A PORTABLE COLD ATOM GRAVITY GRADIOMETER WITH FIELD APPLICATION PERFORMANCE

by

Ben Stray

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

Quantum Matter Research Group
School of Physics and Astronomy
College of Engineering and Physical Sciences
University of Birmingham

January 2021

Abstract

A cold atom interferometer is operated in a gradiometer configuration, to overcome the challenges presented by gravity sensing in a civil engineering environment. Such operation allows for the suppression of common mode noise sources, including vibration and environmental effects. Through integration of the instrument into a sensor and control package, the sensor package may be readily carried by two persons in order to operate in the field. An all-in-fibre laser system is used to reduce effects of mechanical shocks associated with transportation and operation outside. The cold atom system is shown to work repeatably over a broad ambient temperature range, with atom numbers of 10^5 participating in each interferometer, at temperatures of $(2.86 \pm 0.09)\,\mu\mathrm{K}$ and $(3.7 \pm 0.2)\,\mu\mathrm{K}$ for the top and bottom clouds, respectively. The gravity gradiometer demonstrates the same performance in and out of the laboratory, with an outdoor performance of $(466 \pm 8)\,\mathrm{E}/\sqrt{\mathrm{Hz}}$, at a 0.67 Hz measurement bandwidth, integrating to a resolution of 17 E. A sensitivity to tilt three times less than an ideal gravimeter is measured, allowing for operation in more challenging environments. Measurements of changes in gravity gradient are demonstrated through the modulation of test masses, resolving signal changes below 50 E. The sensor is operated in a civil engineering environment, demonstrating point to point repeatability of $(52 \pm 40)\,\mathrm{E}$.

Above all else, the mentat must be a generalist, not a specialist. It is wise to have decisions of great moment monitored by generalists. Experts and specialists lead you quickly into chaos. They are a source of useless nit-picking, the ferocious quibble over a comma. The mentat-generalist, on the other hand, should bring to decision-making a healthy common sense. He must not cut himself off from the broad sweep of what is happening in his universe. He must remain capable of saying: "There's no real mystery about this at the moment. This is what we want now. It may prove wrong later, but we'll correct that when we come to it." The mentat-generalist must understand that anything which we can identify as our universe is merely a part of larger phenomena. But the expert looks backward; he looks into the narrow standards of his own speciality. The generalist looks outward; he looks for living principles, knowing full well that such principles change, that they develop. It is to the characteristics of change itself that the mentat-generalist must look. There can be no permanent catalogue of such change, no handbook or manual. You must look at it with as few preconceptions as possible, asking yourself: "Now what is this thing doing?"

The Mentat's Handbook - Children of Dune

Frank Herbert

Contents

Abstract **i**

1 Introduction **1**

1.1 Measuring the gravitational field 2

1.1.1 Operation of gravity sensors 5

1.2 Measuring gravity gradients 7

1.2.1 Common mode noise suppression 8

1.2.2 Reduction in tilt sensitivity 9

1.3 Measuring gravity with atoms 10

1.4 A gravity gradiometer for civil engineering 12

1.5 Thesis outline 14

1.6 Statement of contribution 14

2 Theory and system requirements **15**

2.1 Cold atoms 15

2.1.1 Magneto-optical trapping 17

2.2 Raman transitions 21

2.2.1 Two-photon stimulated Raman transitions 22

2.2.2 Evolution of a three level atom in a light field 25

2.2.3 Rabi oscillations 29

2.2.4 Atom optics 31

2.2.5 Raman beam configuration 32
2.2.6 Velocity selection 34
2.3 Zeeman shift 35
2.4 Atom interferometry 38
2.4.1 Mach-Zehnder sequence 38
2.4.2 Laser phase 40
2.4.3 Interferometer phase 42
2.4.4 Gravity gradients 43
2.4.5 Sensitivity limit of the atom interferometer 45
2.5 Phase components in a gradient measurement 47
2.6 Extracting gravity gradient using ellipse phase 49
2.7 Coriolis force 51
2.8 Initial demonstrator 52
2.9 Design considerations for a new sensor package 54
2.9.1 Target modelling 56
2.9.2 Determining baseline and noise limits 58
2.9.3 Environmental phase shifts 60
2.9.4 Cold atom requirements 61
2.9.5 Sensor package design overview 64

3 Sensor package **66**
3.1 Vacuum chamber 67
3.1.1 Vacuum pressure requirement 67
3.1.2 Science Chambers 68
3.1.3 Prism mount 71
3.1.4 Interfacing module 73
3.1.5 Bakeout for UHV 76
3.2 Mounting structure 78

3.3 MOT coils . . . 79
3.4 Compensation fields . . . 81
3.5 Bias coil . . . 84
3.6 Beam delivery . . . 88
3.7 Detection and diagnostics . . . 90
3.7.1 Detection photodiodes . . . 92
3.7.2 Light sheet detection . . . 93
3.8 Magnetic shielding . . . 93
3.8.1 Main shield . . . 94
3.8.2 Additional magnetic shielding . . . 96
3.9 Integrated sensor head . . . 97

