transition in a three level system, in this thesis $\Delta_R = 834$ MHz is the detuning of the Raman laser from the $|F=2\rangle \rightarrow |F^{'}=2\rangle$ transition[1]. The Rabi frequencies, Ω_{gi} and Ω_{ie}, are those associated with one of the two transitions via the intermediate state and as such are proportional to the magnitude of the incident light field[89]

$$\Omega_{ge} \propto |\mathbf{E}_{gi}||\mathbf{E}_{ei}| = E_{gi}E_{ei}. \tag{5.5}$$

Where Ω_{ni} is the Rabi frequency associated with the light field resonant with the $|n\rangle \rightarrow |i\rangle$ transition. In order to measure the Rabi frequency in the system a Rabi oscillation is performed by taking roughly 15 measurements with different length pulses and fitting a damped oscillation to the result. Repeating this measurement multiple times and monitoring the value of I_6 it is possible to see the impact of intensity

---

[1] In much of this thesis the Raman detuning is given as 1.1 GHz, this is relative to the $|F=2\rangle \rightarrow |F^{'}=3\rangle$ cooling transition as this is the frequency to which the Raman laser is offset locked. However as it is not possible to perform a transition from $|F=1\rangle \rightarrow |F^{'}=3\rangle$ it is less valid to refer to in this context.

variation on the π-time of the cloud. Results of this measurement are presented in Fig. 5.10 where the value of I_6 is given as a percentage variation about the mean.

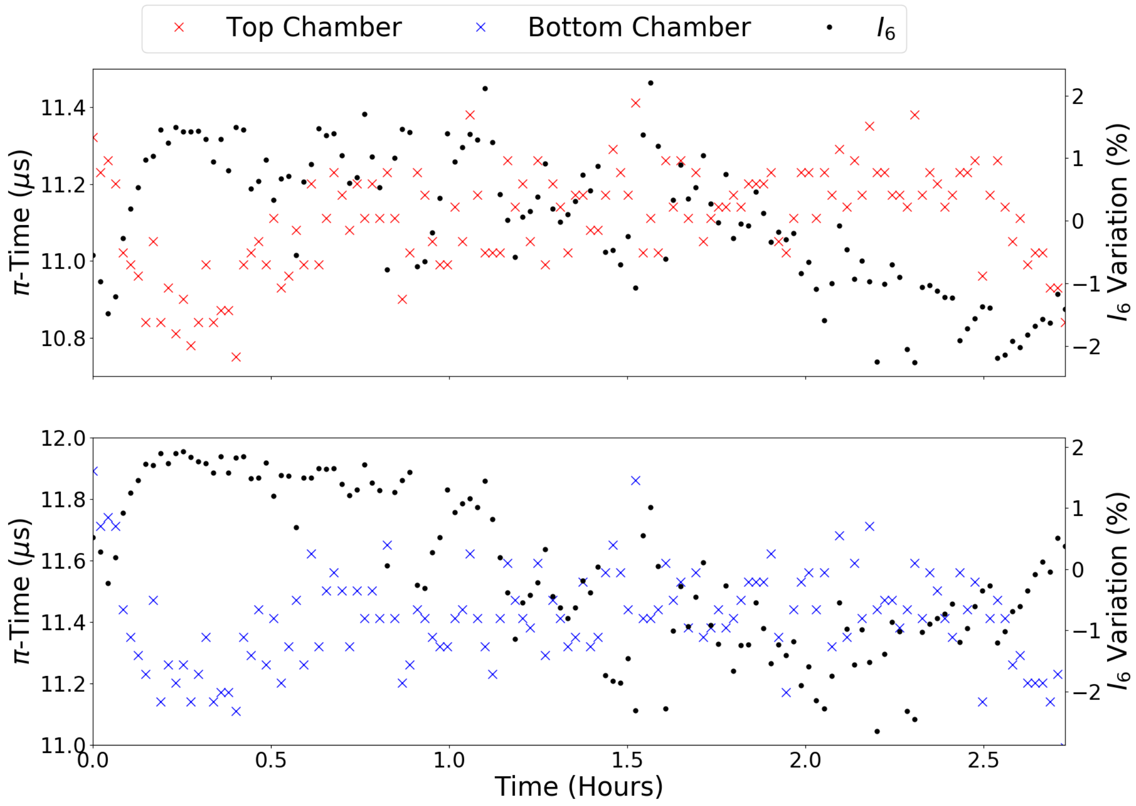

Figure 5.10: Variation of π-time without intensity stabilisation overlayed with the average value of I_6 for the 15 oscillation. There is a clear inverse correlation between the two, as expected.

From this data we see the expected result, as the Raman light into the chamber varies the Rabi frequency measured varies inversely. Using the process applied to the gravimeter data in 2.3.2 it is possible to gauge the impact this π-time variation would have on the percentage of atoms transferred by a pulse. Considering the Bottom chamber with an average temperature of 7.1 μK (from Fig. 5.7 in the absence of a mirror) and π-time range of 11.0 - 11.9 μs this will result in a variation in transferred atom number of 8.6 – 8% respectively. This means that a variation in Raman light power of ~ 4% can lead to variation in the value of I_1 of almost 10%.

With sufficiently good intensity stabilisation this variation in π-time can be reduced. However Eq.

5.5 shows that the effective two-photon rabi frequency will scale with the optical power in each of the contributing Raman frequencies. The light seen by the stabilisation diode is a combination of the carrier and all the sidebands and can only provide a measurement of total power. That is to say

$$E_{\text{Total}} = E_{gi} + E_{ei}, \tag{5.6}$$

is kept constant by the intensity stabilisation. It cannot account for any drift in the relative size of these fields, if $E_{gi} \to E_{gi} - \Delta E$ and $E_{ei} \to E_{ei} + \Delta E$ then Eq. 5.6 will stay constant but Eq. 5.5 will vary. This means that a sideband ratio variation will lead to a noise on the Rabi frequency even if the total light intensity is kept constant.

There is an additional nuance that arises from the optical delivery system; due to independent variations in polarisation, coupling efficient and splitting ratio there will be intensity noise on the two Raman input arms that cannot be removed by the stabilisation. In this system both telescopes feature light that is resonant with two Raman frequencies and an appropriate choice of chirp rate determines which paring is used to probe the transition (see 5.3.1). For now take E_{gi} to be stable but impose the 2% uncommon intensity noise seen in Fig. 4.7 onto E_{ei}. This will impose a direct 2% noise onto the Rabi frequency due to Eq. 5.5, the effective Rabi frequency is not only sensitive to the sideband ratio of the light but also the relative intensity out of the two telescopes.

This sensitivity to sideband ratio is problematic as it has been shown to vary with both the frequency of the RF input to the IQ modulator (Fig. 5.11a) and with time (Fig. 5.11b). These noise sources will map directly onto the effective Rabi frequency within the chamber and will later be shown as the leading order noise source in gradient measurements (6.4.1). As the varying sideband ratio is not visible in the value of I_6 it is monitored with an optical cavity. The labels on the peaks in Fig. 5.11a are associated with which of the frequencies in the modulated light it is associated with: R_1 is the carrier, R_2 is the desired first order sideband, $R_2^{'}$ is the unwanted first order sideband and R_3 is a higher order sideband (Fig. 5.12).

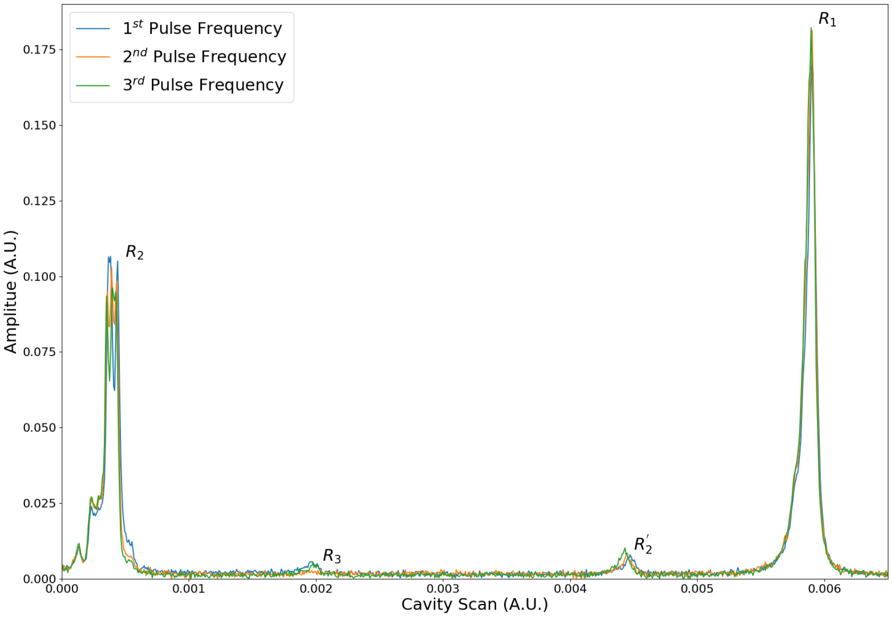

(a)

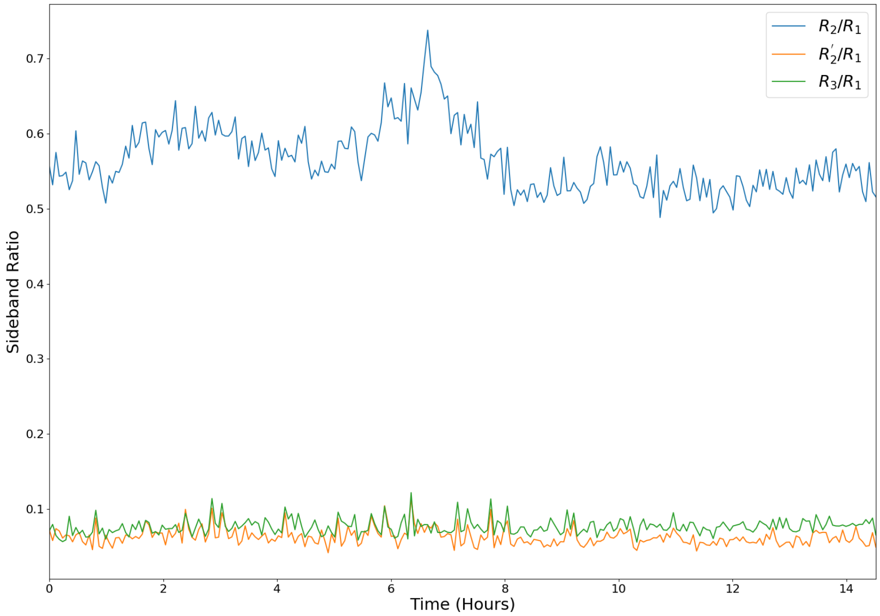

(b)

Figure 5.11: (a) Shows the relative optical power of the sidebands for optimised IQ voltages (see Fig. 5.12 for the peak definitions). (b) Shows how these peak sizes vary over the course of 14 hours.

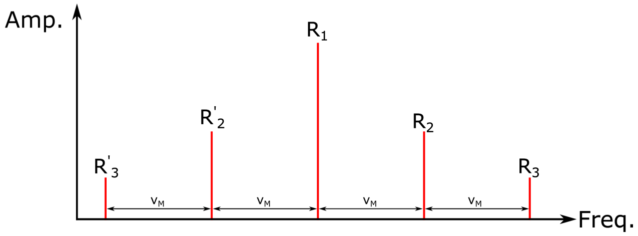

Figure 5.12: When an EOM is driven with modulation frequency ν_M it creates an infinite set of sidebands separated by ν_M.

5.3 Chirp Rate

The chirp rate, α, is frequency change per unit time applied to the driving frequency of the IQ modulator in the Raman laser to ensure the modulation frequency keeps the two-photon transition on resonance as the cloud drops. This rate of frequency change can be used to perform a precise measurement of gravitational acceleration[32], this can be seen from Eq. 2.27 (restated here for convenience)

$$P_{|E,p+\hbar k_{\text{eff}}\rangle} = \frac{1}{2}\left[1 - \cos\left(\left(\alpha - k_{\text{eff}} g\right) T^2 + \phi_{\text{laser}}\right)\right]. \tag{5.7}$$

If $\alpha = k_{eff} g$ then the only variation in the excited state population will arise from the laser phase. A measurement of α is made by scanning its value for a given T-time; this is then repeated for two other values of T. Three sets of measurements are required to identify the point at which all oscillations will intersect allowing for precise determination of the chirp rate. This should only be possible for non-harmonic values of T in the case that $\alpha - k_{eff} g = 0$, providing an accurate value of gravity for a known value of k_{eff} up to the $1/T^2$ resolution of the chirp rate scan.

Performing this measurement for the gradiometer system provides the correct chirp value that will maximise the signal from a given sensor. The measurement shown in Fig. 5.13 has a noticeably worse signal-to-noise ratio for the Bottom chamber as this is dataset was captured during a period of low signal in the Bottom chamber. This set of historic data uses the normalisation scheme set out for by Eq. 5.2.

High frequency measurement noise has been filtered by performing a binned averaging process on sets of 20 nearest neighbour points. It is not completely clear from the measurements presented here but there does appear to be an approximate overlap point with a value of $\alpha \approx$-25.177 MHz/s. This is taken from

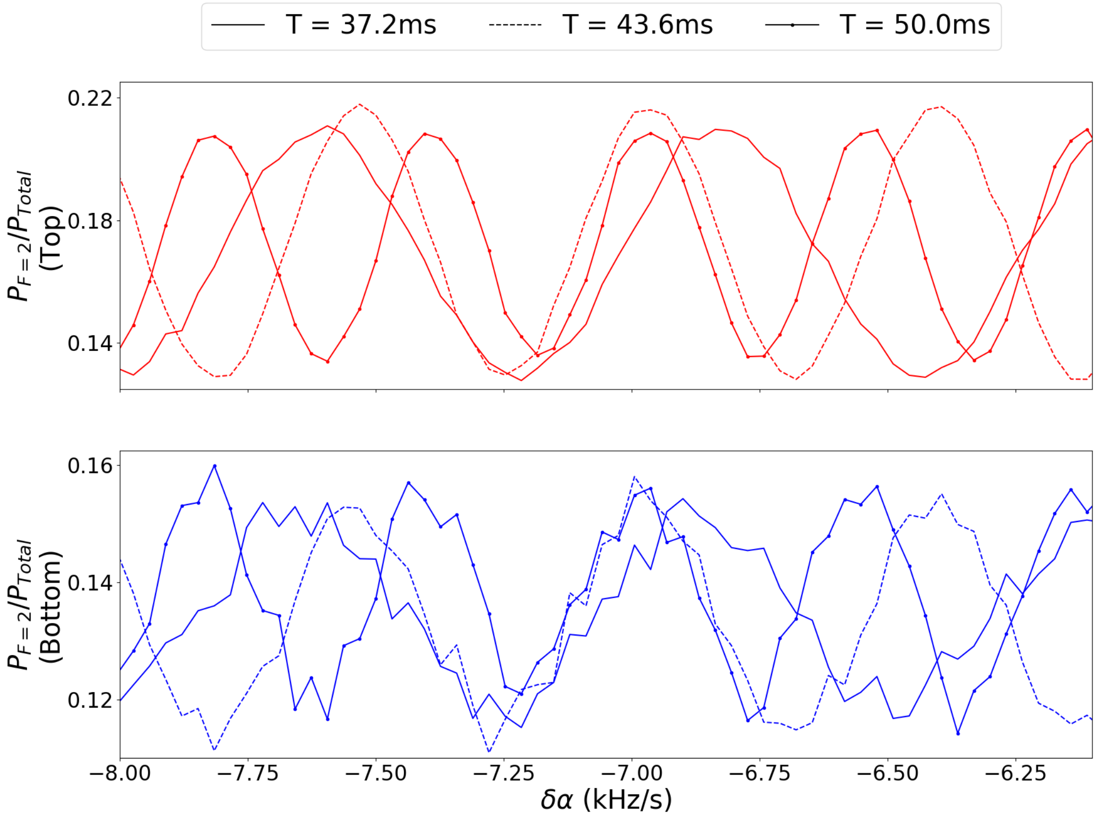

Figure 5.13: Chirp rate scans performed on the two clouds, the true chirp rate $\alpha = -25.17 \times 10^3 + \delta\alpha$ kHz/s with the offset applied to make the data more legible. This is a historic data set (December 2017) and uses the older normalisation style as mentioned in 5.1.3.

the data for the Top chamber, in the data for the Bottom no definitive value can be taken, but it does not outright disagree. At the University of Birmingham there is a well optimised gravimeter with a T-time of 120 ms that has performed these measurements and given a value of α_{ref} =-25.1649 MHz/s[32]. The difference is expected as any systematic offsets between the two setups will manifest here.

5.3.1 Choice of Raman Beams

It is useful to stop here briefly and consider what this chirp rate means in terms of the system. As the single light source is used for both telescope inputs the same R_1 and R_2 are present in both beams. This means that there are two possible orientations that each of the Raman contributions can take, either R_1 delivered by the Top telescope and R_2 delivered by the bottom or the inverse. As will be explained the choice or chirp rate sets which set is used.

In this system the frequency difference between R_1 and R_2 is set by the modulation frequency of the IQ modulator, ν_M. This is generated by beating a variable DDS source against a constant 7 GHz reference and passing the result through a low pass filter. The chirp rate is applied to the variable DDS frequency so the modulation frequency at any given time is thus given by

$$\nu_M(t) = 7000 - (\nu_0 + \alpha \Delta t) \text{ MHz}. \tag{5.8}$$

Where ν_0 is a reference frequency at reference time t_0 and $\Delta t = t - t_0$ is the time between this reference and the time the pulse is performed. The chirp rate α is the rate of frequency change and can be either positive or negative; the modulation frequency will increase if $\alpha < 0$ and decrease if $\alpha > 0$. To explain how the orientation of the Raman beams is chosen let us consider the system in Fig. 5.14, the two possible cases are represented by the coloured arrows.

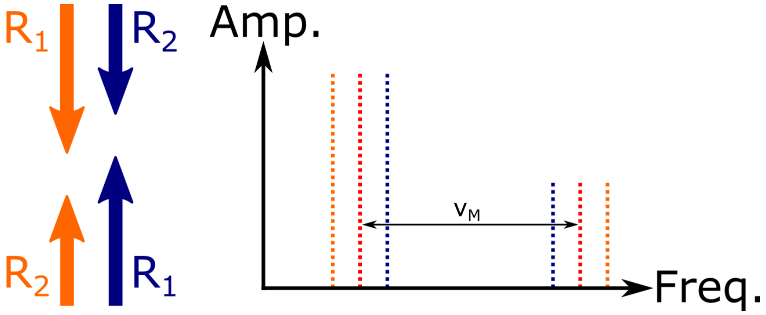

Figure 5.14: Frequency shift depending on chirp rate choice

In this system only the carrier and the higher frequency sideband are used. Let us consider the two possible orientations for R_1 and R_2 in turn. In the case of the blue arrows (R_1 from the Bottom and R_2 from the top) the falling atoms will see the frequency of R_1 increase and R_2 decrease due to the Doppler shift, to stay on resonance with the two photon transition ν_M must therefore increase, requiring $\alpha < 0$, similarly the left hand (orange) case will require $\alpha > 0$.

The standard for this experiment is to take $\alpha < 0$ meaning that it operates with R_2 delivered by the Top telescope and R_1 delivered from the Bottom. This can also be used to determine in which direction the atoms receive their momentum kick. As the atoms are prepared in the $|F = 1\rangle$ state and the frequency of

R_2 is greater than that of R_1 they must first absorb an atom from R_2 and then emit one into R_1, producing a net momentum kick in the direction of gravity.

5.4 Spatial Variation

A stable atom source is required to create an atom interferometer that must then be interrogated by three repeatable Rabi oscillations in a uniform environment. Variations in the conditions under which these oscillations occur introduce systematic offsets or noise sources onto the atomic readout. In this system there are spatial inhomogeneities that arise from the magnetic field profile within the chamber and a spatial variation of the Rabi frequency.