4 Control package **99**
4.1 Laser and RF system . . . 99
4.1.1 Cooling and spectroscopy laser arm . . . 100
4.1.2 Raman laser arm . . . 103
4.2 Timing and electronic control . . . 105
4.3 Data analysis using ellipse-specific fitting . . . 106

5 System characterisation **111**
5.1 Cold atom preparation . . . 112
5.1.1 MOT characterisation . . . 112
5.1.2 Optical molasses . . . 119
5.2 Detection scheme . . . 125
5.3 Spectrum scan . . . 128
5.3.1 Chirp rate calculation . . . 130
5.3.2 Magnetic field profile . . . 132
5.3.3 Bias field switching speed . . . 137

5.4 Rabi oscillations . . . 138
5.5 Sub-level state selection . . . 142
5.5.1 Creating an F=1 blowaway . . . 143
5.5.2 State characterisation and velocity selectivity . . . 144
5.6 Ellipse phase shift . . . 147
5.6.1 Applied relative phase shift . . . 148
5.6.2 Ellipse parameters vs. interferometer T time . . . 149
5.7 Characterisation summary . . . 155

6 System performance and measurement campaigns 156
6.1 Instrument sensitivity . . . 156
6.1.1 Short term vs long term sensitivity . . . 157
6.1.2 Sensor evaluation . . . 158
6.2 Noise sources . . . 159
6.2.1 Detection scheme evaluation . . . 161
6.2.2 Data processing optimisation . . . 165
6.3 Systematics due to external effects . . . 167
6.3.1 Tilt . . . 167
6.3.2 Ambient temperature . . . 168
6.3.3 Coriolis measurement . . . 170
6.4 Measurements of test masses . . . 172
6.5 Outdoor measurements . . . 176
6.5.1 Comparative performance . . . 176
6.5.2 Out of lab performance and repeatability . . . 177
6.6 Road to a civil engineering detection . . . 181

7 Conclusion 184

Appendices

A Rubidium-87 data **187**

B Additional theory **189**

B.1 Propagation phase 189

B.2 Calculating an Allan deviation 192

C Vacuum assembly **194**

C.1 Cleaning for UHV 194

C.2 Indium sealing 195

C.3 UHV epoxy testing 196

D Additional characterisation data **199**

D.1 Time of flight 199

D.2 VRP settings 201

D.3 Atom number in detection 202

List of Figures **205**

List of Tables **209**

List of References **210**

Acknowledgements **225**

Chapter 1

Introduction

Measurements of the local gravity, and gravity gradient, of the Earth, have long been highlighted as a means to measure sub-surface mass changes and density distributions, enabling non-destructive imaging of the underground. The first transportable gravity, and gravity gradient, sensors, developed in the 19th century, derived their measurements from the oscillation time of a double pendulum [1], and the rotation of a torsion balance [2]. Devices of this construction were used ubiquitously until the mid-20th century, contributing to the field of geodesy, as well as oil and mineral prospecting, where fractional changes of g at the 10^{-7} were sufficient to measure large geophysical phenomenon [3]. Today, these devices are largely the remit of museums and collectors, with advances in technology turning focus towards devices measuring objects in free-fall [4], or applying a force to keep a test mass in equilibrium with gravity [5, 6].

With the advent of atom interferometry demonstrating sensitivity to gravity in the 1990's [7, 8], new gravity sensors based on interfering cold atoms began development, with the significantly lower sensitivity limits meaning they quickly outperformed the leading commercial classical sen-

sors. A push in the last two decades has been to bring instruments of this type from the well controlled conditions of the laboratory, to operation in a field environment, whilst maintaining the performance improvements given by using interferometry [9–11]. Use of cold atom systems for applications outside of fundamental physics is often termed "quantum technology", with effort focused on making instruments of this type more robust and compact. Particular emphasis has been put on creating laser systems that can function reliably in the adverse conditions presented by outdoor operation [11–14].

Systems based around quantum technology have been developed to target wider geophysics applications [15], with measurements of large geological features being made [16]. Although the performance of these instruments are promising, any relevance to civil engineering has not been demonstrated in a representative environment. The hurdles to achieving such a measurement are worth tackling; the potential sensitivity increase of quantum devices opens up the possibility for detection of smaller feature sizes and anomalies at lower depths, enabling more widespread use of gravity and gravity gradient measurements in applications such as civil engineering [17, 18].

Here, a cold atom gravity gradiometer is developed for use in a field environment, with a sensitivity relevant to the detection of typical civil engineering features, such as pipes and tunnels. An overview of the thesis may be found at the end of this chapter.