5.4.1 Magnetic Field Profile

When Rabi flopping is induced between the two hyperfine ground states of ^{87}Rb, there are a number of possible transitions to transfer the atoms between the $|F = 1\rangle$ and $|F = 2\rangle$ states (see Appendix A of [26]). In the absence of a magnetic field all these paths are degenerate and will contribute to the measured Rabi oscillation at the same two-photon resonance frequency. Introducing a bias field breaks this symmetry depending on the way in which the states are aligned with respect to the applied field. This field induces a frequency shift, $\Delta\nu_{\mathrm{Z}}$ on the state proportional to the value of m_f [90]

$$\Delta\nu_{\mathrm{Z}} = \frac{\mu_B g_F m_f B}{h}. \tag{5.9}$$

Where μ_B is the Bohr magneton and $g_F = -\frac{1}{2}, +\frac{1}{2}$ is the hyperfine Landé g-factor for the $|F = 1\rangle$ and $|F = 2\rangle$ states respectively. The opposing signs on the frequency shift of the two hyperfine states means the net frequency shift for a transition between them is twice the value given in Eq. 5.9. As discussed in 2.3.2 a Raman transition has a Fourier limited linewidth so a pulse of length τ_R has a transition linewidth of $\Delta\nu_R = 1/\tau_R$. A sufficiently strong magnetic field allows probing of this Zeeman dependent substructure with a two-photon Raman pulse. If the splitting induced by the field is sufficiently large relative to the linewidth of the Raman transition it is possible to detect this substructure by scanning the two frequency difference of R_1 and R_2. For a typical $\pi/2$-pulse duration, $\tau = 5\,\mu$s, a field of <285 mG would allow this

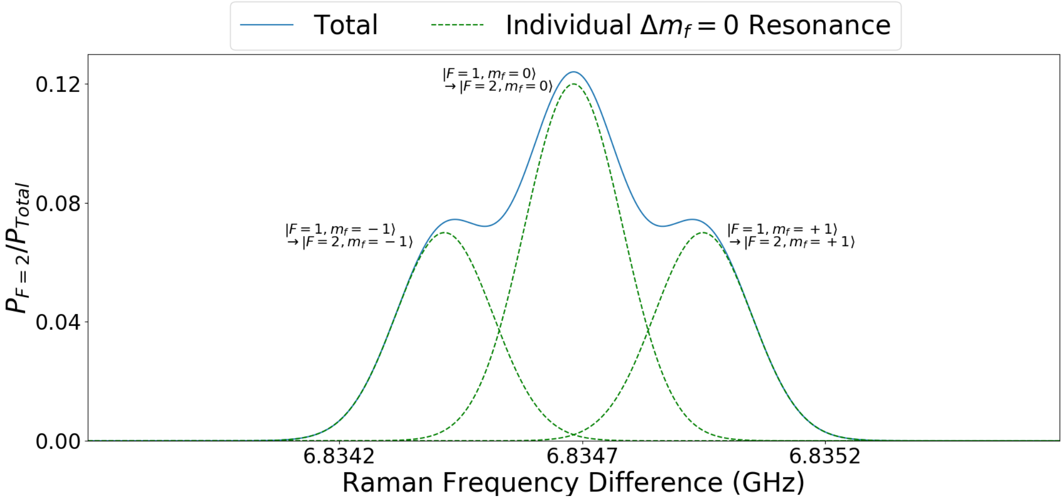

Figure 5.15: Simulated Zeeman splitting in 190 mG magnetic field an a Raman pulse of τ =5 μs.

structure to be probed with less than a 1% contribution from the other m_f states within 3σ ($\sigma = 1/\tau$) of the central $m_f = 0$ frequency. This creates a problem in a sensitive experiment. Consider the case presented in Fig. 5.15, where scanning the frequency difference between the two Raman lasers allows identification of the Zeeman substructure.

In making a measurement of gravity the central $|F = 1, m_f = 0\rangle \rightarrow |F = 2, m_f = 0\rangle$ resonance is used to ensure that any measurement is independent of the first order Zeeman effect giving reduced sensitivity to external magnetic fields. It can be seen that in this very low field (B =190 mG) case choosing the central resonance will still give a contribution to the interferometer from the magnetically sensitive states, approximately 14.2% contribution within 3σ of the central peak frequency. Over the course of a survey external magnetic fields are likely to vary from site-to-site and shot-to-shot and will be an uncommon contribution to the two clouds. In the above example changing the applied field from 190 mG to 191 mG changes the contribution from the side peaks to 13.9%. The effect of this can be mitigated by applying a well defined bias field that provides good splitting between the m_f states.

In gravimeters the uniformity of this field over the course of the interferometry region is important as inconsistencies will lead to small changes in the second order Zeeman shifts on the $m_f = 0$ resonance and field gradients will induce acceleration onto the atoms. Any offset which is constant shot to shot will change the measured value of gravity proportionally, in a gradiometer system this will manifest if the gradient profile is different for the two clouds. In which case it will manifest as a phase difference, if this difference is constant then it will be systematic on all measurements and not affect the output from a survey.

If a strong bias field is applied it is possible to use this frequency shift to map the magnetic field profile within the interferometry region. Figure 3.9 presents the field profile of the gravimeter bias coil measured with a Hall probe before the coils were mounted onto the system. If the atoms don't fall centrally through the coils or there is a secondary source of magnetic fields then this would skew the profile seen. Making measurements of the m_f spectrum over the course of the drop it is possible to map the precise magnetic field profile seen by the atoms. With this it will be possible to identify the locations at which the peak splitting seen in Fig. 5.15 will be large enough to allow a (first order) magnetically insensitive Raman transition.

The results in Fig. 3.9 are for the 18 turn gravimeter bias coil powered with 200 mA to produce a maximum field of ~260 mG. The gradiometer features two identical pairs of coils both powered with 300 mA, this would give an expected magnetic field strength of ~400 mG. This can be converted into an expected frequency shift of ~530 kHz on the $m_f = \pm 1$ transitions using Eq. 5.9. In order to see the resonances of the three Zeeman sensitive transitions a scan should be performed over a ~1.5 MHz range.

The scan range in Fig. 5.16 is much larger than the suggested 1.5 MHz as the frequency shift given by Eq. 5.9 is an incomplete model. As the atoms fall they accelerate under gravity causing the frequencies of the two Raman beams to be Doppler shifted away from resonance. To maintain the resonance condition the difference between the two beams must be chirped (Eq. 2.26), allowing the probing of the Doppler sensitive transition. To be resonant with one of the non-degenerate Raman transitions the frequency difference between the two Raman beams, $\nu_{\Delta R}(t, B)$, is

$$\nu_{\Delta R}(\Delta t, B) = \nu_0 + \alpha \Delta t + \Delta \nu_Z(B)\,. \tag{5.10}$$

Where Δt is the length of time the atom has been in free fall and $\Delta\nu_Z(B)$ is the frequency shift due to the local magnetic field B. The value of $\Delta\nu_Z(B)$ can be positive, negative or zero depending on the m_f state of interest. In Fig. 5.16 there are two large amplitude peak triplets that are symmetric about 0 MHz. These arise from the two possible choices of chirp rate as described in 5.3.1 with the left hand triplet

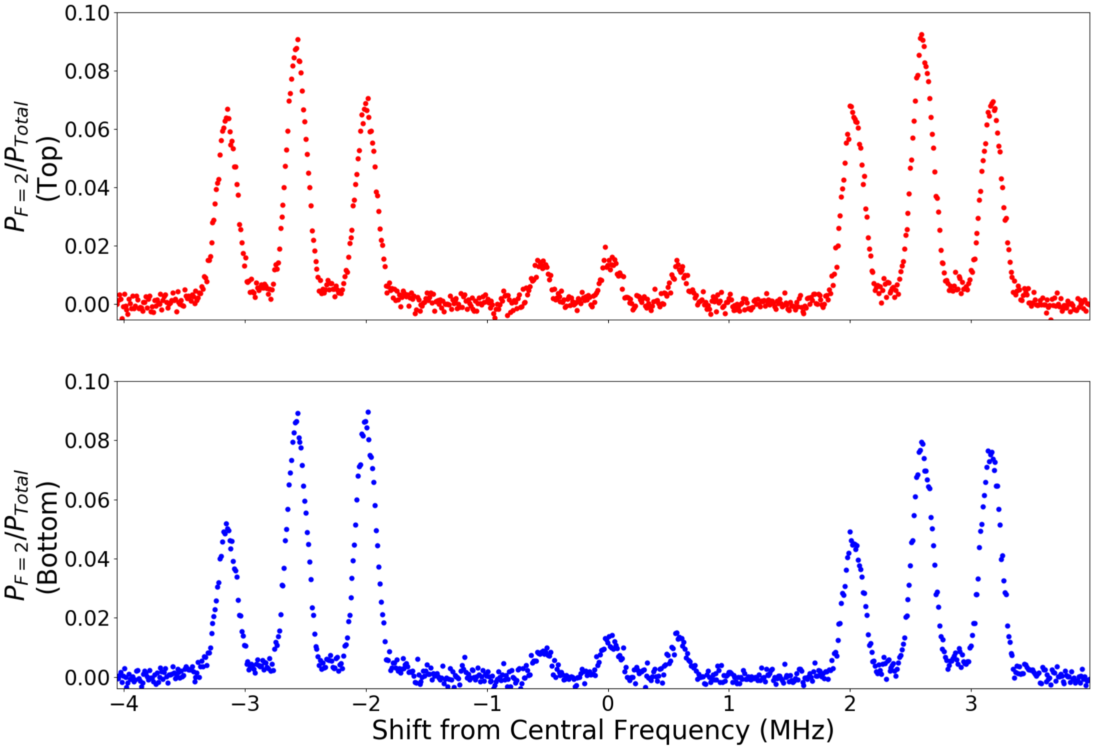

Figure 5.16: Scan of the Raman sideband frequency difference after ~100 ms of free fall. The larger sets of triplets on the far left/right extremes of the plot arise from probing the Zeeman sensitive transitions with counter-propagating linearly polarised light. The central triplet is the same Zeeman sensitive transitions probed with the small amount of Doppler insensitive circularly polarised light present in the input beams.

corresponding to $\alpha < 0$ and the right hand corresponding to $\alpha > 0$. The Zeeman splitting is applied on top of this effect resulting in the two triplets of possible transitions.

The central triplet arises from a Doppler free transition as its position does not change over the course of the drop. This arises from the co-propagating Raman transition where R_1 and R_2 are both delivered from the same telescope, the same Doppler shift has been applied to both frequencies and thus vanishes in the difference. In a well defined bias field this transition can only be driven by circularly polarised light present in the co-propagating beams, as the light in the telescopes is intended to be linear the presence of these peaks tells us that the Raman light in the chamber does not have a clean polarisation. The lower amplitude of these peaks is a result of the lower amount circularly polarised light, resulting in a slower

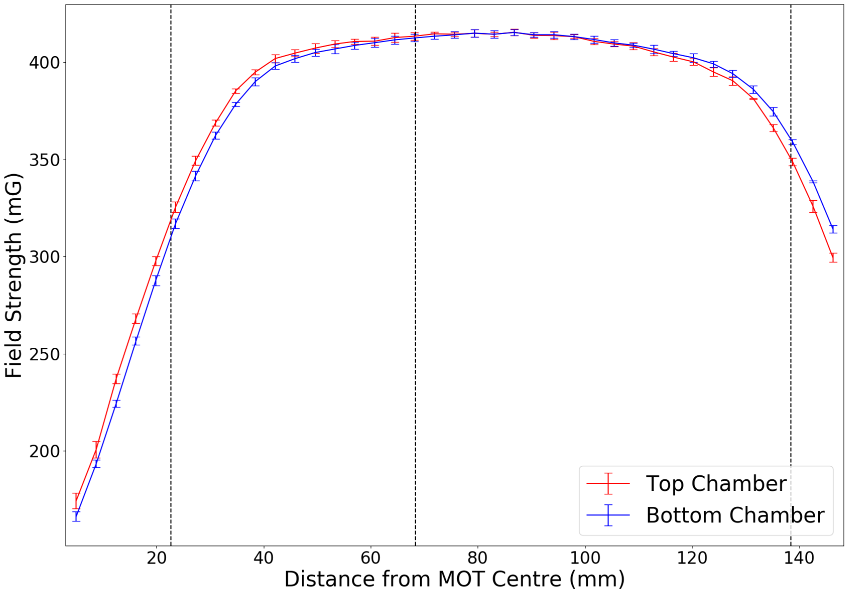

Figure 5.17: Magnetic field profile inside the science chamber as measured by the atoms, the black dashed lines represent the locations of the Raman pulses used in the current experimental sequence.

Rabi oscillation and lower transfer from a pulse of the same duration.

All three triplets have the same dependence on magnetic field as the Zeeman frequency shift is independent of light polarisation. This provides six measurements of the magnetically induced Zeeman shift from a single set of data. The results shown in Fig. 5.17 were found by taking the frequency difference between the $m_f = \pm 1$ peaks and the $m_f = 0$ peak of a given triplet and averaging, with the errorbar set by the standard deviation of these values.

From this it is now possible to assess whether the field generated here will be able to sufficiently separate the magnetically sensitive peaks. From the data in Fig. 5.17 it can be seen that the field strengths at the times of the pulses are 315 mG, 413 mG and 354 mG for the first, second and third pulses respectively. Modelling these fields in the same way as the results in Fig. 5.15 it is possible to determine the impact of external field noise on these measurements. Considering only the contribution from the $m_f = \pm 1$ peaks that falls within 3σ of the $m_f = 0$ peak frequency it can be found that the impact of the magnetically sensitive states in these fields is 0.243%, $\sim 1 \times 10^{-3}$% and 3.7×10^{-2}% respectively. The deviation arising from a change about the field strengths of ±5 mG is highest for the first pulse giving a difference on ∼ 0.11%. A π-pulse transfers roughly 22% of the atoms (Fig. 5.4) this magnetic sensitivity will not be the leading order noise on a measurement.

Figure 5.17 also shows that the 0.2% variation in field strength seen in the bias field when characterised off the system (Fig. 3.9) is increased to ∼ 3.5% over the same region. The non-finite cloud size means the result here is the average value in all three dimenions, but could also arise from the presence of additional magnetic field sources not present in the original measurement. The background field in Fig. 5.18 can be measured by reversing the direction of the bias field and repeating the same measurement as in Fig. 5.17, the resulting fields can then be subtracted. This is preferable to removing the bias field entirely as unless the background fields provided a frequency shift greater than 2σ (200 kHz corresponding to a field strength of ∼145 mG) it will not be able possible to distinguish the peaks.

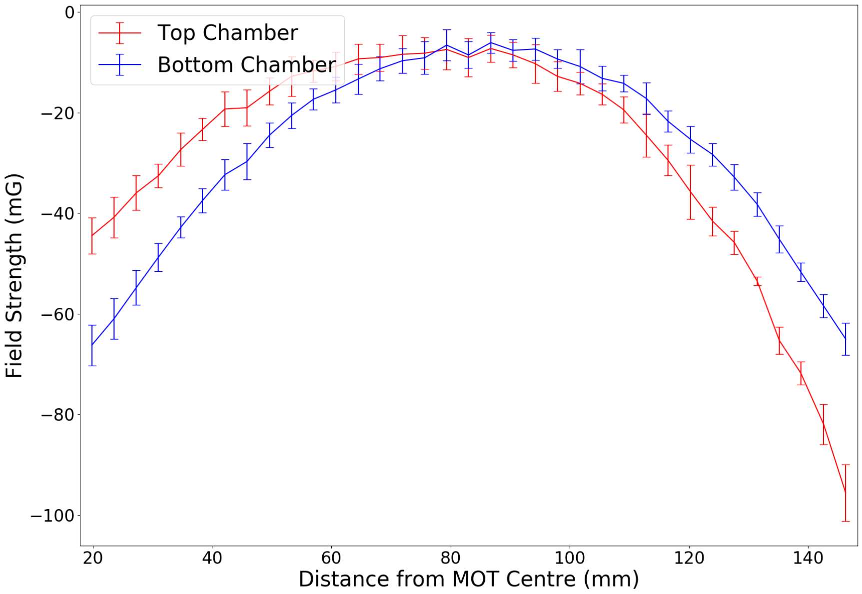

Figure 5.18: Background field profile.

Figure 5.18 shows a strong residual field with a fairly interesting shape. It is suspected that this is due to the magnetisation of the inner experimental shield, as it lies only 2 mm from the wire creating the bias field. The shape of the residual field matches that expected close to the edge of the bias coil given the higher concentration of windings near the edges that decreases near the centre. The difference in the two field profiles shown in Fig. 5.17 will present an offset to the measured gradient value but will be a constant systematic and as such will not present a problem.

5.4.2 Spatially Dependent Rabi Frequency

In addition to the varying magnetic field within the chamber there is a spatial effect that arises from the generation of unwanted sidebands in the Raman light (Fig. 5.12). Using frequency modulator such as an EOM or IQ modulator generates an infinite number of frequencies offset from the carrier by the modulation frequency ν_M such that

$$\nu_{R_i} = \nu_{R_1} + n\nu_M. \tag{5.11}$$

Where $n = \pm 1, \pm 2, \pm 3, ...$ is the modulation order of the ith sideband. It has been explained that the modulation frequency is set to keep the difference between R_1 and R_2 on resonance with the splitting of the hyperfine ground states of ^{87}Rb. As all sidebands have the same modulation frequency this results in an infinite number of frequency pairs that are resonant with this splitting, each ν_M further detuned than the last. In the case of weak modulation the majority of the optical power remains in the carrier and the primary sidebands, given high detuning and low optical power the higher order sidebands can be considered negligibly small.

It has been shown[35] [91] that all these pairs contribute to the effective Rabi frequency, in this weakly modulated case the effective Rabi frequency can therefore be assumed to arise from the carrier and primary sidebands only. This results in two frequency beatnotes that are resonant with the hyperfine splitting, the ν_{R_1}:$\nu_{R_1} + \nu_M$ pair and the ν_{R_1}:$\nu_{R_1} - \nu_M$ pair. These two beatnotes will have a different phase evolution in space meaning there will be a spatial dependency on the effective Rabi frequency[92]

$$\Omega_{eff}(l) = \frac{\Omega_{R_2}\Omega_{R_1}}{2\Delta_{R_1 R_2}} e^{ik_M l} + \frac{\Omega_{R_2'}\Omega_{R_1}}{2\Delta_{R_1 R_2'}} e^{-ik_M l}. \tag{5.12}$$

Where $k_M = 2\pi\nu_M/c$ is the spatial phase evolution of the modulation frequency and l is the distance between the cloud and the mirror. This spatial variation can be characterised by performing a Rabi oscillation measurement over the course of the drop, the resulting spatial variation is presented in Fig. 5.19. The finite spatial extent of an atom cloud has long been known of as a problem in interferometry, as the cloud spreads over the Gaussian intensity profile of a Raman beam there will be a large number of Rabi frequencies present. A vertical dependence on the measured Rabi frequency resulting from the second sideband introduces an additional component to this noise. The period of the spatial scillation is ~2.2 cm, the wavelength of the RF modulation frequency.

The cloud size overlaid on the results in Fig. 5.19 has assumed a starting diameter of 2 mm and expansion velocity in the vertical direction of 6.2 mm/s based on velocity selection by a 5 μs pulse, a typical $pi/2$-pulse duration in the experiment. Taking the most extreme instance of π-time variation across the cloud (the middle pulse) results in a difference in π-time of 1.1 μs for both chambers.

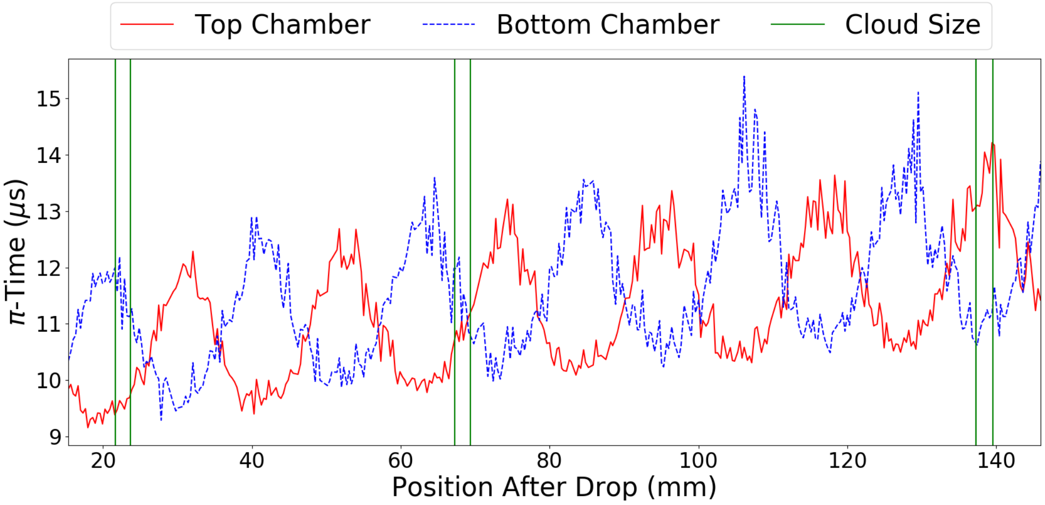

Figure 5.19: Rabi variation with height using an EOM, the cloud size is set by assuming an initial diameter of 2 mm and a velocity selective Raman pulse of 5 μs.

This spatial variation couples into an issue seen earlier in this chapter, the time-of-arrival variation. The 0.4 ms variation seen in Fig. 5.8 has been shown to be equivalent to a start position variation of 0.7 mm. Again considering the middle pulse, this noise of 0.7 mm would lead to a change in the π-time for the centre of the cloud of ~0.5 μs. This is a smaller effect than the 1.1 μs change due to the spatial extent of the cloud but unlike the gradient across the cloud it is not a systematic effect and will result in a phase noise on the gradient measurement. This issues is exacerbated by the fact that the oscillation is approximately π out of phase between the two clouds. Given that time-of-arrival variations are seen to be

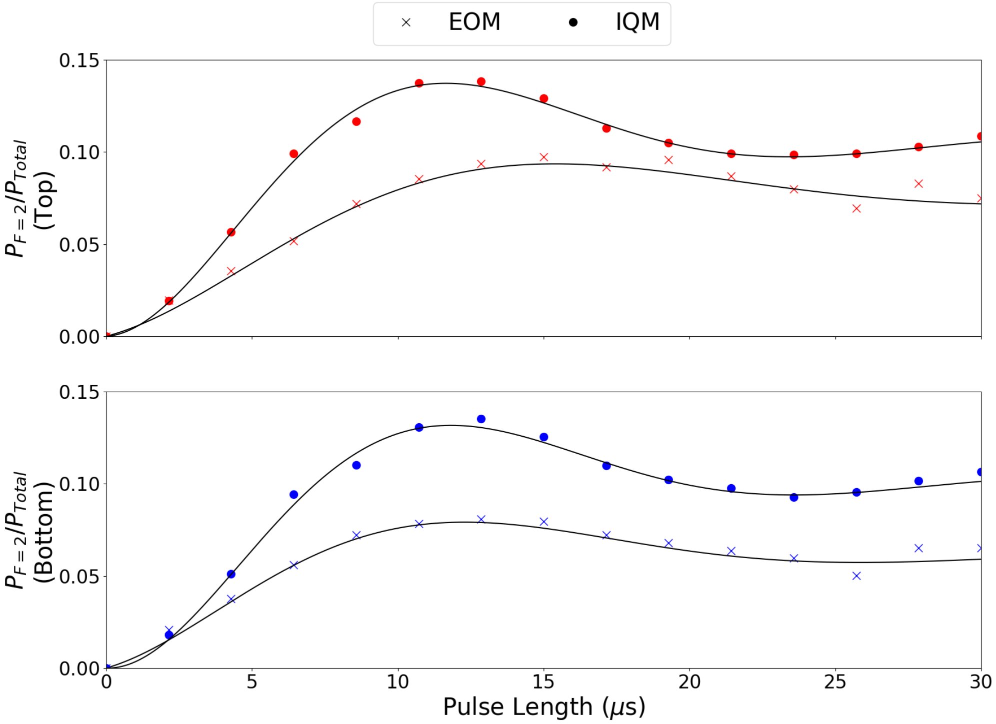

Figure 5.20: The results of Rabi oscillations for the two chambers first using and EOM to generate the Raman sideband and then and IQ modulator.

common (Fig. 5.8) the effect of this change on the relative π times for the two clouds will have opposite effects meaning it will not be removed by the common mode suppression.

From these considerations it is clear that there is much to gain by removing the parasitic transition through the use of an IQ modulator in the Raman laser, also giving arbitrary control of the sideband-to-carrier ratio. Figure 5.20 shows a comparison of two Rabi oscillations taken at the position of the final pulse, in one case using an EOM and the other and IQ modulator. There has been a clear gain in both amplitude and contrast due to the increased uniformity in the light seen by the cloud removing the less desirable Rabi frequency components.

Figure 5.21 has repeated the same Rabi frequency with height measurement now using the IQ modulator to produce only the desired sideband. Again considering the central pulse the new variation of π-time across the cloud has reduced to ~0.2 μs and the effect of the time-of-arrival noise has been reduced to ~0.1 μs. An interesting feature of this data is that as the atoms get lower in the chamber the

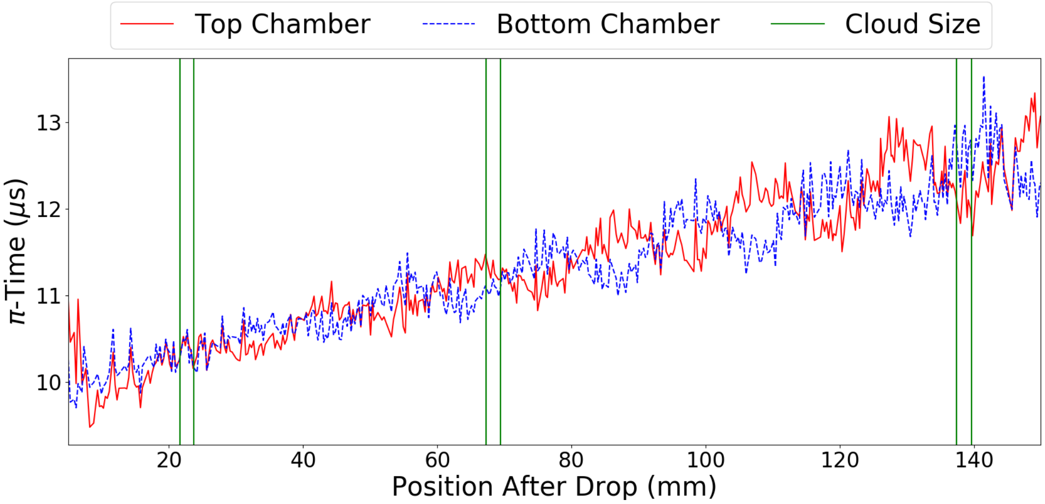

Figure 5.21: Rabi variation with height using IQ modulation and common input beams for both clouds.

effective suppression of the IQ appears to reduce. This is the manifestation of the effect seen in 5.11a, as the suppression changes with the driving frequency of the RF. As the atoms fall the modulation frequency is chirped to keep them on resonance so over the course of the drop the IQ modulator will be supplied with different driving frequencies. The unwanted sideband suppression was optimised for the time of the first pulse, high in the chamber, so the largest deviation from this optimised case would be expected at the bottom of the drop.

In this chapter certain characteristics of the system have been identified and presented along with their potential impacts. There has been some discussion of how changes made have either reduced instabilities or minimised the sensitivity of the system to those instabilities. In the next chapter the results of measurement campaigns in both real world and laboratory environments will be presented alongside the data processing used as part of their analysis.

CHAPTER 6

MEASUREMENT CAMPAIGNS

As outlined in 1.2 the overarching goal of this thesis was to create an atom interferometer that is highly portable whilst maintaining a useful level of sensitivity. To demonstrate this measurement campaigns were performed with anomalies of comparable signal size, one in a laboratory and one al fresco. This chapter will address the results and analysis involved in turning measurements from these campaigns into usable measurements of gravitationally induced phase. Once these have been considered the methodology and results from the two campaigns will be presented. The first of these was a three point measurement campaign that hoped to prove the durability and sensitivity in highly adverse conditions (see figures 6.11 and 6.10) by detecting an anomaly of 265 E. The second was a stationary measurement in a laboratory environment where a signal of ~280 E was modulated in time between three orientations to introduce a varying signal that could potentially be seen above phase drifts.

6.1 Systematic Offsets

When constructing an absolute gravity sensor, systematic offsets such as those mentioned in 5.3 will add an offset to the measured value of gravity. If we consider Eq. 2.27, the value of α is used to compensate for the Doppler shift experienced by the cloud as it accelerates in free fall. During a drop the cloud is decoupled from the mechanical components of the system and as such does not experience this vibration. However during the Raman pulse the effect of this vibration is imparted onto the cloud through the instantaneous velocity of the light reference, i.e. the mirror or the input telescope. This velocity acts to apply an additional Doppler component to the light seen by the atoms, moving it further from resonance.

It is impossible to separate the velocity component due to the motion of the mirror from the acceleration of the cloud induced by gravity so would look like an offset in the measured chirp rate. The measured chirp rate can therefore be broken down into a number of components

$$\alpha_{\text{measured}} = \alpha_{\text{gravity}} + \alpha_{\text{vibration}} + \alpha_{\text{other}}. \tag{6.1}$$

The α_{other} term is included to account for any sources of spurious chirp other than vibration, for instance if there is a strong magnetic gradient this will impart an acceleration onto the atoms. Most sources of acceleration will be constant, for instance the field gradient at the edge of the interferometry region as shown in Fig. 5.17 is a systematic offset so can be characterised and corrected. For vibration this will not be possible as it will vary over the course of a single shot and require compensation during a measurement, this is the reason most gravimeters include a substantial source of vibrational isolation or utilise long interrogation times, whilst it is intrinsically suppressed for gradiometers. Let us consider Eq. 5.7 for two gravity sensors where it has been assumed that $\alpha_{\text{vibration}} = \alpha_{\text{other}} = 0$

$$P^{T}_{|E,p+\hbar k_{\text{eff}}\rangle} = P^{T}_{off} - \frac{1}{2}\cos\left(\left(k_{\text{eff}} g - \alpha_{gravity}\right) T^2 - \phi_{\text{laser}}\right) \tag{6.2a}$$

$$P^{B}_{|E,p+\hbar k_{\text{eff}}\rangle} = P^{B}_{off} - \frac{1}{2}\cos\left(\left(k_{\text{eff}}\left(g + \delta g\right) - \alpha_{gravity}\right) T^2 - \phi_{\text{laser}}\right). \tag{6.2b}$$

The value of ϕ_{laser} will be the same for both clouds up to any delay introduced by the propagation time of the light. The factor of 1/2 has been absorbed into the offsets P^{T}_{off} and P^{B}_{off} which arise from experimental concerns such as decoherence effects and the presence of background atoms. The gravity difference term, δg, arises from the position dependence of gravity relative to its source, it is the difference in phase arising from this gradient that is being observed in an ellipse. Scanning the value of ϕ_{laser} from 0 to 2π will trace out a sinusoid for each sensor (referred to as a fringe) with some phase offset between them as shown in Fig. 6.1a. Creating a Lissajous figure with these two sinusoids results in a set of values that have an elliptical relationship; for discrete values of ϕ with uniform spacing they will correspond to points with uniform spacing on the ellipse (the white faced markers on Fig. 6.1c).

If some common phase noise is introduced onto both clouds, i.e. vibration, the effect on the fringe can quite clearly be devastating. The noise present in Fig. 6.1b is less than $2\pi/5$ (i.e. 20%) but the fringe is clearly indistinguishable in the data. However as this phase is common for both clouds this will act to move the points to some arbitrary phase whilst staying constrained to the ellipse, conserving the phase difference between the sensors. Common phase noise on the two clouds can be more than a single 2π and the common mode suppression will be maintained.

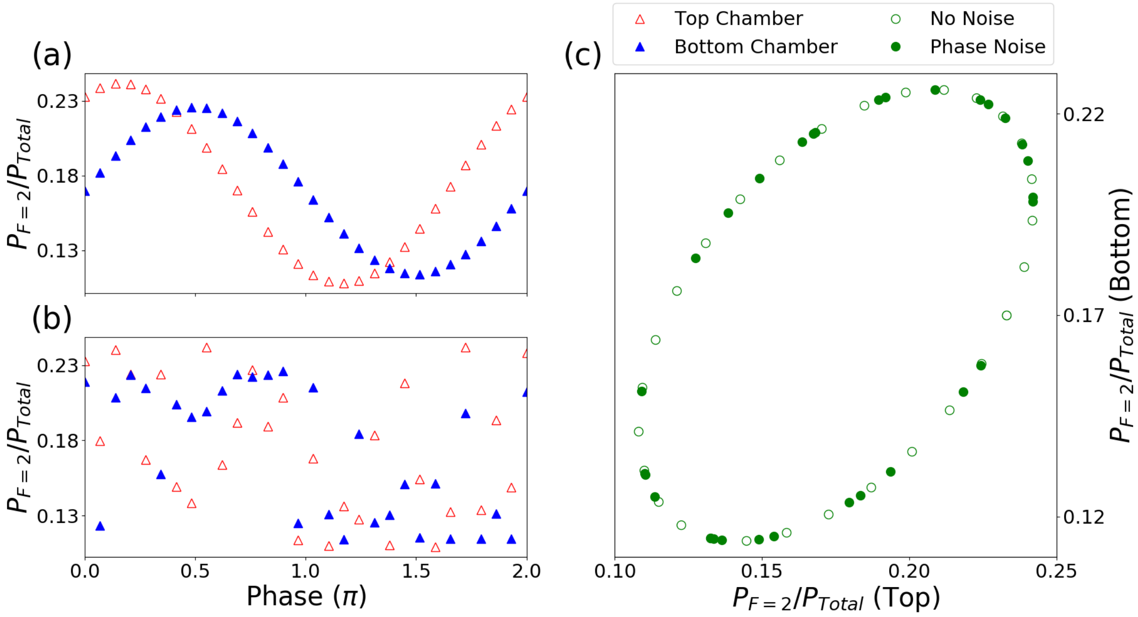

Figure 6.1: The evolution of the relative state populations with laser phase: 6.1a in the absence of any phase noise, 6.1b in the presense of 20% common phase noise. 6.1c the Lissajou figure produced by the two sinusoids demonstrating the effect of the common mode noise suppression.

For a gradiometer with the aim of performing surveys in unstable conditions, it is useful to know the exact value of α to provide the best ellipse contrast (i.e. the amplitude for the ellipse in the x or y directions). However as the aim of this experiment is not to give an absolute measurement of the gradient but to track local spatial changes, any systematic phase offset between the two clouds is unimportant as long as it is constant in time. The offset between the field profiles as seen in Fig. 5.17 or any deviation from the true value of α will manifest as a change to this phase offset, but not adversely affect the measurement. This will introduce, potentially large, arbitrary phase offsets between the two clouds that cause the ellipse to give an artificially small or large gradient value compared to the true Earth gradient, leading to a more “open" or “closed" ellipse. It has been found that a larger phase difference can help provide a lower error on the ellipse fitting as a too small phase difference causes the ellipses to tend to a line introducing an ambiguity as in which side of the ellipse should be associated with a given measurement point, a larger ellipse makes this determination easier.

6.2 Ellipse Analysis

A gradiometer performs a measurement simultaneously on two interferometers against some common reference, these measurements have a phase offset, here due to gravity gradients. By plotting the results parametrically any common phase shifts manifest only by moving a given point around an ellipse (see Fig. 6.1), conserving the phase relationship between the measurements. Spatial dependence of the phase offset ($\delta\phi_g = \delta g k_{eff} T^2$) is the quantity of interest when performing a gravity gradient survey. Tracking the changes in the value of $\delta\phi_g$ as the sensor moves over an anomaly it is possible to see changes this anomaly introduces to the gradient as $\delta\phi_g$ will change depending on the displacement between the sensor and the anomaly. Comparing different measurement points it is possible to remove any temporally stable offsets, like those mentioned in the previous section, as they will feature equivalently in all measurements.

6.2.1 Fitting the Ellipse

Acquiring a value for $\delta\phi_g$ is non-trivial, as the equation for an ellipse is non-linear in both x and y. For an ellipse with radius along the x-axis of A, radius along the y axis of B and centred at (x_0, y_0) the relation is

$$\frac{(x-x_0)^2}{A^2} + \frac{(y-y_0)^2}{B^2} = 1. \tag{6.3}$$

This makes using standard linear fitting algorithms difficult. The fitting algorithm used in this thesis is a method developed[2] utilizing Markov chain Monte Carlo algorithms to fit to the parameter space, which will be briefly explained here. For the ellipse data presented in this chapter the convention is to place the value of $P_{F=2}/P_{Total}$ for the Top sensor on the x axis and that for the Bottom sensor on the y. Instead of using Eq. 6.3 to define the relationship between the two parameters it is simpler to define x and y parametrically with respect to a variable θ

$$x = \frac{P^T_{F=2}}{P^T_{Total}} = A\sin(\theta) + x_0 \tag{6.4a}$$

$$y = \frac{P^B_{F=2}}{P^B_{Total}} = B\sin(\theta + \Delta\phi_g) + y_0. \tag{6.4b}$$

Fitting these equations still presents a problem for off the shelf fitting algorithms and as such a Bayesian Monte Carlo approach is used. The problem is now one of finding the set of values of parameters, $\boldsymbol{\gamma}$, that best match the data, $\mathbf{z}$. Contained within $\boldsymbol{\gamma}$ are all the free parameters of the system described in Eq. 6.4 such that

$$\boldsymbol{\gamma} = \left[A, B, x_0, y_0, \delta\phi_g\right] \tag{6.5}$$

All variables within $\boldsymbol{\gamma}$ have a range of possible values they can take. This parameter space is explored through $15,000$ iterations whereby a Bayesian method is used to iteratively move towards the most likely set of values for $\boldsymbol{\gamma}$. This is done by evaluating Bayes' theorem for each set of $\boldsymbol{\gamma}$ in turn[93]

$$p\left(\boldsymbol{\gamma}|\mathbf{z}\right) = \frac{p\left(\boldsymbol{\gamma}\right)p\left(\mathbf{z}|\boldsymbol{\gamma}\right)}{p\left(\mathbf{z}\right)}, \tag{6.6}$$

where $p\left(\boldsymbol{\gamma}|\mathbf{z}\right)$ is the probability that for some set of data $\mathbf{z}$ there will be a given set of values for $\boldsymbol{\gamma}$, similarly for $p\left(\mathbf{z}|\boldsymbol{\gamma}\right)$. The probability distributions of $\boldsymbol{\gamma}$ and $\mathbf{z}$ are $p\left(\boldsymbol{\gamma}\right)$ and $p\left(\mathbf{z}\right)$ respectively. The numerator of Eq. 6.6 gives you the likelihood that a given set of parameters $\boldsymbol{\gamma}$ fits the data and the denominator is the prior probability distribution of the data.

To find a good fit this calculation is iterated through many values of $\boldsymbol{\gamma}$ and compared to the current saved value. If the current data is labelled with $\boldsymbol{\gamma}$ and the comparison set with $\boldsymbol{\gamma}^{'}$ it is possible to determine which set best matches the data by taking the ratio of the two values of $p\left(\boldsymbol{\gamma}|\mathbf{z}\right)$

$$\frac{p\left(\boldsymbol{\gamma}|\mathbf{z}\right)}{p\left(\boldsymbol{\gamma}^{'}|\mathbf{z}\right)} = \frac{p\left(\boldsymbol{\gamma}\right)p\left(\mathbf{z}|\boldsymbol{\gamma}\right)}{p\left(\boldsymbol{\gamma}^{'}\right)p\left(\mathbf{z}|\boldsymbol{\gamma}^{'}\right)}. \tag{6.7}$$

If this ratio is greater than one then the current set of parameters, $\boldsymbol{\gamma}$, is a better fit to the data and if it is less than one the comparison set of data, $\boldsymbol{\gamma}^{'}$ is a better fit. At each comparison stage the best fit to data is saved and used as the reference for the next point then by selecting different sets of values from the known distribution of $\boldsymbol{\gamma}$ and repeating this process as many times as possible the values for $\boldsymbol{\gamma}$ will converge towards the most likely solution. This process also provides the way the error in the phase difference ($\delta\phi_g^{Err}$) is quantified. After a large number of iterations there should be a small, localised parameter space for the values of $\boldsymbol{\gamma}$ that will be explored in future iterations, meaning that there are many different possible values for all entries of $\boldsymbol{\gamma}$ with similar probabilities. All quoted ellipse parameters are the most likely value of γ (i.e. the “saved value" for the final iteration) and the error is the standard deviation in the parameter values for the last 25% of the iterations. For a standard run this is 3750 parameter choices. This is deemed as a representative measure of the error on the phase as it will reflect the range of possible values for $\delta\phi_g$ that could fit the data, with a more noisy data set having a large spread of possible values at similar probabilities.