1.1 Measuring the gravitational field

Gravity is the weakest of the fundamental forces, with objects as large as the Earth creating a small attraction between its inhabitants. However, whist the forces are small, the effects of gravity cannot be attenuated - changes in the gravitational field are always measurable provided an instrument is

sensitive enough. To understand the signal sizes associated with gravity, first consider the force due to gravitational attraction between two objects. This is given by the equation for universal gravitation [19], and depends on the masses of the two objects m and M, inversely scaled by the square of the distance r between them

$$\mathbf{F} = \frac{GmM}{r^2}\hat{\mathbf{r}}. \tag{1.1}$$

To measure the strength of the gravitational field, referred to simply as gravity, from a mass M, the substitution $\mathbf{F} = m\mathbf{g}$ is made. This results in the field equation (1.2), where additionally the mass M, is expressed in terms of a spatially varying density ρ, integrated over the volume V. The value of gravity g, at a given point, is the magnitude of the infinite sum of all contributions of mass, scaled by the square of their relative distances

$$g = |\mathbf{g}| = \left|\frac{GM}{r^2}\hat{\mathbf{r}}\right| = \left|G\int \frac{\rho \cdot \hat{\mathbf{r}}}{r^2} dV\right|. \tag{1.2}$$

For the gravitational field of Earth, there is a moderate variation over the surface due to a variety of global and local effects. The oblate spheroid shape of the Earth provides the largest positional change, with the average value of g changing with latitude, from $9.78\,\mathrm{m\,s^{-2}}$ at the equator, to $9.83\,\mathrm{m\,s^{-2}}$ at the poles [20]. Interest in temporal and spatial perturbations to these values form the field of geodesy, where geological, topographical, and environmental features all shift the local value of g by amounts dependent on proximity and size. Typically these shifts change the value below the order $10^{-5}g$ [3]. Ocean loading, ground water levels, and tides all add a time dependence to the local value of g, with different characteristic times. Shifts of this kind may be readily measured by gravity sensors operating at the $10^{-8}g$ level

[3, 21].

All measurements of gravity suffer from the same unavoidable issue as a consequence of the equivalent principle - acceleration due to the gravitational field is indistinguishable from acceleration produced by an inertial force. The consequence of this is that any acceleration, such as vibration, will be measured equivalently to the gravitational field. Vibration and micro-seismic noise is dominated by low frequencies, of order 0.1 Hz [17, 22]. Without appropriate compensation, a gravity measurement will require averaging over many cycles to reduce vibrational noise, in order to be able to measure the small signals produced by a gravitational anomaly. For features smaller than the gravity signals of tidal effects, the resolution of such an instrument relies heavily on an accurate tidal model being subtracted from the raw data.

From a fully characterised field profile, where the value of gravity is known at all spatial positions and times, the mass distribution responsible for the field may be determined with a unique solution. In general, measurements of gravity using a gravity sensor, often termed gravimeter, are limited to measurements across a small area of the surface, or even just a single line, due to long measurement times and accessibility. Since the characterisation of the field is incomplete, the inversion of the gravity map leads to degenerate solutions of the mass distribution; the same signal may be produced by a massive object far away, or a small object close by. Prior information may be input to the calculation, such as expected density contrasts, in order to reduce the number of possible solutions and calculate the most likely gravitational anomaly [3, 23, 24]. The use of complementary measurement techniques may also be used to rule out solutions and lift the degeneracy. Complementary information may come from additional measurements of the gravity gradient, or through other modalities, such as magnetometry.

1.1.1 Operation of gravity sensors

Measurements of gravity may be characterised by the terms "absolute" or "relative". An instrument that independently measures acceleration due to gravity, without external calibration, may be termed an absolute gravimeter. If, however, the value of gravity may only be derived through calibration to a known gravity reference, then the instrument is termed a relative gravimeter. The use case tends to determine which technology is used, and thus the type of measurement outputted.

For applications targeting monitoring, the instruments tend to be transportable systems that are installed for long measurement campaigns. Often, long term precision is prioritised in order to measure changes in the gravity signal with time. To achieve this, measurements of this type require low instrument drift [25]. Applications requiring accuracy may prioritise use of an absolute gravimeter, but many applications are interested in spatial or temporal changes compared to a reference point, meaning a relative measurement of gravity is usually sufficient. Examples of gravimeters used for monitoring purposes are instruments based on a levitated superconducting sphere [25], and the free-fall of a corner cube reflector [4]. The most widespread absolute gravimeter used for monitoring, measures the change in fringe phase of a Michelson interferometer, during the free fall of a corner-cube reflector in one of the interferometer arms [4]. Instruments of this type have demonstrated accuracy to to 2 ng [26]. In general, mechanical systems tend to lower the measurement bandwidth to 1 drop every 10 minutes, in order to reduce mechanical wear of the free-fall mechanism. Free-fall corner cube gravimeters already operate close to their noise floor, meaning that orders of magnitude improvements in performance are unlikely to be made [27].

Gravity sensors may also be used in the context of fundamental physics and the determining of physical constants. Here, accuracy is required in order to calculate other quantities, for example the Kibble balance, where the weight of the balanced test mass is dependent upon the value of local g [28].

In order to achieve the best results for instruments whose function is monitoring and accuracy, the conditions of the measurement space must be prepared to provide a low noise environment, shielded from environmental and vibrational noise. The alignment of a gravimeter must also be carefully controlled, to avoid changing the measured value of gravity. An instrument aligned at an angle to the local gravitational acceleration will measure the projection of the gravity vector $\mathbf{g}$, scaled by the cosine of the angle. Signal sizes for gravitational changes are small, and the precision of a gravimeter will be fundamentally limited by the ability to control the tilt of the instrument. A signal size of order $10^{-9}g$ is typical for civil engineering applications, and would be obscured by a tilt of 2.5 mdeg.* An additional challenge for gravimeters is thus providing a stable, low noise measurement location in order to reach an instrument's best performance.