Verifying Data Normalisation

This error determination method allows the verification of the normalisation technique outlined in the detection sequence (5.1.3) in the context of the impact it has on the phase error. Using data recorded over the course of almost 10 hours[1] the data was processed in two different ways before an ellipse was fit. One set, labelled as “Normalised", uses the same process of shot-to-shot normalisation as shown in Eq. 5.1. The other, labelled as “Avg. Normalised", was normalised to the average value of I_2 (see Fig. 5.6 for the definition) for bands of 90 measurements that are used to fit the ellipse, this scaled down the data to allow for an easier comparison but also removed any slower run-to-run variations.

Figure 6.2 plots the value of $\delta\phi_g^{Err}$ against time for the two approaches. It is clear here that not normalising the data shot-to-shot results in a much larger parameter space for $\boldsymbol{\gamma}$ and thus there must be a larger error on the result. The fact that the error on the “Avg. Normalised" tends towards the value of the “Normalised" data, suggests that the shot-to-shot variations become less prominent in the system over time. This is potentially due to the system reaching an equilibrium state as the dispensers thermalise after activation at the start of the day or the system reaches a stable rubidium pressure given the duty cycle of the measurements.

That the average normalised data has a higher error can be explained by observing the shot-to-shot data and fitted ellipses that result from the analysis as in Fig. 6.3. There is no clear ellipse shaping in

---

[1] 168 discrete ellipse runs of 90 shots each.

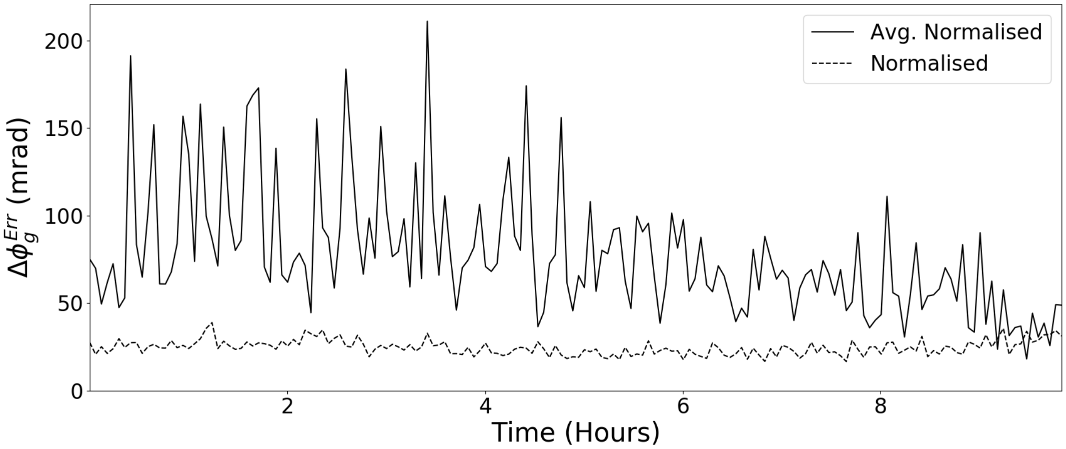

Figure 6.2: Effect of normalisation style on ellipses fitted to sets of 90 measurement shots over the course of 10 hours.

the average normalised data and there is a substantially higher spread, to be expected if there is a noise source that has been suppressed. The data normalised in the standard way has a clear shape that matches the fit well, making it clear by eye that normalising shot-to-shot reduces the noise level on the data. This validates both the normalisation method and the use of the standard deviation of the fit parameters as a reasonable metric for the error as it clearly increases for worse data.

6.2.2 Phase Corrections

In a gravity survey there will be spatially dependent variations on any given measurement that do not arise from the anomaly of interest. They arise from both large and small scale geography changes and should be constant in time for a given measurement point. There are a well established set of corrections for these effects in gravimetry, however they will need to be adapted for use by a gradiometer. Any given measurement of gravity g^M relative to some reference plane[1] is a combination of a number of correction terms and the “true" value of gravity g[94]

$$g^M = g + \delta g^L + \delta g^B + \delta g^{FA} + \delta g^{Ter}. \tag{6.8}$$

---

[1] In gravimetry this reference point is usually taken as the geoid. In a survey there is a freedom of choice for this reference to be placed at any point at which the sensor has performed a measurement.

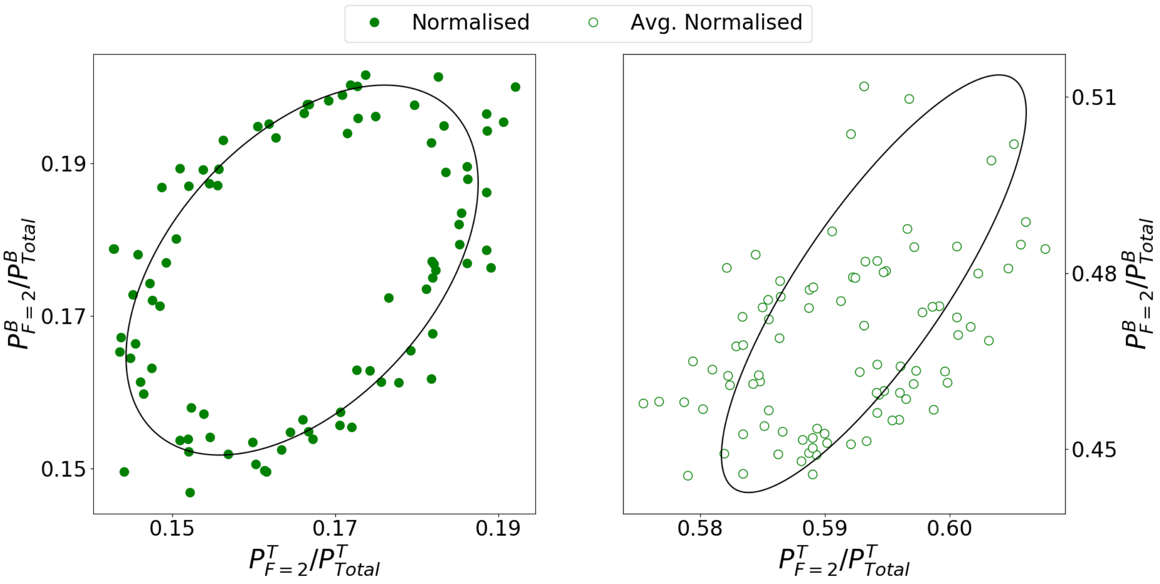

Figure 6.3: Left: Data normalised with Eq. 5.1 that shows a good correlation with the ellipse. Right: Data normalised to the average background for all the data points that shows a much larger spread and a clearly worse match to the data.

- **Latitude and Bouguer corrections:** The latitude correction, δg^L, accounts for the fact that the Earth is not a perfect sphere, with a larger radius near the equator. For a geoid measurement a correction is required to account for this bulging to ensure all measurements are made relative to the same reference radius. In gradiometery this correction will be identical for both sensors and so will subtract out. The Bouguer correction, δg^B, arises from large changes in local terrain, such as the presence of hills or valleys. This will also manifest near identically for both sensors so will subtract out. It is worth noting that if a local survey is performed rather than a reference geoid measurement these corrections would subtract out both in the Lissajous measurement and when comparing measurement points to see changes.

- **Free air correction:** δg^{FA} arises from a change in the vertical distance with respect to the anomaly due to changes in local topography, but does not account for any material that has been introduced into the space under the sensor. It is as if the sensor has been levitated and the void under filled with air (hence the name). The value of δg^{FA} can be found by differentiating Eq. 1.6 over space for a given change in height Δh

$$\delta g^{FA} = \frac{\partial g}{\partial r} dr = \frac{GM}{r^3} \Delta h. \tag{6.9}$$

For a gradiometer it is the difference in this value for the two sensors that is of a concern. As they have a different distance from the centre of mass (CoM) from the anomaly this change will manifest differently, the phase shift for each of the sensors is therefore

$$\phi_B^{FA} = k_{eff} \frac{GM}{{R_B^{FA}}^3} \Delta h T^2 \tag{6.10a}$$

$$\phi_T^{FA} = k_{eff} \frac{GM}{{R_T^{FA}}^3} \Delta h T^2 \tag{6.10b}$$

$$\delta\phi_{FA} = \phi_B^{FA} - \phi_T^{FA} = k_{eff} GM \left(\frac{1}{{R_B^{FA}}^3} - \frac{1}{{R_T^{FA}}^3} \right) \Delta h T^2. \tag{6.10c}$$

The distances R_B^{FA} and R_T^{FA} are the distance of each sensor from the centre of gravity (CoG) of the anomaly with this height change (for the Bottom and Top sensors respectively). Assuming[1] that this small change in height does not produce a substantial change to the values of R_B and R_T calculated in Fig. 6.5 then

$$\begin{aligned} R_B^{FA} &\approx R_B + \Delta h \qquad (6.11) \\ R_T^{FA} &\approx R_T + \Delta h. \end{aligned}$$

As the height of the sensor changes during a survey Eq. 6.10c will therefore result in a variation in the phase difference of the two sensors.

- **Terrain correction:** δg^{Ter} can be viewed as complimentary to the free air correction. The Terrain correction accounts for the new mass introduced to hold the sensor at the new elevation that was negated in the free air correction. If the distance from the bottom sensor to the CoG of the new

---

[1] A more rigorous treatment would include this analysis but if Δh is sufficiently small it should have minimal impact on the location of the CoG.

terrain is R^{Ter}, the average density of the material is ρ and the volume of material introduced is V then there will be an additional gravity contribution of

$$\delta g^{Ter} = \frac{G\rho V}{R^{Ter^2}}. \tag{6.12}$$

Applying this to two sensors separated by a baseline, b, the phase shift introduced by this change in mass will result in a differential phase shift

$$\phi_B^{Ter} = k_{eff} \frac{G\rho V}{R^{Ter^2}} T^2 \tag{6.13a}$$

$$\phi_T^{Ter} = k_{eff} \frac{G\rho V}{\left(R^{Ter} + b\right)^2} T^2 \tag{6.13b}$$

$$\delta\phi^{Ter} = k_{eff} G\rho V \left(\frac{1}{R^{Ter^2}} - \frac{1}{\left(R^{Ter} + b\right)^2} \right) T^2. \tag{6.13c}$$

It has been assumed that R^{Ter} is small such that the CoG lies at the same point as the CoM of the new material. If the sensor is lowered then the density term will be negative, as mass was removed, $\Delta\phi_{Ter}$ will therefore be signed depending only on change of mass.

- **Tilt noise:** As mentioned in 1.1.1 gravity sensors are sensitive to tilt as only the projection of gravity on the internal reference is measured. As expanded on in 1.2, this effect also manifests on gradiometers but is substantially suppressed. For a gravimeter this projection is

  $$g^M = g\cos\theta. \tag{6.14}$$

  As this is not constant in time but will vary as the external influences on the sensor change it cannot be characterised as a single offset. By monitoring the value of the tilt, θ, over the course of a measurement the effect on ellipse phase can be corrected. Consider that the phase measured by either of the two sensors is

  $$g_B^M = g\cos\theta + \delta g_B^{FA} + \delta g_B^{Ter} \tag{6.15a}$$

$$g_T^M = \left(g + \delta g\right)\cos\theta + \delta g_T^{FA} + \delta g_T^{Ter}. \tag{6.15b}$$

As the sensors will both have the same tilt applied and the true gravity value for the top sensor is just that of the bottom with some difference δg. Taking the difference of the Eq. 6.15a and 6.15b then converting to a phase gives the full correction term

$$\delta\phi_g = \frac{\delta\phi_g^M - \delta\phi^{FA} - \delta\phi^{Ter}}{\cos\theta}. \tag{6.16}$$

Where $\delta\phi_g^M$ is the result of the ellipse fitting algorithm and $\delta\phi_g$ is the true local gravity from an anomaly. This correction term will be used in the processing the survey phase difference data.

6.3 Tunnel Survey

In order to test the validity of the system as a portable sensor, field trials were required in an appropriate real world environment. A site was chosen located on the University of Birmingham campus where a near surface maintenance tunnel provided the perfect target to test the capabilities of the sensor. As will be outlined in this section this meant working in highly averse conditions with the sensor protected from the rain, wind and snow by nothing more than a gazebo[1] intended for summer use. Before the discussion of the result the site will be presented in greater detail and there will be some discussion of the methodology used over the course of the measurement campaign.

6.3.1 The Survey Site

In March of 2018 the performance of the system was such that detecting an anomaly seemed feasible, a survey location was identified near the Physics West workshop on the University of Birmingham campus. The potential anomaly (see Fig. 6.4) was a service tunnel that came within 30 cm of the surface, giving the subsurface void a vertical cross section of 2.11×2.61 m. The signal from this anomaly would depend on the density of the surrounding material, but making the reasonable assumption that it is concrete, then

---

[1] Tesco's finest *Outsunny 3m x 3m Pop Up Gazebo Water Resistant - Blue*

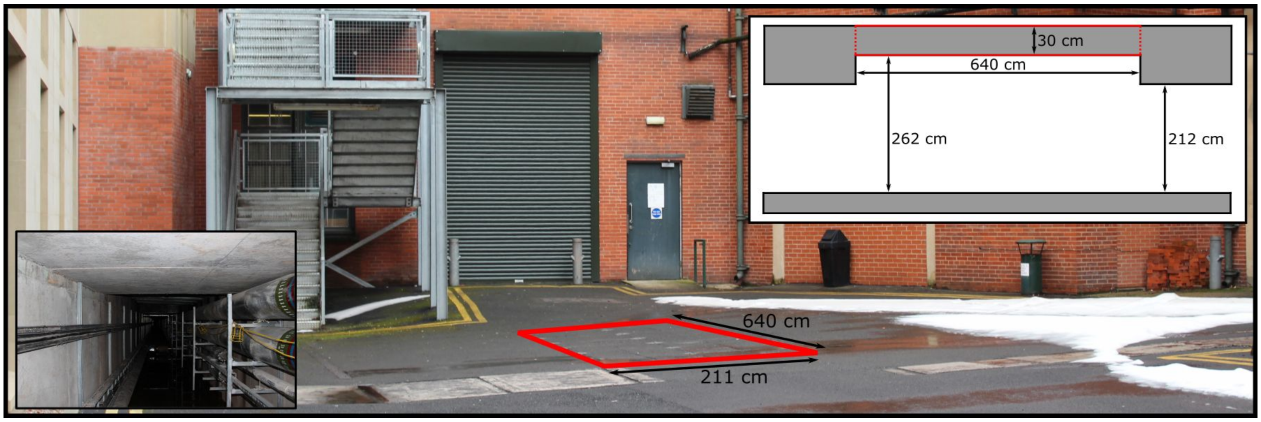

Figure 6.4: The site before setup on the first day of the survey. The red rectangle shows a region in which the distance to the surface drops to 30 cm, leading to an increase in the detected signal.

a 1 m slice would have a mass equivalent to more than 13 tonnes. This site had additional merit as there was good access to power and overnight storage facilities whilst being sufficiently close to the lab that transport and staff availability would not be an issue.

In advance of attempting the survey a calculation of the expected signal size was performed to ensure that the anomaly would present a large enough gradient to be detected by the sensor. The simple way to do this is to model the 13 tonnes as acting at the CoM of the void, the issues with this is that the distance from the sensor to the top of the anomaly is less than the height of the anomaly. Due to the inverse squared relationship of gravitational acceleration with distance this means that the CoM and CoG of the anomaly would not lie at the same point. To calculate the location of the CoG for each sensor a numerical simulation was performed. To do this we return to a similar case to that of 1.3 where the two sensors detect a gravitational acceleration given by

$$g_B = \frac{GM}{R_B^2} \tag{6.17a}$$

$$g_T = \frac{GM}{R_T^2} \tag{6.17b}$$

Each infinitesimal element of the anomaly will contribute its own infinitesimally small force to the total experienced by the atoms. If the density is assumed constant across the anomaly region, a good approximation given it is an air filled void, then the only varying contribution will arise from the distance

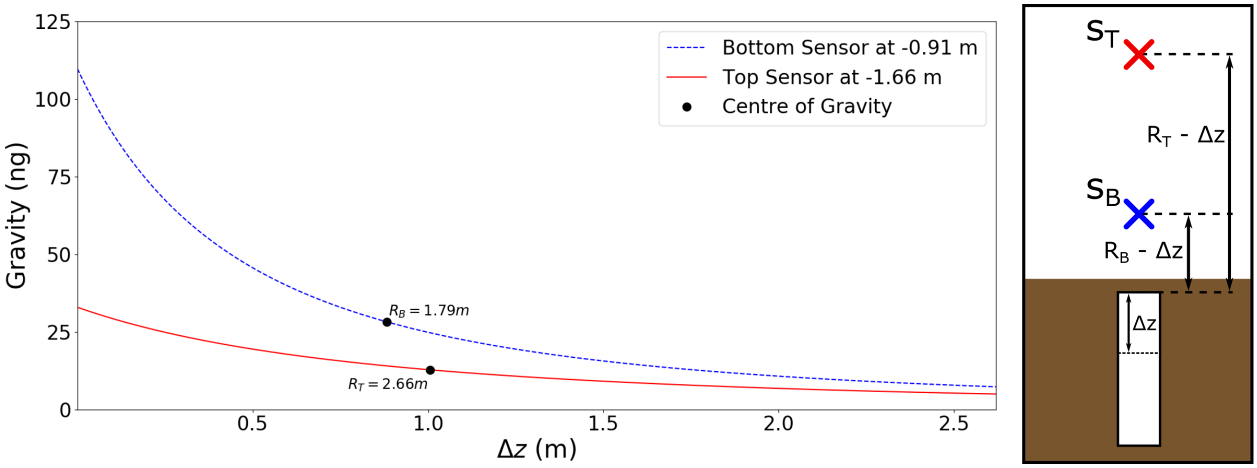

Figure 6.5: Given the physical parameters of the anomaly and those of the sensor, the CoG relative to each chamber is labelled after integration of the signal across the anomaly, the integrated values of Eq. 6.17a and Eq. 6.17b are the dashed (blue) and solid (red) lines respectively. The top of the anomaly is located at Δz =0 m. The values of R_B and R_T include systematic offsets of the system: the thickness of the tarmac above the anomaly, the distance from the ground to the bottom sensor and the separation between the two sensors.

term. For the purpose of this model the mass was set to be equivalent to a 1 m vertical slice of the void. By calculating the contribution of a test mass at each of these positions it is possible to determine where the CoG can be set for each of the two sensors. In Fig. 6.5 the Bottom sensor is offset from the top of the anomaly by a total 91 cm with 30 cm from the depth of the anomaly and 61 cm from the sensor geometry, the Top sensor head is offset from this by the baseline of 75 cm.

In Fig. 6.5 the sensors have been placed in the negative Δz region as per the diagram. By finding the average of the gravity signal across this region, the CoG will be located at the value of Δz at which matches this. Performing this analysis for each seosnr in turn gives R_B = 1.69 m R_T =2.58 m. These are offset from the CoM of the anomaly by 45 cm and 31 cm respectively. Given that the baseline of the sensor is 75 cm these are non-trivial offsets. Using these distances it is possible to model the 1 m slice as a point mass at R_B and R_T for the two sensors and use that to approximate the phase difference expected by the gradiometer.