Geophysical and civil engineering applications often require operation of instruments outside of designated measurement facilities, increasing the challenges presented by gravity sensing. Additional noise sources are introduced, such as: wind noise, where a field instrument is buffeted by the wind; and temperature fluctuations. Wind noise induces both inertial noise and changes to instrument tilt, especially when surveying near trees, where the motion of the trunk causes the roots to change the ground level [29].

*For measurements of angle, degrees will be written as the contraction "deg" with standard SI prefixes of order. This is to avoid ambiguity when discussing in conjunction with temperature.

For spatial mapping, an instrument will be moved multiple times, requiring fast measurements and point to point repeatability. Field instruments also require additional protection from environmental effects, and require rugged packaging to shield against temperature fluctuations, wind noise, magnetic fields, and more. They also need to be readily transported, and without long setup times. For field applications, spring based gravimeters, such as the *Scintrex* CG-6 relative gravimeter, are a popular choice. They offer a compact size and usability, whilst providing a measurement repeatability of 5 ng [30]. Spring based gravimeters suffer from drifts of order 100 ng/day in the measured value of g [30], but repeated measurements at a base station point can allow for linear corrections to be applied to the data in post-processing [3]. This means that survey practice can often be used to overcome some shortcomings of an instrument.

1.2 Measuring gravity gradients

Measurements of gravity, using two sensors separated by a vertical distance, can be used to calculate the vertical gradient of the vertical component of the gravitational field, which shall be colloquially referred to as gravity gradient.* As well as providing additional information about the field profile, certain implementations of the gradient measurement can be used to overcome some of the identified issues with gravity sensing; namely, the subtraction of common noise sources, such as vibration and microseismic noise, as well as a reduction in the sensitivity to tilt. To achieve such an implementation, two measurements of the gravitational field must be taken

*To measure the gravity gradient completely, a full 3×3 tensor is required, providing the gradients between each spatial axis. Here, the component T_{zz} is discussed only, which provides the most information about changes in the local gravitational field due to subsurface voids.

simultaneously, and in a well coupled system, such that the noise measured at one position is measured equally at the other. When taking the subtraction of the measurement, required to calculate the gradient, the common noise sources are cancelled, leaving only the gradient of the gravitational field.

The resulting gradient signal is reduced when compared to the signal sizes associated with gravity; however, since the gradient can be changed by increasing the baseline, the signal size may be tuned. The benefits of gradiometry over gravimetry mean that when measuring equivalent anomalies, the signal to noise ratio is often better, since offsets and systematic shifts are suppressed or reduced in a coupled gradient measurement [17].

1.2.1 Common mode noise suppression

In the case of the well coupled system and the two measurements occurring simultaneously, the common mode noise sources can be suppressed. Gradiometry can therefore allow for an instrument to operate in environments with increased seismic activity, natural or man-made, which generally are not explored using gravimeters. Without the requirement for long integration times, a gradiometer with common mode noise suppression may provide faster measurements in the field, increasing the use of gravity as a civil engineering tool.

The mechanism for a coupled system determines how well the common mode noise suppression can be achieved. High sensitivity in a gravity gradiometer comes from both the performance of the acceleration measurement at each spatial point, and the size of the baseline between each sensor. Achieving high performance in each sensor requires all systematic effects to be accounted for and the relative alignment of each sensor to be fixed.

Furthermore, coupling of inertial noise must be exact between each measurement point, which becomes more difficult as the baseline increases into the metre scale. If noise sources within the sensor can be made to be common between each sensor measurement, then the gradiometer can also reduce some of the technical noise constraints, which would have introduced noise or systematic shifts in an equivalent gravimeter.

1.2.2 Reduction in tilt sensitivity

Instead of measuring the tilt projection onto the radius of the Earth, the gravity gradient measurement instead measures the tilt with respect to the baseline of the instrument. The angle with respect to the gravitational field is common between both measurements, with the tilt taken as a projection on the baseline instead of the Earth's radius. The resulting shifts require a larger tilt in the gradiometer to achieve the equivalent shift in the gravimeter. This is shown in figure 1.1. Where a tilt of 2.5 mdeg was enough shift the measured value of gravity by 1 ng, a comparable gradient signal size would require an instrument tilt of 1.46 deg to reduce the signal size by 1 E.*

To illustrate the point further, assuming that the tilt can be controlled to within 2.5 mdeg in the field, the gradiometer measurements will have an uncertainty due to tilt of $3 \cdot 10^{-6}$ E, significantly lower than typical signal sizes of order 10 E. Again, this suppression in instrument tilt, when compared to a gravimeter, may only be achieved in a well coupled system, where both measurement points experience the same tilt.

*Eötvös is a typical unit for gravity gradient, where $1\,\mathrm{E} = 1 \cdot 10^{-9}\,\mathrm{s}^{-2}$

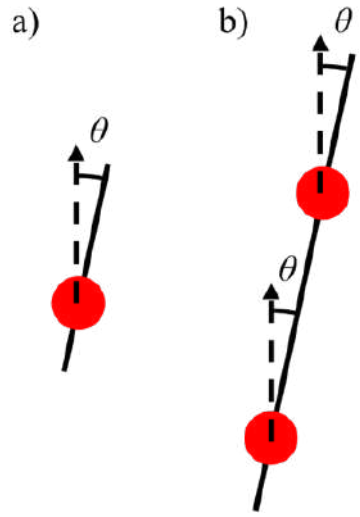

Figure 1.1 Instrument tilt causes a measurement of the projection of gravity. a) In the case of a gravimeter the tilt is projected onto the radius of the Earth, causing small tilts to significantly change the measurement of g. b) The effect of tilt is significantly reduced in a gradiometer configuration, where a common instrument tilt is projected on to the distance between the two sensors, of order meters, and evaluated at the average gravity gradient of Earth, $3086\,\mathrm{E\,m^{-1}}$ [3].

1.3 Measuring gravity with atoms

Atoms provide an ideal test mass for measurements of gravity. Each atom is indistinguishable from others of the same isotope, making a universal reference mass that does not depend upon manufacture tolerances. Through interactions with optical and RF frequencies, as well as magnetic fields, ensembles of atoms can be trapped, cooled, and manipulated. Using the techniques of atom interferometry, the sensitivity limits of an atom based gravity sensor far exceed any classical sensor, of which the best already operates at its classical noise limit [27]. Even with technical noise, cold atom based gravimeters have been shown to outperform the classical standard in sensitivity, even when operating in the standard interferometry scheme - the Mach-Zehnder type interferometer [11, 31–33].

Atom interferometry puts a free falling atom into a quantum superpo-

sition of two momentum states, creating two spatial paths for the atom to travel over. When the paths are recombined, a phase difference between each path can be determined, from which a value of gravity can be measured. Clouds of cold atoms are used to perform atom interferometry, allowing for millions of individual atoms to undergo the same sequence at one time. Large statistical significance may therefore be placed on the resulting gravity measurements. The sensitivity of an atom interferometer depends its area - the integrated space enclosed by these two paths, and can be increased through lengthening the time the atoms evolve for, or by separating the two paths by a larger distance through techniques such as large momentum transfer (LMT) [34–37]. The current performance of cold atom gravimeters shows promise as a competitor to current technologies, with a trajectory to orders of magnitude improvements to the performance through techniques such as LMT. The potential improvement over current technologies makes the use of cold atoms in an inertial sensor appealing; the remaining challenges are mostly of a technical nature.

For use in measurements of gravity gradient, atoms provide an ideal medium since spatially separated atoms can be directly coupled with a common laser beam. The common laser beam is used as the reference between the two measurements and noise sources, such as vibration or laser phase noise, are imparted on the atoms equally. The gradient measurement removes these common effects, as discussed, leaving only uncommon systematic shifts that effect each atom differently. Suppression of common mode noise has been measured to below 140 dB in a cold atom gradiometer [38], demonstrating the suitability for using atom interferometry to construct a gravity gradiometer.

Gravity gradient measurements, using atom interferometry, have been

performed in laboratories to measure fundamental constants such as: the fine-structure constant [39], the gravitational constant [40], as well as tests of the equivalence principle and general relativity [41, 42]. Cold atom gravity gradiometry has also been proposed as a means to measure mid-band gravitational waves, with construction of initial experiments underway [43–45]. These instruments are designed to be very targeted in their application, enabling the high precision required, but in a form factor which is not directly translatable into other applications.

1.4 A gravity gradiometer for civil engineering

In civil engineering, gravity, and gravity gradient, maps may be used to understand the sub-surface. The most prevalent identifiable features are those which demonstrate a large contrast in mass, or density, compared to the surrounding area, which for civil engineering typically means a soil to air boundary. Thus, gravitational anomalies of interest usually fall in the category of: mine shafts, tunnels, sinkholes, pipes, etc.

Current gravity sensing in the field is usually performed by the CG-5/CG-6 relative spring based gravimeters. Mapping an area for microgravity, before construction, can reduce additional costs due to unexpected sub-surface voids. However, this is rarely performed due to the length of time such a measurement set takes. Current gravity maps take 10-15 minutes per measurement point, making gravity mapping for large scale civil engineering targets prohibitive, due to time and associated costs.

For typical signal sizes in a civil engineering field application, it has been shown, through simulation, that atom interferometry based gravity gradiometry could offer a significant benefit through reduction in inertial noise, leading to an overall higher signal to noise ratio compared to gravimetry [17].