In a survey the factor of interest is the difference in the average ground density between two points, rather than the true mass of the anomaly. To understand this consider the two cases represented in Fig. 6.6, in the left hand case a measurement is performed in the presence of an anomaly, the surrounding

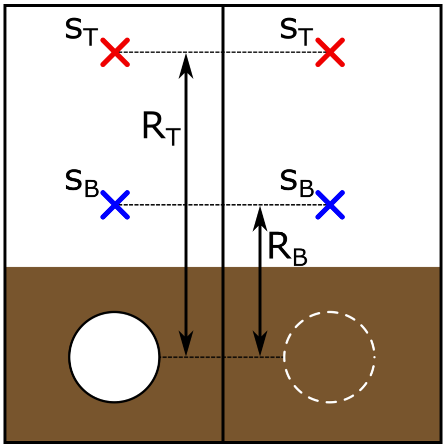

Figure 6.6: Sketch of the impact of taking a background measurement. Assuming everything else is constant it simply means there is a comparison between an anomaly with or without material present.

concrete therefore contributes to the acceleration measured and whilst the void contributes nothing. In the alternate case the void is filled with concrete so this provides an additional contribution to the measured acceleration that was not present in the original measurement. If the right hand case is considered zero and the left hand as a positive gravity value it is reasonable to see that the void could be defined as having a negative density that has the same magnitude as concrete (Eq. 6.18). Due to this symmetry it is easy to see how a survey could identify both regions of high density, such as buried resources, or regions of low density, such as pipes and tunnels.

$$\rho_{\text{void}} = -\rho_{\text{concrete}}. \tag{6.18}$$

Using Eq. 6.18 and the values of R_T and R_B from the model, it is possible to find an approximate phase change induced by the void. One must first find the difference between the signal measured by the two sensors as in 1.3

$$\delta g_{\text{void}} = g_B - g_T = GV\rho_{\text{void}}\left(\frac{1}{R_B^2} - \frac{1}{R_T^2}\right). \tag{6.19}$$

This is then subtracted from a background level

$$\delta g = \delta g_{\text{background}} - \delta g_{\text{void}} = GV\rho_{\text{concrete}}\left(\frac{1}{R_B^2} - \frac{1}{R_T^2}\right). \tag{6.20}$$

Where $\delta g_{\text{background}} = 0$ and the situation in Fig. 6.6 in the absence of a void is considered as $\delta g = 0$. The phase shift expected from the void relative to the background is therefore

$$\delta\phi_g = k_{eff}T^2GV\rho_{\text{concrete}}\left(\frac{1}{R_B^2} - \frac{1}{R_T^2}\right). \tag{6.21}$$

This gives a value of 6 mrad for the 1 m slice considered in this example, a more thorough treatment[1] of the system predicted a gradient value of ~8 mrad. In the weeks preceding the survey measurements were taken in order to characterise the state of the system and if the value of the phase error $\delta\phi_g^{Err}$ was sufficiently small to see the signal. The ellipse featured in Fig. 6.7 is representative of the state approximately two weeks before the survey. If the error was maintained at 27.86 mrad over the course of the survey and if the number of ellipses produced was N_E then with the $1/\sqrt{N_E}$ reduction in error it should be possible to see a 6 mrad signal after integrating over 22 ellipses, with a signal-to-noise ratio of 1. With 90 shots per ellipse and a 0.3 Hz measurement rate this should take approximately 110 minutes, this becomes 13 ellipses and 65 minutes for an 8 mrad signal.

The final verification before performing a live survey with the sensor was to perform a survey with a CG-5 gravimeter, verifying the simulated 8 mrad signal and provide information on the gravity profile of the area that may have not been present in this initial treatment. For instance the presence of an additional smaller void, a region of high density buried material or the proximity of buildings would all impact the true signal of the area. Given that the edges of the ~12 m long survey region were within 2 m of the buildings it was important to determine if these would introduce a contribution to the signal that has not been accounted for in the simulation.

One of the main drawbacks of the CG-5 is the presence of a drift with time as the mechanical stress on the spring causes it to gradually stretch. This results in a drift with both linear and non-linear components. To ensure an accurate measurement this is calibrated against a base station over the course of the measurements and allows for the drift to be removed. Figure 6.8 shows the outcome of the CG-5 survey

---

[1] For this treatment the anomaly was considered as a large number of small cubes and the same calculation as above was performed on a finer mesh.

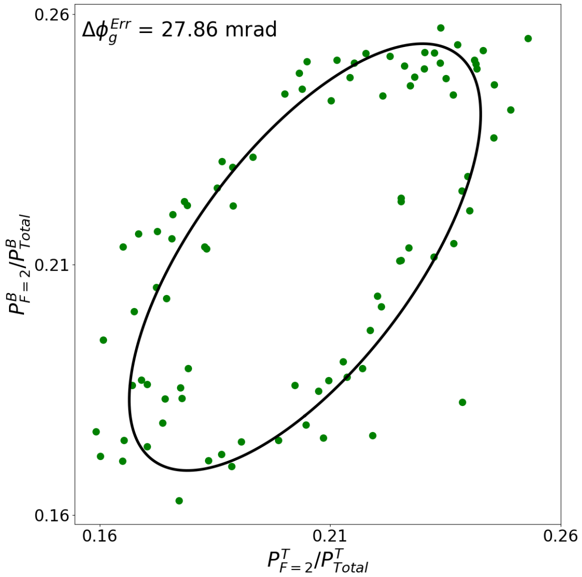

Figure 6.7: An ellipse and fit representative of the state of the system approximately two weeks before the survey.

with both the raw and corrected data. This provides a value of gravity over the survey region but will need to be converted into a gradient signal.

From Eq. 6.17 if one takes the ratio of two gravity values then what remains will be the ratio of the square of their distances to the anomaly. As such it is reasonable to say that

$$g_T = g_B \frac{R_B^2}{R_T^2} \tag{6.22}$$

Following the same analysis as in Fig. 6.5 with the CG-5 geometry the value of R_{CG5} =1.25 m can be found. Using this value with Eq. 6.22 and a gravity difference of ~65 μGal from Fig. 6.8 then it can be found that g_T = 14.4 μGal and g_B =31.7 μGal. Using these values with Eq. 6.21 then the phase shift expected due to the gravity gradient from the tunnel will be

$$\delta\phi_g = 6.98\,\text{mrad}. \tag{6.23}$$

Performing the same analysis for the smaller background signal seen at -4 m a phase difference of

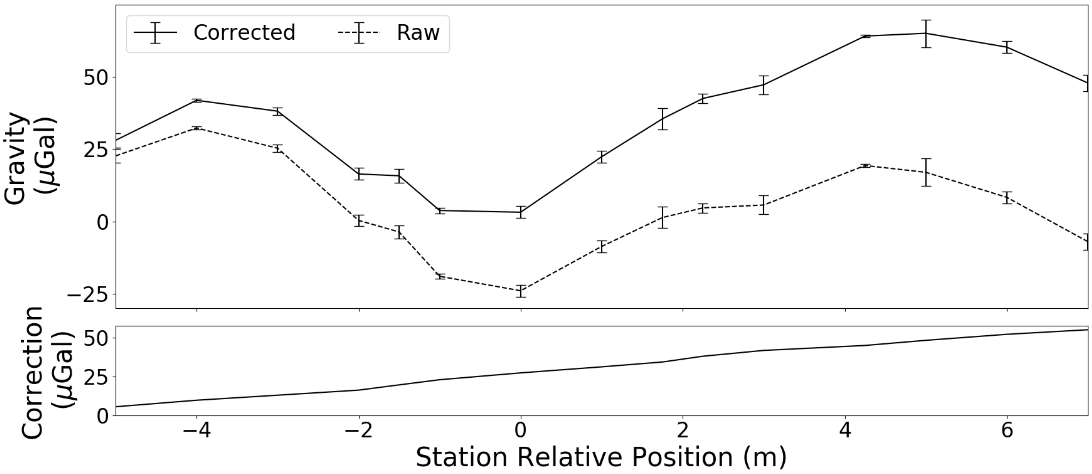

Figure 6.8: Results of a preliminary survey of the site performed with a Scintrex CG-5, 0 m denotes the location of the centre of the anomaly. The linear component correction is the drift in the CG-5 measurement as referenced against a reference base station.

$\delta\phi_g \approx 4.5$ mrad can be found. This may potentially arise from the proximity of the 4.5 mrad point to a building at the edge of the survey site, this would introduce an increased local mass above the sensor not present in the preliminary analysis. The result in Eq. 6.23 is similar to both the simple and more thorough models considered when analysing the survey site. If the system performance remains at that seen in Fig. 6.7 then the larger of these signals should be usable after averaging over the course of 16 ellipses at 90 shots each. However the smaller signal would require the averaging of 39 ellipses.

6.3.2 Methodology

An anomaly which is highly localised will present a good contrast with respect to the background points, in this case the simplest survey which will make the feature evident involves three measurements. One of these needs to be at the point of highest signal, i.e. directly above the anomaly, the other two need to be either side as far from it as possible to present the lowest signal. For the data presented in Fig. 6.8 this would result in a shift in the ellipse phase of ~7 mrad or ~4.5 mrad for the outermost points of the survey, both signals that should be detectable with minimal averaging time.

Due to the highly limited testing of the system in external environments it was unclear what level of

Figure 6.9: The author providing indisputable proof that the sensor is person portable if appropriate care is taken.

noise to expect by removing it from a laboratory and how this would impact the long term sensitivity. To mitigate the effect of this the survey was performed in two different ways over two different days to determine if either method produced a lower sensitivity to long term noise. Survey 1 consisted of four sets of 600 shots taken on each of the three measurement points, distributed randomly over the course of 10 hrs of measurement time, resulting in a total of 2400 individual shots that should present minimal bias with time. Survey 2 involved a single ~4 hr measurement on each of the survey points, performing 4200 shots, the increase in measurement data allowed for a longer integration period to potentially reduce the impact of noise. It was found in the analysis that due to slow drifts experienced it was impossible to see any discernible trend from Survey 2 hence only the results from Survey 1 will be presented.

For a portable sensor, even one that stands at ~1.7 m tall and weighs ~50 kg, the simplest way to move the sensor from point to point is to transport it by hand. The ultimate proof that the sensor is portable is that this can be done by a single person, as in Fig. 6.9, although it should be noted that training is required to do this safely. This also works as proof of the system robustness, with minimal finesse taken in moving and placement of the sensor it was sufficiently stable to be moved from point to point in a matter of minutes needing only re-levelling to put it back in a functional state. The levelling was performed manually using the fine machine threads of the tripod base of the sensor and a readout from the same tilt sensor used to determine the data in Fig. 6.12.

After every set of 600 shots (roughly every 30 minutes) readings of the optical power, the Raman sideband to carrier ratio (i.e. the amplitude of R_2/R_1) and the Rabi oscillations at the three pulse heights were taken. This allowed the sideband ratio and the unwanted sideband suppression to be kept reasonably constant over the course of all measurements, minimising the impact of the spatially dependent Rabi frequency (see 5.4.2). These steps were taken in an attempt to minimise the influence of some of the

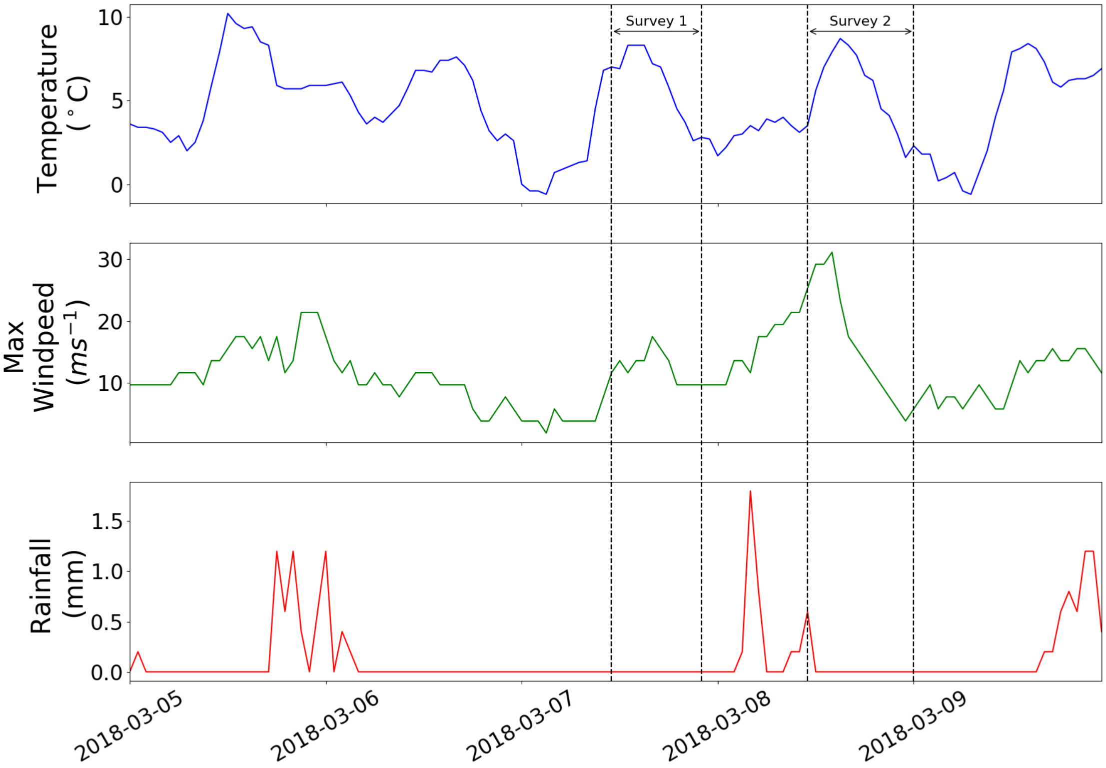

Figure 6.10: Weather data courtesy of the UK Met Office weather station at Coleshill.

noise sources that are present on the sensor both inside and outside the lab.

The conditions seen over the course of the survey were fairly adverse, taking place in the middle of “The Beast from the East"[1] one of the most severe and coldest storms experienced in the UK for decades, this would present a challenge even for well established systems such as the CG-5. Figure 6.10 shows the weather data over the survey period from the nearby Coleshill weather station supplied by the UK Met Office. Here it can be seen that over the course of the survey there was a substantial temperature variation from -0.5 ° to 10.2 ° as well as large spikes in wind speed. For a gravimeter these conditions would have been fairly devastating with the wind alone introducing sufficient tilt and vibration noise to invalidate the data.

The environmental conditions presented in Fig. 6.10 and Fig. 6.11 demonstrate just how far from

---

[1] `https://en.wikipedia.org/wiki/2018_Great_Britain_and_Ireland_cold_wave`

Figure 6.11: The gazebo which acted as home over the course of the survey. Through the sodden window it is possible to see the glow of the computer screen and in the background is what remained from the recent snow storm. A true “We’re not in Kansas anymore." moment.

ideal lab conditions this. With the temperature fluctuations introducing noise on power of the laser, frequency noise arising from temperature sensitive RF components and light intensity noise in the chamber arising from changes to the polarisation in the fibres. The impact of the rain was mitigated with the system’s magnetic shields and blackened plastic covers acting to make it splash proof, whilst the effect of heavy rain was removed with a $3 \times 3\,\mathrm{m}^2$ gazebo. The gazebo acted to provide some level of comfort for the survey team that was highly necessary on some survey days when the weather decided to be particularly unkind as is quite clear from Fig. 6.11. This was potentially the most important tool of the survey, other than the sensor itself, as in the winter of 2018 the UK experienced freak storms generating high winds, low temperatures and heavy rainfall.

6.3.3 Results

In all the survey period lasted 5 days between 5^{th} and 9^{th} March 2018. Initial teething problems arising from lack of survey experience meant that the first two days were spent characterising the system and planning the survey methods. The data of interest comes from Survey 1 on the 7^{th}, with the data from Survey 2 discarded as the noise on the measurement was seen to be worse using this measurement technique.

Rigidly mounted to the base of the sensor is a tilt meter that provided tilt data of the system relative to it's internal x and y axes (θ_x and θ_y). By treating the measurement plane of the sensor as rigid and perpendicular to the sensor itself it is trivial to show that the conversion to tilt of the sensor θ is

$$\theta = \arccos\left(\cos\theta_x + \cos\theta_y\right) \tag{6.24}$$

This is the θ that will be used in the correction given by Eq. 6.16. The tilt correction can only be applied post-process to fitted ellipse phase data with Eq. 6.16, meaning it is not possible to perform a shot-to-shot correction. Instead the average value of this tilt is applied as correction to the ellipse phase after fitting, this will result in a phase noise that cannot be corrected. To determine if this would be problematic, consider the maximum tilt recorded over the course Survey 1 of ~1.8 mrad. Equation 6.16 shows that this impacts the phase as $1/\cos\theta$, for the expected $\delta\phi_g$ =7 mrad this will introduce a shift of 11.3 nrad ($\sim 1\times10^{-4}$%) so will be negligibly small and ignored as a result.

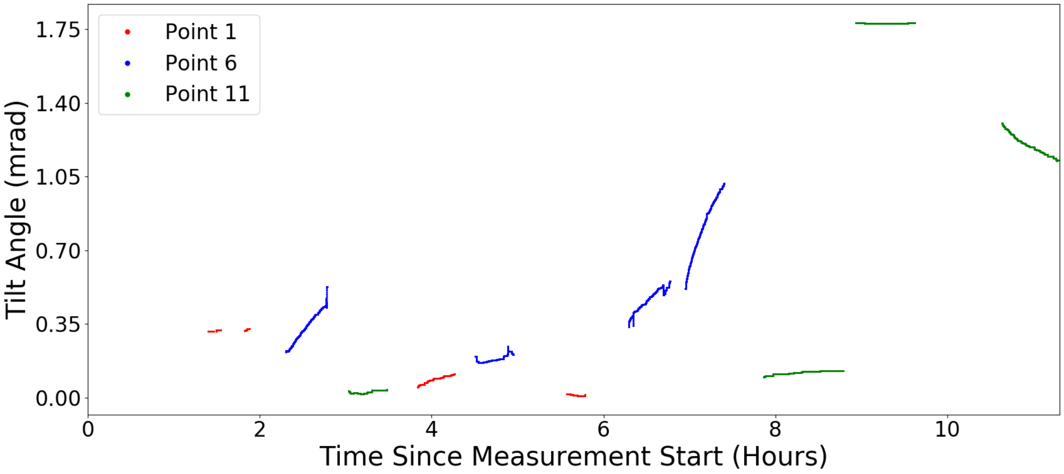

Figure 6.12: The tilt data from Survey 2, during the breaks in the data the sensor is being moved between measurement points.

Looking at individual ellipses (such as that in Fig. 6.13a) an interesting feature can be noted, the single ellipse performance appears to be better in survey conditions, with an almost twofold decrease in $\delta\phi_g^{Err}$ (as compared to Fig. 6.7). No quantifiable reason for this has ever been found, however it has been seen that many of the RF components in the flight case are sensitive to temperature variation, particularly during hot summer months. The flight case features 12 box mounted fans to remove heat from the box, however the airflow near the thermally sensitive components has been measured to be negligible. Figure 6.10 shows that temperatures during the survey were substantially lower than would ever be seen in the lab, the working theory is that this allowed for the temperature of the flight case to stay sufficiently cool for the components to operate more consistently.

Even with this improvement the ~14.2 mrad noise on the single ellipse fit too high to see the tunnel without integrating, theoretically requiring only 450 measurements (~25 minutes) to produce a signal-to-noise ratio of 1. Figure 6.13b combines all 2400 shots taken on the central data point. For this much data the error would be expected to reduce to 2.75 mrad but the result seen here is that expected from averaging over only 990 shots, roughly 2.5 times less data. This suggests that a drift is present on the data that is sufficiently slow it cannot be removed by integrating over such short time scales or it is a noise source that does not follow a normal distribution. The noise level of 4.28 mrad should still be sufficiently low that the 7 mrad and 4.5 mrad signals are visible with a signal-to-noise ratio of more than 1.

The analysis reveals a weakness of the fitting algorithms based on the least square of residuals when used to fit an ellipse, there is a heavy bias towards outliers. In Fig. 6.13b this manifests as a poor fit to the data at the thinnest portion of the ellipse, near its semi-minor axis, whereas near the upper right edge the fit appears to have been pulled out slightly to match the outermost points. Correlated noise on the sensor outputs that does not act on the phase would manifest by applying amplitude or offset noise to the ellipse, these would both act to make the ellipse more noisy along the semi-major axis. Offset noise would act to move the centre of the ellipse, introducing a spread on both edges equivalently, amplitude noise however would scale with signal amplitude, i.e. it would have a greater impact in the top right corner of the plot, this appears to be the case here.

For all three data points the 2400 shots were combined to give a single ellipse fit. Noise causing a spread of the data along the semi-major axis of the ellipse appears in all three data sets and a similar bias in the fitting algorithm appears. Applying the phase correction from Eq. 6.16 on the results of these fits provides the result seen in Fig. 6.14. A topographic scan revealed that point 1 was 15.8 cm below point 6 and point 11 was 18.5 cm below and the tilt term has been negated. In this data point 6 has been defined as the zero point for both elevation and phase.

The expected phase values from the preliminary CG-5 survey are marked on the graph for easy ref-

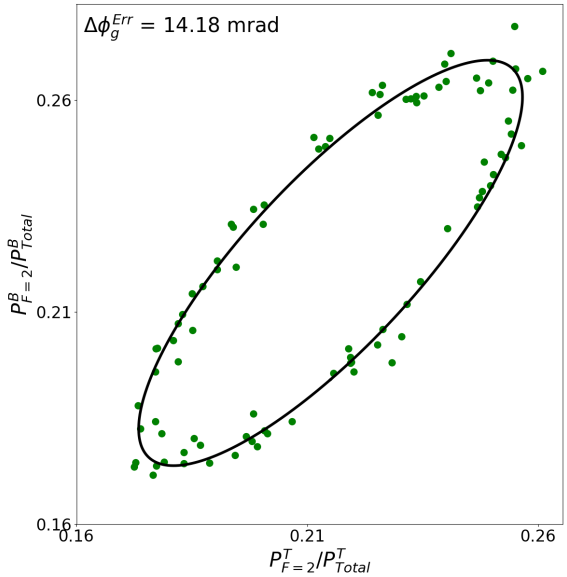

(a) A typical 90 shot ellipse from over the course of the survey, it is interesting to note that the error on the fit for this ellipse is lower than a comparative one from lab based measurements.

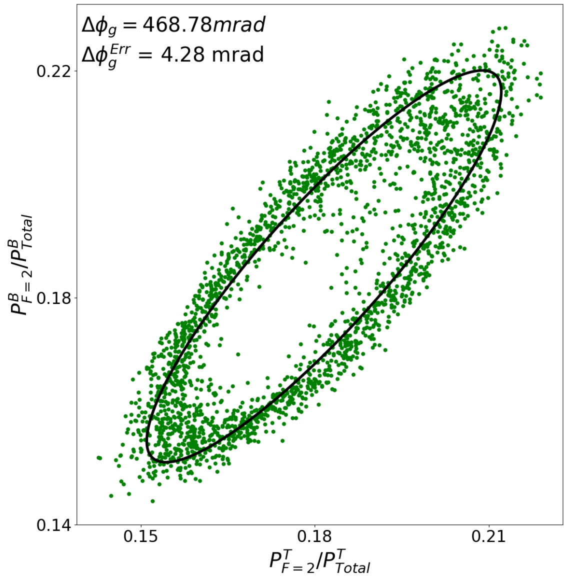

(b) The combination of almost all 2400 shots for the central survey point directly above the anomaly. Obvious outliers have been removed from the data to improve the fit of the ellipse.

Figure 6.13

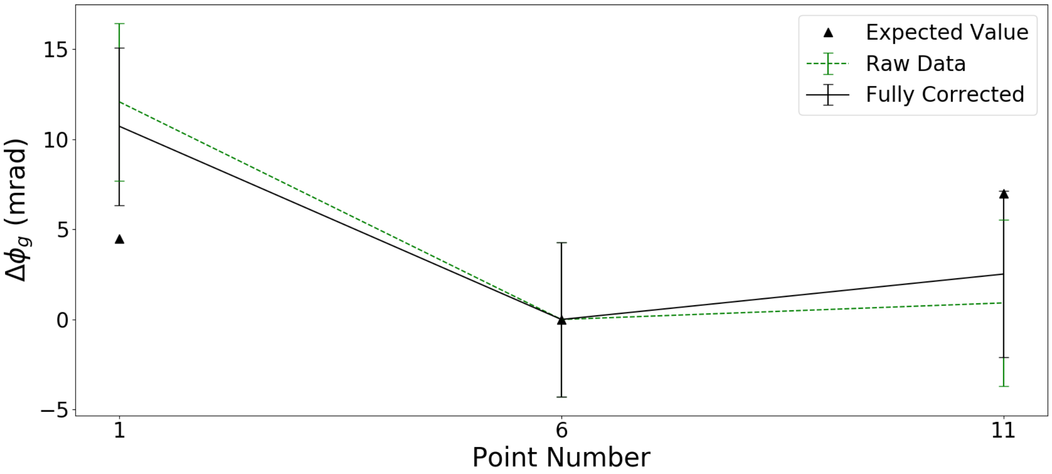

Figure 6.14: The phase data arising from Survey 1. To compare this result to Fig. 6.8 point 1 will be located at ~4.5 m, point 6 at ~0 m and point 11 at ~-4 m. The expected value is that determined using the CG-5 (Fig. 6.8) modified using Eq. 6.22.

erence. In spite of the substantial error on the signals there is a distinct difference between the phase as measured at the central point and that seen on the edge points. This data suggests that after having applied appropriate corrections to the data a trend is visible in the phase which is indicative of the expected signal from the tunnel, but does not allow a definitive claim of detection. This result shows that the sensor performed above expectations, as the thermal stability was a well known source of noise and the survey was performed before the introduction of Raman intensity stabilisation. From the same data we can see that the errorbars present are unlikely to improve with longer integration times as the improvement seen in Fig. 6.13 does not match that expected for the current integration times.

The way the drift and noise manifested provided a direction to focus the efforts and improve the sensor performance. The increased spread of the data points along the common mode axis (i.e. the $y = x$ line) of the ellipse suggests that the source of the noise is something which was common to both sensors but introduces minimal phase noise. Therefore the noise must arise from the light or timing of the system as these are the only two components of the sensor head that are common for both clouds. A potential culprit for this was the detection light, since the intensity was stabilised (Fig. 4.6) this suggests that it would be frequency noise of the cooling light introducing this drift.

Another possibility is a drift in the intensity of the Raman beam, given that the stabilisation as outlined in Chapter 4 had not yet been implemented. Raman intensity would be highly sensitive to temperature given the thermal sensitivity of light polarisation in fibres. Similarly temperature sensitive is the Raman sideband ratio as generated by the IQ modulator as the two RF signals were transmitted through 3 m cables to the modulator. At RF frequencies small changes to cable length or signal propagation speeds can lead to large phase drifts on the signal. If these were experienced differently by the two cables it would manifest as a phase drift on the signal due to the return of the position dependent Rabi frequency (more analysis is done on these effects in 6.4.1). The result of this survey gave reassurance that the sensor was both sensitive and robust enough to survive real field trials but suffers from a lack of stability that limits its ultimate performance.

6.4 Lead Detection

After improving sensor stability with intensity stabilisation on the Raman light, an increase in measurement rate and changes to the RF delivery the sensor was deemed ready for another survey. Given the potential thermal sensitivity it was decided to perform the survey with an anomaly of comparable magnitude but in a more stable laboratory environment.

The setup took advantage of the attractive nature of gravity; placing a mass centrally between the two clouds produces a larger gradient than the same mass placed directly below the system. This was achieved by placing two lead masses symmetrically about the central axis, mid-way between the two sensor heads, (Fig. 6.15). This produced an upward pull on the Bottom cloud and a downward pull on the Top cloud doubling the effective contribution from the lead. The symmetry of the setup also meant that off-axis components would be removed.

Each lead mass consisted of 20 $5\times10\times20$ cm lead bricks in stacks of 5 layers with 4 blocks per layer. The result was an anomaly of approximately solid lead with a total weight of ~225 kg and an identical setup on the other side of the system. The CoG of the lead was assumed as equivalent to the CoM of the block, a reasonable assumption as it will be equidistant from both clouds and thus any offset would be the same. It is possible to find that $x \approx 235$ mm, $R \approx 442.5$ mm and $\theta \approx 32.1\,^{\circ}$ given the known baseline of $b = 750$ mm.

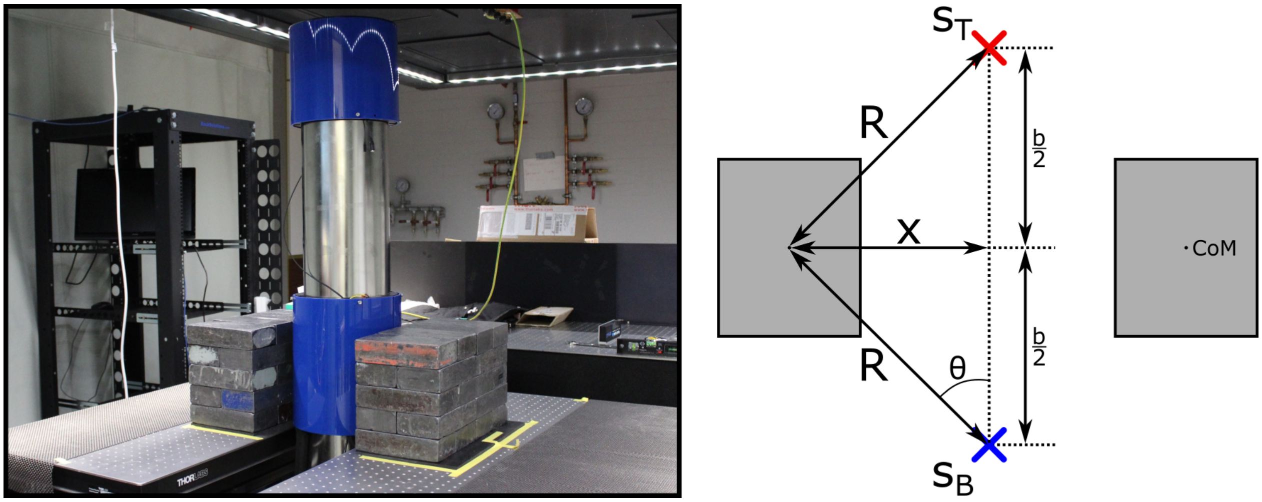

Figure 6.15: Left: The lead placed around the sensor in the “lead" position, the Top sensor is visible and the Bottom is below the optical tables. Right: Schematic of the experimental setup, placing lead at the midpoint of the two sensors it is possible to maximise the gradient signal by producing opposite phase shifts on the two clouds.

Using these parameters with Eq. 1.6 and correcting for theta with Eq. 6.16 the expected phase difference induced by the lead is

$$\delta\phi_{\text{Lead}} = 10.5\,\text{mrad}, \tag{6.25}$$

a signal which is approximately a 50% larger than the tunnel and measured in a more stable laboratory environment. Using 20 smaller lead blocks to construct the anomaly presented an option to modulate the signal to whatever level desired, that is to say anywhere between 0 and 20 blocks could be used. Three cases were considered and are presented in Fig.6.16. The first is the case of maximum signal with the lead placed as close to the sensor head as it is possible and will result in the phase difference in Eq. 6.25. The second is the case of minimum signal, the lead is moved as far from the sensor as possible whilst keeping it on the optical tables; this is the “null" case and will result in a phase shift on the sensors of

$$\delta\phi_{\text{Null}} = 2.1\,\text{mrad}. \tag{6.26}$$

From Eq. 6.25 and 6.26 it can be seen that $\delta\phi_{\text{Lead}} - \delta\phi_{\text{Null}} = 8.4\,\text{mrad}$, which is a signal slightly larger than that expected from the tunnel. The final orientation is a mid-point between the two extreme cases,

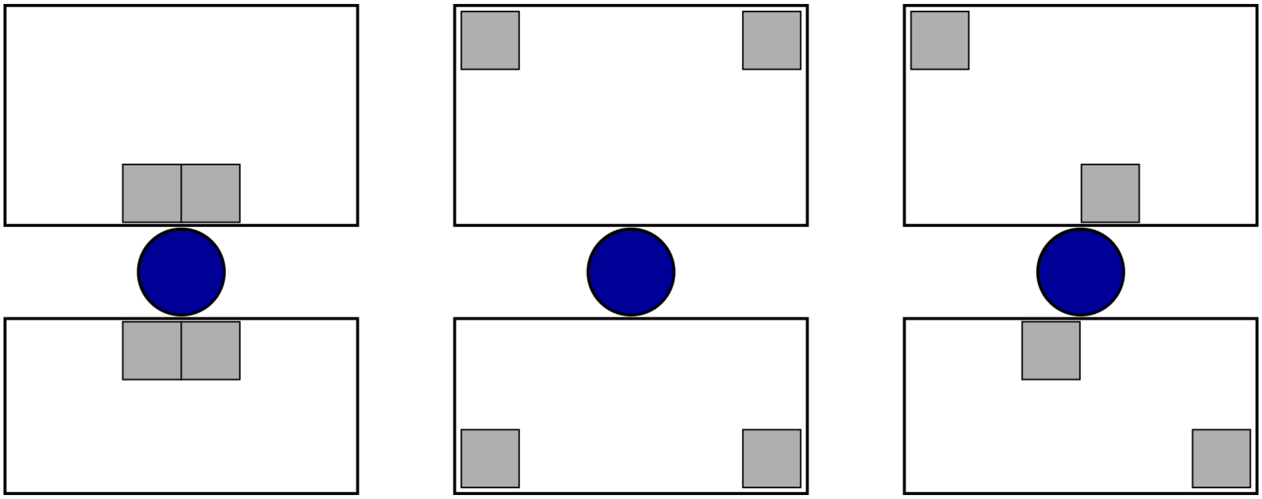

Figure 6.16: The three orientations of lead shown would have one of three effects (from left to right): maximise the gradient signal, minimise the gradient signal or produce half maximum gradient signal on the ellipse. By modulating between these three regularly and in a random order it would allow the effect of the lead to be seen above any drifts the system may experience.

half the lead is placed in the “lead" position and half in the “null" position (this mixed case is labelled as “half"). This should lead to a result that is the midway between those in Eq. 6.25 and 6.26 giving a value of

$$\delta\phi_{\mathrm{Half}} = 6.3\,\mathrm{mrad}. \tag{6.27}$$

Varying between these orientations in a semi-random fashion was hoped to serve the same purpose as randomising the data capture process during the tunnel survey. Taking many measurements in these different orientations in a random order throughout the day was hoped to make it possible to see jumps in the phase above any background drift.

6.4.1 Results

The approach to the stationary lead survey followed a similar approach as the more successful tunnel campaign; a measurement was taken in a given lead configuration for 1620 shots, a process which took approximately 1 hour. The state of the system was then changed and the measurement repeated. Long term measurements were taken overnight in the “half" and “null" configurations to allow for potentially long integration times. Limited access to the lead used to create the anomaly meant that it was not possible to perform overnight measurements for all three configurations. Figure 6.17 presents the result of the ~48 hour measurement campaign. Over the course of the survey there is a drift on the phase of the ellipses at of roughly 120 mrad, enough to dwarf the expected 8.5 mrad signal from the lead.

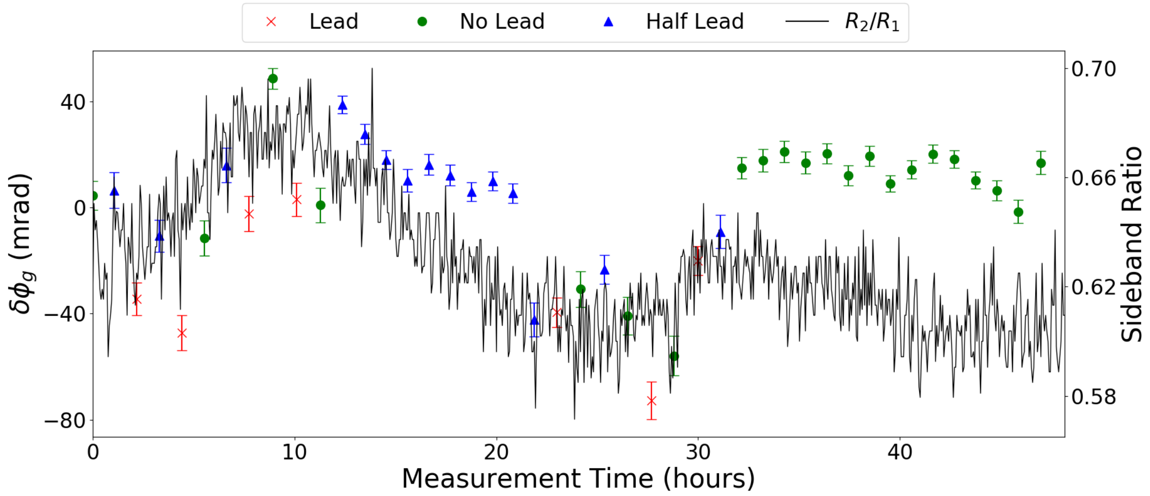

Figure 6.17: The result of the two day lead modulation survey. Each of the phase value is the result of 1620 measurements and the sideband ratio is recorded after every 90, the drift in the phase shows a good correlation with the sideband ratio.

Also shown on Fig. 6.17 is the primary Raman sideband ratio (R_2/R_1) which was monitored after every 90 shots. Given the ~ $\pi/2$ offset in the spatial variation of the Rabi frequency as seen in Fig. 5.19 it was conceivable that if the sideband suppression varied then this would induce a phase change on the ellipse. The data presented here shows a definite correlation between the sideband ratio and the ellipse phase. A measurement was performed by manually adjusting the sideband ratio and taking 90 shots (~4 minute) to determine if there is a direct causative relationship without longer term drifts impacting the result. Figure 6.18 shows that the sideband ratio has a direct impact on the phase of the ellipse. For each of these measurements the power in the Raman beam was seen to fluctuate by ~1 mW over the course of 270 measurements.

Given the impact the sideband ratio has on the phase of the ellipse it was important to identify the origin of this drift. Using the tilt-meter mounted to the sensor an ambient temperature reading was recorded over the course of the lead measurement campaign alongside the tilt data. Figure 6.19 shows

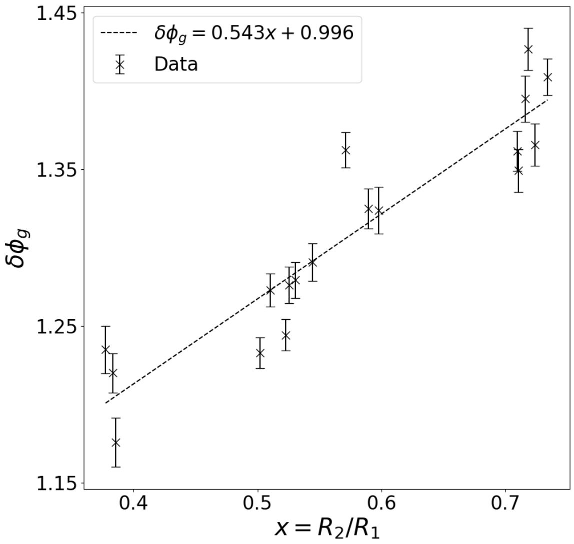

Figure 6.18: The sideband ratio was manually changed whilst maintaining constant total Raman power. A 90 shot ellipse measurement was performed at each value, the result is a clear correlation to which a linear relationship was fit.

this temperature variation overlayed with the IQ sideband ratio, showing another very strong correlation. This suggests that in spite of the more thermally stable laboratory the sensor is still highly sensitive to temperature drifts.

The likely source of this temperature sensitivity comes from the two RF inputs to the IQ modulator, they are generated on one side of the flight case and then connected to the other with two separate 3 m coaxial cables. This identified two potential sources of temperature sensitivity: the length of the cable can change as the conducting material expands or contracts with temperature and the dielectric constant of the PTFE insulator can change producing a substantial phase delay on the signal. It is fairly simple to model the magnitude of these two effects over the relevant temperature range considered here, this phase change can then be used to model the change in the sideband ratio with Eq. 4.4 from [32]. The relevant equations will be reproduced here for simplicity[1]

---

[1] In [32] the carrier is labelled as P_0 and the two sidebands as $P_{\pm1}$ where P_{+1} is the desired sideband. The convention in this thesis has been to use R_1 to refer to the carrier, R_2 to refer to the desired sideband and $R_2^{'}$ to refer to the undesired sideband. The thesis will remain internally consistent but its important to note the difference to avoid confusion.