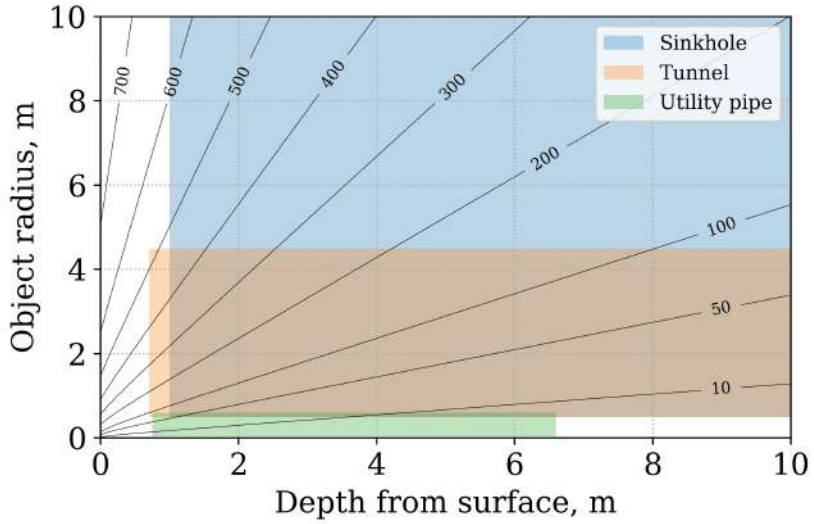

Figure 1.2 Contours showing signal sizes in Eötvös for a density contrast of soil to air, for a gradiometer with a 1 m baseline. Shaded regions show the parameter space for some typical civil engineering features [46].

This is the case for both current classical gravimeters, or future quantum technology based gravimeters.

Modelling a gravitational anomaly as an infinite cylinder of different radii and depths from surface, allows the parameter space of civil engineering to be explored. This is shown in figure 1.2, where contours show the expected signal in units of gravity gradient, and approximate dimension of civil engineering features are overlain. Despite the promising detection depths predicted in the literature [17], the parameter space for anomalies are restricted to near surface at current instrument performances [10]. For an instrument to cover all applications in civil engineering, the sensitivity and repeatability should be such that a signal size of a single Eötvös is resolved. Since current gravity measurements are limited due to integration time, a cold atom gravity gradiometer should aim to reach a comparable resolution in a shorter time.

1.5 Thesis outline

An overview of the theory for atom interferometry, using cold atoms, is introduced in chapter 2, as well as the parameter space a field capable gravity gradiometer should sit within. This is used to help design a new sensor package, using an existing control package, to reach application ready sensitivities. The design and construction of this new sensor package is discussed in chapter 3, and an overview of the supporting control package in chapter 4. Once integrated as a complete system, the characterisation and optimisation of the gravity gradiometer are discussed in chapter 5. The final performance of the instrument is examined in chapter 6, where initial characterisation of systematic effects has begun, determining the instrument's suitability as a field instrument.

1.6 Statement of contribution

Unless otherwise stated or referenced, the data, analysis, and work are the authors own.

Chapter 2

Theory and system requirements

The gravity gradiometer is based on Raman atom interferometry, using cold atom clouds. The cold atom preparation is first introduced, then the framework of the Raman transition is discussed, to enable the parameter space of the gradiometer to be explored. From this, the design of a gravity gradiometer is proposed, which meets the requirements of operation in a field environment.

2.1 Cold atoms

By reducing the average velocity of thermal atoms to only a few mm per second, ensembles of atoms may be described by temperatures of only a few μK. The mean velocity of atoms with mass m, may be related to temperature, T, through equating its thermal and kinetic energy with the Boltzmann constant, k_B, such that

$$\overline{v} = \sqrt{\frac{k_B T}{m}}. \tag{2.1}$$

Since the temperature is derived from the Maxwell-Boltzmann distribution, the temperature for a single atom cannot be well defined. For an ensemble of atoms, the temperature of the cloud may be estimated from the average velocity. Using standard laser cooling techniques, temperatures proportional to the recoil velocity can be achieved, which gives a minimum thermal velocity of $5.8845\,\mathrm{mm\,s^{-1}}$ for rubidium-87 [47]. Thus, for atom evolution times of less than a second, a cooled atom will travels only a few mm from where it was cooled and trapped. Ensembles of atoms can therefore remain localised over the evolution time of an experimental sequence.

Experiments requiring evolution times of many seconds may wish to use colder samples, forming atomic ensembles into a Bose-Einstein condensate (BEC), where the ensemble acts as a single wavefunction and at a reduced temperature. In this work, creating a BEC would produce significant technical challenges for operation in a portable system, though it has been successfully demonstrated in a sounding rocket and drop tower experiments [48–50]. Instead, operation at application relevant sensitivities may be achieved using thermal atom ensembles, cooled to temperatures of single µK.

Generation of cold atoms has been performed using a standard combination of background vapour loaded, 3D magneto-optical trapping, with an optical molasses to generate atomic ensembles of millions of atoms, cooled to a few µK. Atoms in a thermal cloud have typical densities of order $1 \cdot 10^{10}\,\mathrm{cm^{-1}}$ [51], with little atom-atom collision. Each cloud may therefore be treated as the statistical average of millions of individual atom-light interactions. Details of the trapping and cooling techniques are well doc-

umented and detailed in the literature [52, 53], but an introduction to the magneto-optical trap is included for completeness.

2.1.1 Magneto-optical trapping

Mechanism

The physics of a magneto-optical trap (MOT), is an extension to laser cooling, where an atom, moving with a relative velocity to a laser beam detuned from resonance, experiences a scattering force proportional to its velocity vector. The MOT adds an additional force through the introduction of a quadrupole magnetic field, creating a potential well for atoms to become trapped in.