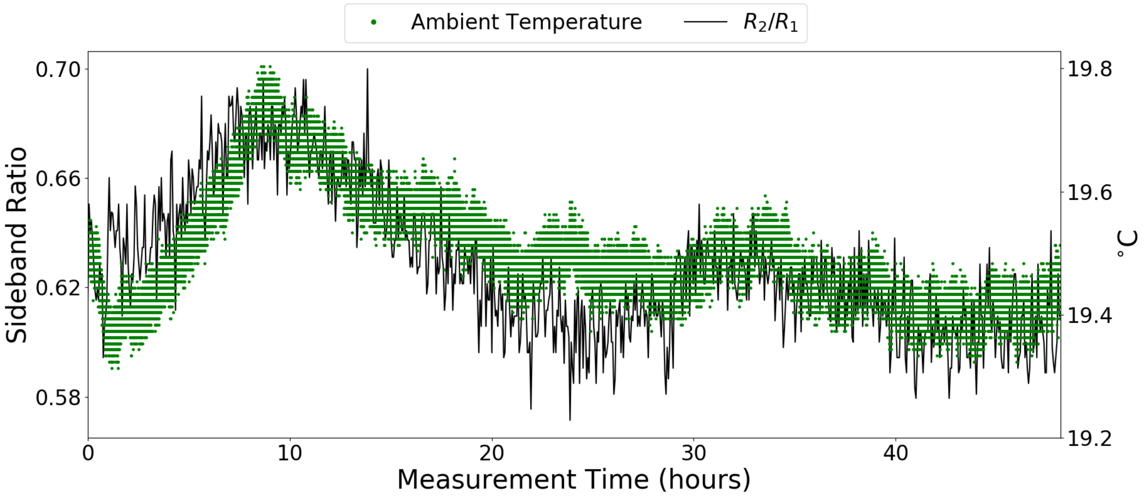

Figure 6.19: The change of primary sideband ratio against time as compared to the temperature variation measured at the sensor head. Discretisation of the temperature data is due to the limited resolution of the temperature sensor and makes it impossible to see any rapid variations. A correlation is clear despite some short term deviations around the 2 hour and 20 hour marks.

$$R_1 \propto J_0^2(\beta)\left(\cos^2\chi_1 + \cos^2\chi_2 + 2\cos\chi_1\cos\chi_2\cos\chi_3\right) \tag{6.28a}$$

$$R_2 \propto J_1^2(\beta)\left(\sin^2\chi_1 + \sin^2\chi_2 + 2\sin\chi_1\sin\chi_2\cos\left(\chi_3 - \chi_e\right)\right) \tag{6.28b}$$

$$R_2^{'} \propto J_1^2(\beta)\left(\sin^2\chi_1 + \sin^2\chi_2 + 2\sin\chi_1\sin\chi_2\cos\left(\chi_3 + \chi_e\right)\right). \tag{6.28c}$$

Where $J_n(\beta)$ is the n^{th} order Bessel function, β is the modulation index, $\chi_{1,2,3}$ are the voltage controlled phase differences of the internal interferometers and χ_e is the phase offset between the two input RF signals. The common leading order constants have been dropped as only the ratios are of interest.

A change to the primary sideband ratio will not directly impact the phase of the ellipse as this effect will manifest as a change to Rabi frequency common to both clouds (Eq. 5.5). However it does impact the spatial Rabi frequency, in Eq. 5.12 changing the value of $\Omega_{R_2}\Omega_{R_1}$ will also change the $\Omega_{R_2'}\Omega_{R_1}$ term will. This will lead to the spatial Rabi frequency varying with time, introducing a phase dependent component.

Phase noise induced by changing cable length can be determined by modelling the cable as a solid piece of copper, with a thermal expansion coefficient of $\kappa_{\mathrm{TM}} = 16.5\,\mu\mathrm{m}/\mathrm{m\,K}$. The phase change on the RF induced by the change of cable length, χ_{length}, can be calculated

$$\chi_{\mathrm{length}}\left(T_K\right) = \frac{2\pi\kappa_{\mathrm{TM}} L_{\mathrm{cable}} \Delta T_K}{\lambda_{\mathrm{RF}}} \tag{6.29}$$

For a cable of total length L_{cable}, a temperature change ΔT_K and an RF wavelength in the cable of $\lambda_{\mathrm{RF}} = c_{\mathrm{copper}}/\nu_{\mathrm{RF}} \approx 0.71c/\nu_{\mathrm{RF}}$. This allows the change in phase over the full temperature range to be approximated, a 0.5 K change as seen in Fig. 6.19 will result in a length change of $24.75\,\mu$m. For a signal with $\nu_{\mathrm{RF}} = 6.83$ GHz this shift can be maximised by assuming it only applies to one of the cables resulting in a 4.7 mrad phase shift.

The phase delay in coaxial cables as induced by temperature change has been presented in [95]; the paper shows that different cable designs can have vastly different effects on their phase delay they introduce. Figure 3(a) and 4(a) of [95] present the response for SS402 cable which has a similar design to the *MIL-DTL-17* cable used in the system, a tape based PTFE insulator and is manufactured by the same company[1] so will have a similar response. Over the 18-20 °C temperature range it is shown that there is a phase delay coefficient of $\kappa_{\mathrm{PD}} \sim 1\,\mathrm{ps}/\mathrm{m\,K}$. This produces a phase shift due to the change in dielectric constant

$$\chi_{\mathrm{PD}}\left(T_K\right) = 2\pi\kappa_{\mathrm{PD}} L_{\mathrm{cable}} \Delta T_K \nu_{\mathrm{RF}} \tag{6.30}$$

Applying this over the maximum temperature range of 0.5 K produces a range of induced phase delays of 31.6 mrad.

Assuming that the only time dependence arises from the temperature dependence on the RF signals due to the above effects it is possible to determine the temperature dependence of the sidebands. Taking the ratios of the Eq. 6.28b and Eq. 6.28c with respect to the Eq. 6.28a the following relations can be found for the temperature dependent sideband ratios

$$\frac{R_2}{R_1} = \frac{J_1^2(\beta)}{J_0^2(\beta)} \tan^2\chi \frac{\left(1+\cos\left(\chi_3 - \left(\pi/2 + \delta\chi_e(t)\right)\right)\right)}{\left(1+\cos\chi_3\right)} \tag{6.31a}$$

---

[1] *Harbour Industries*

$$\frac{R_2^{'}}{R_1} = \frac{J_1^2(\beta)}{J_0^2(\beta)} \tan^2\chi \frac{\left(1+\cos\left(\chi_3+\left(\pi/2+\delta\chi_e(t)\right)\right)\right)}{\left(1+\cos\chi_3\right)}. \tag{6.31b}$$

Given that there is freedom of choice in the phase values the simplification has been made that $\chi_1 = \chi_2 = \chi$. The phase offset between the RF inputs has been modelled as a constant shift of $\pi/2$ with a time dependent modulation

$$\chi_e(t) = \frac{\pi}{2} + \delta\chi_e(t), \tag{6.32}$$

where it is assumed that $\delta\chi_e(t=0) = 0$. With these conditions and the knowledge from measurements that $\frac{R_2}{R_1}(t=0) = 0.615$ and $\frac{R_2^{'}}{R_1}(t=0) = 0.061$ the values of the free parameters were found numerically to be

$$\begin{aligned} \beta &= 0.552 \\ \chi &= 4.056\,\mathrm{rad} \\ \chi_3 &= 2.206\,\mathrm{rad}. \end{aligned} \tag{6.33}$$

These match the initial conditions with a deviation of less than 1%. Taking these as constant in time it is now possible to determine the effect the temperature variation has on the sideband ratios if it only impacts $\delta\chi_e$, producing the result in Fig. 6.20. When compared to Fig. 6.19 it is clear that the simulation closely follows the ambient temperature variation and correlates with the measured sideband ratio.

These results reinforce the belief that phase noise in the delivery cables will introduce a thermal sensitivity on the sideband ratio, which scales linearly with cable length. The order of magnitude difference of the two effects means that this is not the leading order contribution to the variation, it must arise from an effect not considered in this analysis. The obvious source for would be the values given by Eq. 6.33 which have been assumed constant over the course of the measurement. The two phase values χ_3 and χ (or χ_1 and χ_2 in a more complex model) are set with voltage outputs of a benchtop power supply[1]. This power supply is quoted as having a voltage ripple of 1.5 mW RMS, a setting accuracy of 0.05% +5 mV and

---

[1] Rhode & Schwarz HMP2030

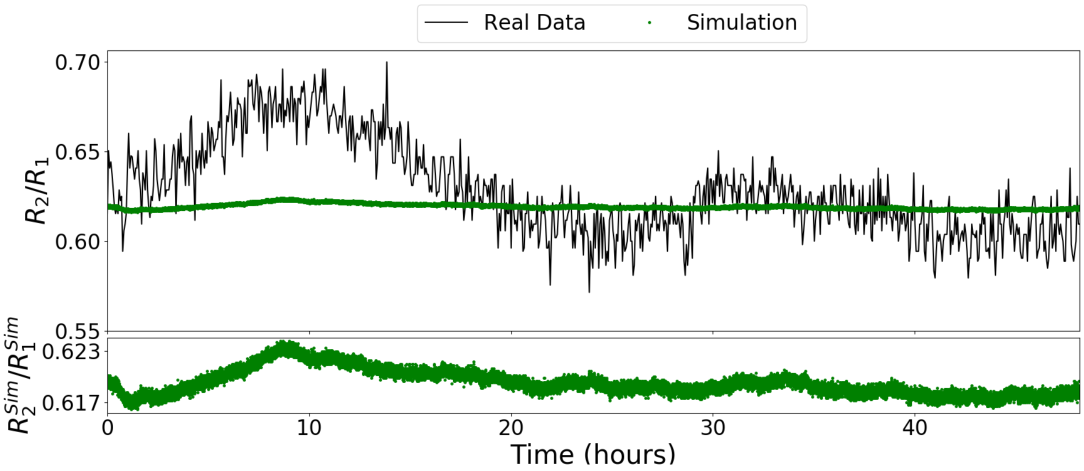

Figure 6.20: The variation of the primary sideband ratio over the course of the 48 hour lead survey. The simulated effect is clearly much smaller (roughly an order of magnitude) than the measured impact, however it can be seen that it follows the temperature variation closely. The lower plot has reduced the y-range of the other to better display the variation in the simulation.

a thermal voltage coefficient of 0.01% + 2 mV/K[96]. For a 0.5 K temperature variation and given an output of ~6.5 V[1] on the three channels this can mean a noise upwards of 12.4 mV on the output of the supply. It is difficult to simulate the effect of this voltage noise on the IQ phase without a far more in depth analysis, however during the tunnel survey changes of tens of mV were used to maintain the sideband ratio so it is feasible that this voltage noise is the problem.

The sideband ratio has also been seen to change with the RF input power, historically this change was used to correct for the change in sideband magnitude with frequency as seen in Fig. 5.11a. The temperature variation can potentially introduce a variable attenuation into the RF delivery chain by changing the efficiency of the amplifiers, the attenuation of the cables or couplers, then this would manifest by changing the value of the modulation index β. It is also possible that the Lithium-Niobate crystal in the IQ modulator itself introduces the temperature sensitivity itself. Finally, the temperature in Eq. 6.19 is the ambient temperature measured at the sensor head, however all these effects would manifest inside the flight case. If the change in ambient temperature was caused by changes in the heat output of the flight case then it is highly likely that the variation would be far higher internally. If this difference were an order of magnitude along the path of the RF cables then this would make the simulated IQ variation

---

[1] The highest voltage output from the power supply for χ_3.

almost identical to the measured.

Here the results and analysis of two sets of survey data, utilising a portable quantum sensor, have been presented. The first measurement campaign placed the sensor in a relevant but highly unstable environment and through manual monitoring and tuning over the course of nearly 12 hours results indicative of the expected signal are visible (Fig. 6.14). The second campaign placed the same sensor in a more stable laboratory environment to detect an artificially created anomaly of comparative magnitude to that of the external survey. Due to long term drifts and substantial temperature sensitivity of the IQ modulator used to create the sideband Raman frequency it is not possible to identify the signal. The potential cause of the drifts has been suggested and analysis performed that identifies a component, but the leading order issue has not yet been identified. These results strongly suggest that with improvements to system stability and reduction in thermal sensitivity the system will be capable of detecting more challenging targets. The next survey is currently planned for December 2018 with the hopes of detecting an anomaly on an active building site.

CHAPTER 7

CONCLUSION AND OUTLOOK

In this thesis the design and operation of a prism MOT gradiometer has been outlined from the initial system prototypes to a survey ready operational sensor. At each stage of development the suggested design has been tested and validated, showing that the single beam delivery system is capable of producing atom numbers in excess of 2.7×10^8 and reaching temperatures below 3 μK, when paired with a compact laser system based on telecoms technology. Single beam MOTs have been shown to provide substantial improvements to the system stability and have been integrated into a design that provides a compact cylindrical form factor for the vacuum system, providing a substantial simplification to system geometry and reduction of system mass. The low footprint design was modified to allow the cylindrical form factor to be implemented in a gravimeter system, which was used to demonstrate comparative performance in both a laboratory environment and outside in direct summer sunlight. The signal-to-noise of the fringes changed from 1.7 to 0.9 from the move outside, demonstrating that with improvements to control measures and utilising the common mode suppression of a gradiometer the design had potential as a survey ready device.

Working from this well validated position the system design has been outlined in detail to showing an efficient implementation of a single beam MOT in a gradiometer orientation which is capable of functioning as a research tool whilst remaining a system that can be manufactured en masse. The sensor features two modular vacuum chambers independently capable of maintaining a pressure an order of magnitude below that required for a $2T = 176$ ms gravimeter, a compact bias coil of only 18 turns that can provide less than 0.2% variation in magnetic field strength over the central 8.5 cm interferometry region and an adaptable mounting structure providing an adjustable 75 cm baseline. Once implemented into the gradiometer it has been shown that atom numbers of $\sim 5 \times 10^7$ and temperatures of ~7.8 μK can be seen in both sensor modules, with the drop in performance relative to the gravimeter likely arising insufficient laser available laser power. The same vacuum infrastructure has been prepared to allow the sensitivity to be quadrupled by implementing an atom launch with only minimal changes needed to the outermost layer of the sensor. Through use of these innovative designs this has all been implemented in a sensor with a total weight less than 50 kg that has been demonstrated to be moveable by a single person.

Pairing this robust sensor design with a portable control package it has been possible to perform surveys in highly adverse conditions. Filtering, actively stabilising and manually removing possible noise source a survey was performed during the month long stormy period in the UK dubbed “The Beast from the East", with a total temperature range of more than 11 °C and maximum wind speed of 30 ms experienced over the course of the survey. Temperature variation of almost 7 °C and wind speeds of more than 10 m/s were experienced over the course of the successful survey with the sensor detecting a result indicative of the expected 230 E signal. Having proved the sensor could provide good results in highly adverse conditions a second measurement campaign was attempted in a more stable environment with an artificially created anomaly of ~280 E. This campaign produced a less promising result as manual correction of the most problematic noise source was no longer provided and the sideband ratio from the IQ modulator was allowed to change by almost 20% over the two day measurement campaign. A clear correlation was seen between the temperature variation, the sideband ratio and the measured ellipse phase with a linear relationship identified between the sideband ratio and ellipse phase. The results of this campaign were inconclusive despite the improved environmental stability. Work is ongoing both on the sensor to prepare it for a new campaign outside the lab in December 2018 and on the data available from the current campaigns as given the clear correlation between sideband ratio, temperature and ellipse phase it is possible that the drift can be removed in post processing if a causative link can be determined.

With the success of the design presented in this thesis large industrial partners have begun to utilise the intellectual property[31] generated to develop their own prototypes for commercial applications. Teledyne e2V have been working with the local group for a number of years to develop their own gradiometer prototype, constantly learning from the progress of this system and have begun the process of building more by harnessing the same core design in two recently funded projects (REVEAL and FREEZE RAY 2). RSK is heading up a consortium, in which the University of Birmingham is the largest academic partner, to develop a separate system with a similar design to demonstrate an improvement in the field over the CG-5 (PIONEER). Additionally the original prism MOT prototypes are in constant demand from industrial partners to act as a test bed for their own technology, with one currently in use by Gooch & Housego to test their laser systems and the other used locally as a teaching tool, demonstrating they are functional for both academic and industrial work.

Work is continuing at the University of Birmingham on this system and those inspired by it. The MOT prototypes are a standard teaching tool and the gravimeter prototype has recently been rejuvenated as part of a new project which has it mounted to a rail. The two new systems mentioned in 1.2 are in the construction phase using modified designs of the components presented in Chapters 3 and 4, exploring the two extremes of the SWAP:sensitivity scale. Both expected to produce their first measurements by March 2019.

The output of this thesis has more than exceeded the expectations that were laid out at its commencement, whilst still in its early days it has proven the potential of the system as a whole to function as a gradiometer in adverse conditions. The system is acting as a testbed for new ideas, models and techniques to advance the readiness of cold atom sensors and has demonstrated to companies the potential for investments, harnessing the design for commercial applications.

APPENDIX A

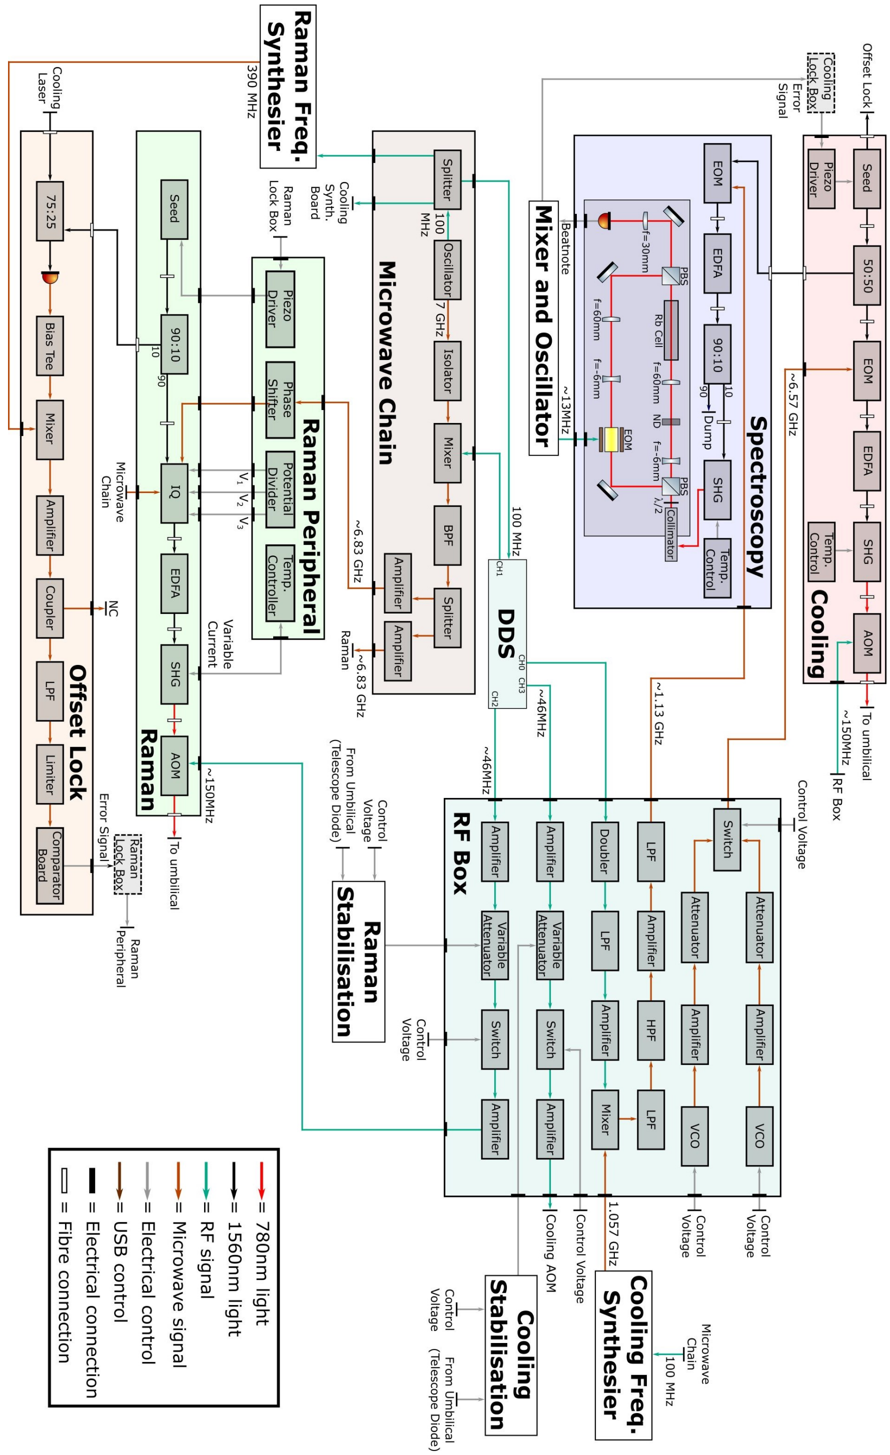

Figure 1: Complete breakdown of the RF and light distribution system from generation till exiting the flight case.