To create a 3D MOT, three orthogonal pairs of beams, detuned below the resonant frequency (often called red detuning), intersect at the magnetic zero of a quadrupole field. With appropriate choice of circular polarisation, atoms travelling through the beams experience a force directed towards the centre of the trap, where the magnetic field is zero. The magnitude of the force is dependent upon the position and velocity of the atoms with respect to the trap centre. The mechanism for this is shown in figure 2.1, where light detuned from the atomic resonance is incident from both directions, and is preferentially scattered when an atom moves towards one of the beams, Doppler shifting the frequency closer to resonance. The magnetic field introduces a position dependent shift to the resonant frequency, increasing the strength of the scattering force with distance from the trap centre. Multiple scattering events are required to remove the kinetic energy of the atom, and the net effect is to cool and push the atom towards the centre of the trap, until the scattering rate from all directions is equal. This represents the Doppler limit where the force from scattering is bal-

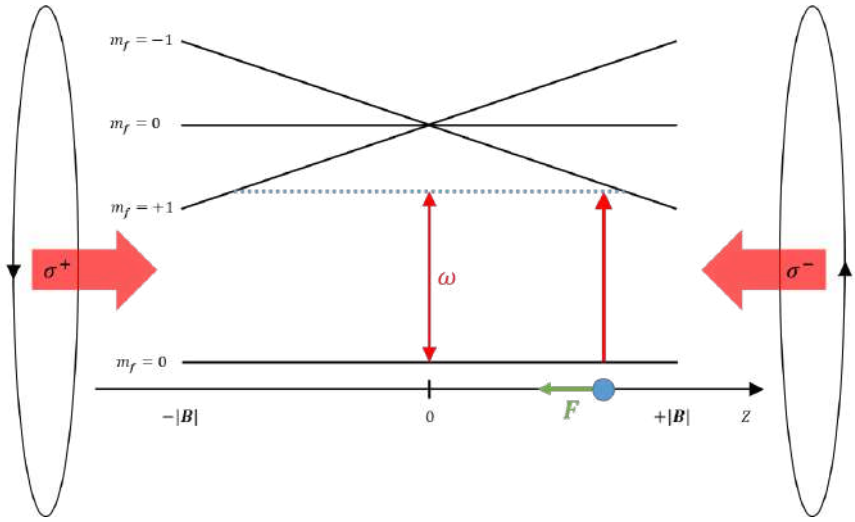

Figure 2.1 Mechanism of the magneto-optical trap in the simplified case of an atom with a transition between $|J = 0\rangle \rightarrow |J' = 1\rangle$. A quadrupole magnetic field centred at zero, spatially splits the degeneracy of the magnetic sub-levels. Light beams are incident from opposite directions with a polarisation which drives $\sigma_{\pm}$ transitions, at a frequency tuned below the resonant frequency of the transition. The resonance condition for the detuned light is satisfied away from the centre of the trap. An atom will preferentially scatter a photon from the direction where the detuning from resonance is smallest. The figure shows the theory of magneto-optical trapping in 1D; extension to 3D may be made by introducing two additional pairs in the orthogonal planes. In 3D, the correct choice of polarisation directs the scattering force towards the centre of the trap for all directions that the atom can move. As well as trapping, the atoms are cooled via the laser cooling mechanism to temperatures close to the Doppler limit.

anced with the heating from random emissions and absorption, providing the lowest temperature achievable in the laser cooling of rubidium, without considering sub-Doppler cooling [54]. The MOT temperature is of order the Doppler limit, and may be expressed as

$$T_D = \frac{\hbar\Gamma}{2k_B}, \tag{2.2}$$

where Γ is the transition linewidth. For rubidium-87, this limit is 146.6 µK, requiring additional sub-Doppler cooling to reach single µK temperatures.

Since the trapping region is finite, atoms which are slowed but reach the edge of the trapping region before stopping will not be trapped. The upper limit of atomic velocity for which trapping will occur may be estimated by the capture velocity [55], given by

$$v_c \approx \sqrt{\frac{\hbar k \Gamma D}{2m}}. \tag{2.3}$$

For a trap with beam diameters, and thus trapping distance, D of 15 mm, the capture velocity is of order $40\,\mathrm{m\,s^{-1}}$ for ^{87}Rb. Atoms with thermal velocity below the capture velocity are readily available from room temperature vapour, allowing for direct trapping from background gases, rather than injection from a slowed atom beam, usually generated from a 2D MOT [56].

MOTs have become a near ubiquitous starting point for cold atom experiments for group 1 alkali and group 2 alkaline earth metals, allowing millions of atoms to be trapped within a few hundred ms, and confined to a trap centre of only a few mm. They have also been successfully applied to sympathetically cool ions [57], as well as produce cold molecules [58].