BIBLIOGRAPHY

[1] N. Metja et al. “Seeing through the Ground: The Potential of Gravity Gradient as a Complementary Technology”. In: *Advances in Civil Engineering* 2011 (2011), p. 9.

[2] A. Rodgers. “Bayesian inference from terrestrial gravimetry measurements of near-surface anomalies using a bespoke reversible-jump Markov chain Monte Carlo algorithm.” PhD thesis. University of Birmingham, 2016.

[3] E. Parseliunas, P. Petroskevicius, and Obuchovski R. “Analysis of Gravimetric Observations Made by Scintrex CG-5”. In: *The 7th International Conference on Environmental Engineering, Vilnius Gedimins Technical University (1422-1428)*. 2008.

[4] E. Parseliunas et al. “Investigation of the Automatic Gravimeters Scintrex CG-5 and Analysis of Gravimetric Measurements”. In: *The 8th International Conference on Environmental Engineering, Vilnius Gedimins Technical University (1422-1428)*. 2011.

[5] L. LaCoste. “The zeroâĂŘlength spring gravity meter”. In: *The Leading Edge* 7.7 (1988), pp. 20–21. URL: `https://doi.org/10.1190/1.1439525`.

[6] W.J. Hinze, R.R.B von Frese, and A.H. Saad. *Gravity and Magnetic Exploration: Principles, Practices and Applications*. Cambridge University Press, 2013.

[7] M. Hauth et al. “Atom interferometry for absolute measurements of local gravity”. In: *Proceedings of International School of Physics “Enrico Fermi” Course 188 “Atom Interferometry”*. 2014.

[8] Scintrex. *CG-5 Autograv Operation Manual*. 2010.

[9] D. Boddice et al. “A novel approach to reduce environmental noise in microgravity measurements using a Scintrex CG5”. In: *Journal of Applied Geophysics* 152 (2018), pp. 221 –235. ISSN: 0926-9851. URL: `http://www.sciencedirect.com/science/article/pii/S092698511730215X`.

[10] GWR Instruments. *Operating Principles of the Superconducting Gravity Meter.* 2011.

[11] Muquans. `https://www.muquans.com/index.php/products/aqg`. Accessed: 2018-06-08.

[12] V. Ménoret et al. “Gravity measurements below 10^{-9}g with a transportable absolute quantum gravimeter”. In: *Scientific Reports* 8.12300 (2018).

[13] AOSense. `https://aosense.com/solutions/gravimeter/`. Accessed: 2018-06-08.

[14] A. Peters, K.Y. Chung, and S. Chu. “High-precision gravity measurements using atom interferometry”. In: *Metrologia* 38.1 (2001), p. 25. URL: `http://stacks.iop.org/0026-1394/38/i=1/a=4`.

[15] I.M. Longman. “Formulas for computing the tidal accelerations due to the moon and the sun”. In: *Journal of Geophysical Research* 64.12 (1959), pp. 2351–2355. URL: `https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/JZ064i012p02351`.

[16] M. Feinberg. *Complemented absolute/relative full-tensor gravity gradiometer system.* Patent No. 6658935 (2003-12-09).

[17] D. DiFrancesco et al. “Gravity gradiometer systems âĂŞ advances and challenges”. In: *Geophysical Prospecting* 57.4 (2009), pp. 615–623. URL: `https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1365-2478.2008.00764.x`.

[18] A. Sugarbaker. “Atom Interferometry in a 10m Fountain”. PhD thesis. Stanford University, 2014.

[19] C. Mahadeswaraswamy. “Atom Interferometric Gravity Gradiometer Disturbance Compensation and Mobile Gradiometry”. PhD thesis. Stanford University, 2009.

[20] X. Wu. “Gravity Gradient Survey with a Mobile Atom Interferometer”. PhD thesis. Stanford University, 2009.

[21] J.M. McGuirk. “High Precison Absolute Gravity Gradiometer with Atom INterferometry”. PhD thesis. Stanford University, 2001.

[22] P. Marte, P. Zoller, and J.L. Hall. “Coherent atomic mirrors and beam splitters by adiabatic passage in multilevel systems”. In: *Phys. Rev. A* 44 (7 1991), R4118–R4121. URL: `https://link.aps.org/doi/10.1103/PhysRevA.44.R4118`.

[23] M.J. Snadden et al. “Measurement of the Earth’s Gravity Gradient with an Atom Interferometer-Based Gravity Gradiometer”. In: *Phys. Rev. Lett.* 81 (5 1998), pp. 971–974. URL: `https://link.aps.org/doi/10.1103/PhysRevLett.81.971`.

[24] A. Ashkin. “Trapping of Atoms by Resonance Radiation Pressure”. In: *Phys. Rev. Lett.* 40 (12 1978), pp. 729–732. URL: `https://link.aps.org/doi/10.1103/PhysRevLett.40.729`.

[25] C.S. Adams and E. Riis. “Laser cooling and trapping of neutral atoms”. In: *Progress in Quantum Electronics* 21.1 (1997), pp. 1 –79. ISSN: 0079-6727. URL: `http://www.sciencedirect.com/science/article/pii/S0079672796000067`.

[26] A. Niggebaum. “Towards Mobile Quantum Sensors for Gravity Surveys”. PhD thesis. University of Birmingham, 2015.

[27] J. Vovrosh et al. “Additive manufacturing of magnetic shielding and ultra-high vacuum flange for cold atom sensors”. In: *Scientific Reports* 8.2023 (2018) (2018).

[28] G. Voulazeris. “Portable Atom Interferometry: Investigations of Magnetic Shielding Techniques for Compact Quantum Sensors”. PhD thesis. University of Birmingham, 2018.

[29] T. Lévèque et al. “A laser setup for rubidium cooling dedicated to space applications”. In: *Applied Physics B* 116.4 (2014), pp. 997–1004. ISSN: 1432-0649. DOI: `10.1007/s00340-014-5788-z`. URL: `https://doi.org/10.1007/s00340-014-5788-z`.

[30] F. Theron et al. “Narrow linewidth single laser source system for onboard atom interferometry”. In: *Applied Physics B* 118.1 (2015), pp. 1–5. ISSN: 1432-0649. URL: `https://doi.org/10.1007/s00340-014-5975-y`.

[31] M. Holynski et al. *Gravity Gradiometer*. Patent No. 1721010.5 (Pending).

[32] L. Zhu. “A Cold Atoms Gravimeter for use in Absolute Gravity Comparisons”. PhD thesis. University of Birmingham, 2017.

[33] C. V. Rammeloo. “Optimisation of a Compact Cold-Atoms Interferometer for Gravimetry”. PhD thesis. University of Birmingham, 2017.

[34] M.A. Kasevich et al. “rf spectroscopy in an atomic fountain”. In: *Phys. Rev. Lett.* 63 (6 1989), pp. 612–615. URL: `https://link.aps.org/doi/10.1103/PhysRevLett.63.612`.

[35] M.A. Kasevich. “Atom Interferometry in an Atomic Fountain”. PhD thesis. Stanford University, 1992.

[36] K.I. Lee et al. “Single-beam atom trap in a pyramidal and conical hollow mirror”. In: *Opt. Lett.* 21.15 (1996), pp. 1177–1179. URL: `http://ol.osa.org/abstract.cfm?URI=ol-21-15-1177`.

[37] P. Berg et al. “Composite-Light-Pulse Technique for High-Precision Atom Interferometry”. In: *Phys. Rev. Lett.* 114 (6 2015), p. 063002. URL: `https://link.aps.org/doi/10.1103/PhysRevLett.114.063002`.

[38] A. Sugarbaker et al. “Enhanced Atom Interferometer Readout through the Application of Phase Shear”. In: *Phys. Rev. Lett.* 111 (11 2013), p. 113002. URL: `https://link.aps.org/doi/10.1103/PhysRevLett.111.113002`.

[39] B. Canuel et al. “Six-Axis Inertial Sensor Using Cold-Atom Interferometry”. In: *Phys. Rev. Lett.* 97 (1 2006), p. 010402. URL: `https://link.aps.org/doi/10.1103/PhysRevLett.97.010402`.

[40] S.M. Dickerson et al. “Multiaxis Inertial Sensing with Long-Time Point Source Atom Interferometry”. In: *Phys. Rev. Lett.* 111 (8 2013), p. 083001. URL: `https://link.aps.org/doi/10.1103/PhysRevLett.111.083001`.

[41] N. Poli et al. “Optical atomic clocks”. In: *Rivista Del Nuovo Cimento* 36 (2013), pp. 555–624.

[42] A.L. Migdall et al. “First Observation of Magnetically Trapped Neutral Atoms”. In: *Phys. Rev. Lett.* 54 (24 1985), pp. 2596–2599. URL: `https://link.aps.org/doi/10.1103/PhysRevLett.54.2596`.

[43] T.W. Hänsch and A.L. Schawlow. “Cooling of gases by laser radiation”. In: *Optics Communications* 13.1 (1975), pp. 68–69. ISSN: 0030-4018. DOI: `https://doi.org/10.1016/0030-4018(75)90159-5`. URL: `http://www.sciencedirect.com/science/article/pii/0030401875901595`.

[44] E.L. Raab et al. “Trapping of Neutral Sodium Atoms with Radiation Pressure”. In: *Phys. Rev. Lett.* 59 (23 1987), pp. 2631–2634. URL: `https://link.aps.org/doi/10.1103/PhysRevLett.59.2631`.

[45] A. Arnold. "Preparation and Manipulation of an ^{87}Rb Bose-Einstein Condensate". PhD thesis. University of Sussex, 1999.

[46] A. Hinton et al. "A portable magneto-optical trap with prospects for atom interferometry in civil engineering". English. In: *Philosophical Transactions of the Royal Society of London Series A* 375.2099 (Aug. 2017). ISSN: 0264-3952. DOI: `10.1098/rsta.2016.0238`.

[47] S.J. Blundell and K.M. Blundell. *Concepts in Thermal Physics*. 2nd. Oxford University Press, 2010.

[48] U. D. Rapol, A. Wasan, and V. Natarajan. "Loading of a Rb magneto-optic trap from a getter source". In: *Phys. Rev. A* 64.2, 023402 (Aug. 2001), p. 023402. eprint: `physics/0204022`.

[49] J. Meija et al. "Atomic weights of the elements 2013 (IUPAC Tehnical Report)". In: *Pure Applied Chemistry* 88 (3 2016), pp. 265–291.

[50] G.L. Weissler and R.W Carlson. *Method of Experimental Physics*. Vol. 14. Academic Press, 1979.

[51] D. Alpert and R. S. Buritz. "Ultra-High Vacuum. II. Limiting Factors on the Attainment of Very Low Pressures". In: *Journal of Applied Physics* 25 (Feb. 1954), pp. 202–209. DOI: `10.1063/1.1721603`.

[52] M. Holynski. "Creating a two dimensional cold mixture experiment". PhD thesis. University of Birmingham, 2012.

[53] G.F. Weston. *Ultrahigh Vacuum Practice*. Butterworth & Co. Ltd., 1985.

[54] Y. Koyatsu, H. Miki, and F. Watanabe. "Measurements of outgassing rate from copper and copper alloy chambers". In: *Vacuum* 47.6 (1996). Proceedings of the 13th International Vacuum Congress and the 9th International Conference on Solid Surfaces, pp. 709 –711. ISSN: 0042-207X. URL: `http://www.sciencedirect.com/science/article/pii/0042207X9600053X`.

[55] Kurt J. Lesker Company. *Torr Seal Specifications*. `https://www.lesker.com/newweb/fluids/pdf/torrsealspecifications.pdf`. Accessed: 2018-06-21.

[56] J. A. Rushton, M. Aldous, and M. D. Himsworth. "Contributed Review: The feasibility of a fully miniaturized magneto-optical trap for portable ultracold quantum technology". In: *Review of Scientific Instruments* 85.12 (2014), p. 121501. URL: `https://doi.org/10.1063/1.4904066`.

[57] Crystan. `https://www.crystran.co.uk/optical-materials/optical-glass-n-bk7-and-others`. Accessed: 2018-06-21.

[58] M. Kasevich et al. “Atomic velocity selection using stimulated Raman transitions”. In: *Phys. Rev. Lett.* 66 (18 1991), pp. 2297–2300. URL: `https://link.aps.org/doi/10.1103/PhysRevLett.66.2297`.

[59] K. Moler et al. “Theoretical analysis of velocity-selective Raman transitions”. In: *Phys. Rev. A* 45 (1 1992), pp. 342–348. URL: `https://link.aps.org/doi/10.1103/PhysRevA.45.342`.

[60] W. Demtröder. *Laser Spectroscopy 1: Basic Principles*. 5th. Springer, 2014.

[61] M. Schmidt. “A mobile high-precision gravimeter based on atom interferometry”. PhD thesis. Humboldt University Berlin, 2011.

[62] T.M. Brzozowski et al. “Time-of-flight measurement of the temperature of cold atoms for short trap-probe beam distances”. In: *Journal of Optics B: Quantum and Semiclassical Optics* 4.1 (2002), p. 62. URL: `http://stacks.iop.org/1464-4266/4/i=1/a=310`.

[63] C. Monroe et al. “Very cold trapped atoms in a vapor cell”. In: *Phys. Rev. Lett.* 65 (13 1990), pp. 1571–1574. URL: `https://link.aps.org/doi/10.1103/PhysRevLett.65.1571`.

[64] A. Kaushik. “Trapping, transport and polarisation of ultracold litium”. PhD thesis. Imperial College London, 2014.

[65] D.A. Steck. *Rubidium 87 D Line Data*. `http://www.steck.us/alkalidata/rubidium87numbers.1.6.pdf`. Accessed: 2018-06-21.

[66] F. Christopher. *Atomic Phyiscs*. Oxford University Press, 2005.

[67] D. Anat et al. *Damping of local Rabi oscillations in the presence of thermal motion*. 2013. arXiv: `1208.1396v2 [quant-ph]`.

[68] Ch.J. Bordé. “Atomic interferometry with internal state labelling”. In: *Physics Letters A* 140.1 (1989), pp. 10 –12. ISSN: 0375-9601. URL: `http://www.sciencedirect.com/science/article/pii/0375960189905379`.

[69] B. Desruelle, P. Bouyer, and A. Landragin. *Cold atom gravity gradiometer.* Patent No. 20140190254 (2014-06-10).

[70] A. Roth. *Vacuum Sealing Techniques.* American Institute of Physics, 1994.

[71] J.F. O'Hanlon. *A User's Guide to Vacuum Technology.* Wiley-Interscience, 2003.

[72] S.G. Cox et al. "Reusable ultrahigh vacuum viewport bakeable to 240âĂŁÂřC". In: *Review of Scientific Instruments* 74.6 (2003), pp. 3185–3187. URL: `https://doi.org/10.1063/1.1574604`.

[73] B. Kaltenhäuser et al. "Low retaining force optical viewport seal". In: *Review of Scientific Instruments* 78.4 (2007), p. 046107. URL: `https://aip.scitation.org/doi/abs/10.1063/1.2723065`.

[74] G. Wilpers et al. "A compact UHV package for microfabricated ion-trap arrays with direct electronic air-side access". In: *Applied Physics B* 111.1 (2013), pp. 21–28. ISSN: 1432-0649. URL: `https://doi.org/10.1007/s00340-012-5302-4`.

[75] M. Kasevich and S. Chu. "Measurement of the gravitational acceleration of an atom with a light-pulse atom interferometer". In: *Applied Physics B* 54.5 (1992), pp. 321–332. URL: `https://doi.org/10.1007/BF00325375`.

[76] NKT Photonics. `https://www.nktphotonics.com/lasers-fibers/product/koheras-basik-low-noise-single-frequency-oem-laser-modules/`. Accessed: 2018-08-01.

[77] N. Meyer. "Building and Characterisation of a Dual Species Quantum Simulator". PhD thesis. University of Hamburg, 2014.

[78] H. Müntinga et al. "Interferometry with Bose-Einstein Condensates in Microgravity". In: *Phys. Rev. Lett.* 110 (9 2013), p. 093602. URL: `https://link.aps.org/doi/10.1103/PhysRevLett.110.093602`.

[79] T. van Zoest et al. "Bose-Einstein Condensation in Microgravity". In: *Science* 328.5985 (2010), pp. 1540–1543. ISSN: 0036-8075. URL: `http://science.sciencemag.org/content/328/5985/1540`.

[80] A. Vogel. "Bose-Einstein condensates for space applications and novel teaching concepts". PhD thesis. University of Hamburg, 2009.

[81] D.J. McCarron, S.A. King, and S.L. Cornish. “Modulation transfer spectroscopy in atomic rubidium”. In: *Measurement Science and Technology* 19.10 (2008), p. 105601. URL: `http://stacks.iop.org/0957-0233/19/i=10/a=105601`.

[82] L. Zhu et al. “Application of optical single-sideband laser in Raman atom interferometry”. In: *Opt. Express* 26.6 (2018), pp. 6542–6553. URL: `http://www.opticsexpress.org/abstract.cfm?URI=oe-26-6-6542`.

[83] G. Ritt et al. “Laser frequency offset locking using a side of filter technique”. In: *Applied Physics B* 79.3 (2004), pp. 363–365. ISSN: 1432-0649. URL: `https://doi.org/10.1007/s00340-004-1559-6`.

[84] L. Yuan et al. “Stress-induced birefringence and fabrication of in-fiber polarization devices by controlled femtosecond laser irradiations”. In: *Opt. Express* 24.2 (2016), pp. 1062–1071. URL: `http://www.opticsexpress.org/abstract.cfm?URI=oe-24-2-1062`.

[85] J. Dalibard and C. Cohen-Tannoudji. “Laser cooling below the Doppler limit by polarization gradients: simple theoretical models”. In: *J. Opt. Soc. Am. B* 6.11 (1989), pp. 2023–2045. URL: `http://josab.osa.org/abstract.cfm?URI=josab-6-11-2023`.

[86] W. Ketterle, D.S. Durfee, and D.M. Stamper-Kurn. *Making, probing and understanding Bose-Einstein condensates*. 1999. arXiv: `cond-mat/9904034v2`.

[87] I. Serre, L. Pruvost, and H.T. Duong. “Fluorescence imaging efficiency of cold atoms in free fall”. In: *Appl. Opt.* 37.6 (1998), pp. 1016–1021. URL: `http://ao.osa.org/abstract.cfm?URI=ao-37-6-1016`.

[88] A.F. Linskens et al. “Two-photon Rabi oscillations”. In: *Phys. Rev. A* 54 (6 1996), pp. 4854–4862. URL: `https://link.aps.org/doi/10.1103/PhysRevA.54.4854`.

[89] H. Metcalf and P. van der Straten. *Laser Cooling and Trapping*. Springer, 1999.

[90] A. Dunning. “Coherent atomic manipulation and cooling using composite optical pulse sequences”. PhD thesis. University of Southampton, 2014.

[91] O. Carraz et al. “Phase shift in an atom interferometer induced by the additional laser lines of a Raman laser generated by modulation”. In: *Phys. Rev. A* 86 (3 2012), p. 033605. URL: `https://link.aps.org/doi/10.1103/PhysRevA.86.033605`.

[92] K. Takase. “Precision rotation rate measurements with a mobile atom interferometer”. PhD thesis. Stanford University, 2008.

[93] A. Gelman et al. *Bayesian Data Analysis*. 3rd. CRC Press, 2014.

[94] W.M. Telford, L.P. Geldart, and R.E. Sheriff. *Applied Geophysics*. 2nd. Cambridge University Press, 1990.

[95] K. Czuba and D. Sikora. “Temperature Stability of Coacial Cables”. In: *Acta Physica Polonica A* 4 (2011), pp. 553–557.

[96] Rhode & Schwarz. *R&S HMP Power Supply Family*. Accessed: 2018-09-18.