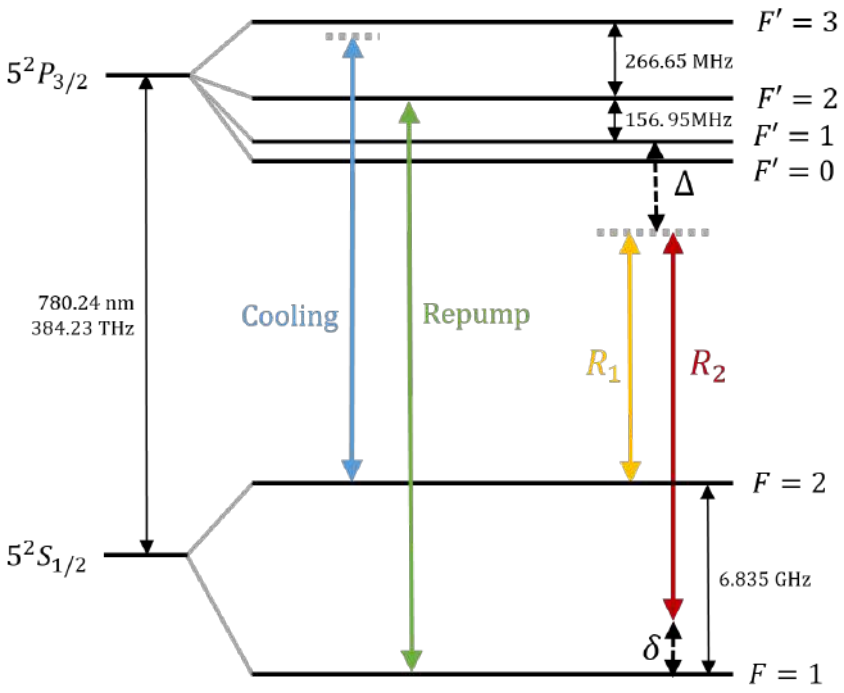

Figure 2.2 Energy level diagram of ^{87}Rb D_2 line. Transitions of the cooling, repump, as well as those of the Raman transitions, are shown. Adapted from [47].

Trapping and cooling ^{87}Rb

Rubidium is often referred to as the workhorse of cold atom physics, being used to achieve the first demonstration of a BEC [59], and many successful cold atom experiments. The element comes in two isotopes: a more abundant ^{85}Rb, and the widely used ^{87}Rb. Each isotope has a relatively simple energy level structure, including resonances with frequencies in the near infra-red, where a large variety of lasers exist. Operation of trapping, cooling, and Raman interferometry may be performed on the D_2 line of ^{87}Rb, which corresponds to transitions from the $5^2S_{1/2}$ ground state to the $5^2P_{3/2}$ upper manifold. The energy level diagram may be seen in figure 2.2, with additional atomic properties in appendix A.

Laser cooling relies on multiple scattering events, since the momentum

change from a single photon scattering is small. This can be achieved through the use of a cycling transition, in which an atom is returned to the original state after the scattering process. ^{87}Rb has a cycling transition between its ground hyperfine state $|F = 2\rangle$, and the excited hyperfine state $|F' = 3\rangle$. Transitions of this type will be denoted by $|F = 2\rangle \leftrightarrow |F' = 3\rangle$, where the prime notation is used to denote the hyperfine level of the excited state manifold $5^2P_{3/2}$. This transition is termed the cooling transition, with light exciting transitions between these states colloquially referred to as cooling light.

As there is a finite detuning between the transition $|F = 2\rangle \leftrightarrow |F' = 2\rangle$, there is a scattering rate associated with exciting an atom to $|F' = 2\rangle$ instead of $|F' = 3\rangle$. If excited to the state $|F' = 2\rangle$, the electron may decay to either state $|F = 1\rangle$ or $|F = 2\rangle$, with the probability determined by branching ratios, calculated from the Clebsch-Gordan coefficients [47]. Electrons that fall into the state $|F = 1\rangle$ may no longer participate in the cooling cycle and thus have no force confining them to the MOT. To mitigate this, atoms are pumped back into the cooling cycle by the introduction of an additional frequency, resonant with $|F = 1\rangle \leftrightarrow |F' = 2\rangle$, from which they may decay back into the $|F = 2\rangle$ state. This transition will be termed the repump transition, with a complimentary definition of repump/repumper light.

2.2 Raman transitions

An atom interferometer uses light pulses to create a quantum superposition of atomic states.* Such an interferometer relies upon the ability to coherently transfer population probability between two internal states of

*Atom interferometers can be constructed using the same internal atomic state, such as in the case of double-diffraction [60].

an atom, and may be achieved through the use of two-photon stimulated Raman transitions. Atom-light interactions of this type may be tuned to transfer populations between ground states completely, or create superposition states, used to form the atom optic equivalent of beam splitters and mirrors.

To maximise interferometer area, stimulated Raman transitions will be used in a counter-propagating configuration to impart $\approx 2\hbar\mathbf{k}$ photon momentum to the atom. An abridged version of the formalism is introduced to allow the motivation of the interferometry scheme used, and then extended to that of a gravity gradiometer. The derivation follows closely the thesis of Rammeloo [9], with additional references from [11, 61, 62].

2.2.1 Two-photon stimulated Raman transitions

To construct the atom optics used in the interferometer, coherent manipulation of atomic states is required. Two-photon stimulated Raman transitions allow the coherent coupling of two internal energy states by two optical photons, with a frequency difference equal to the energy state splitting. Such a transition is preferential in Rubidium, where the momentum imparted to the atom can be of order 10^5 times larger than the single photon microwave transition between the two ground states. The sensitivity of the interferometer is proportional to the effective wavevector, discussed later in section 2.4.5, making the two-photon optical transition favourable. Atomic species with long lifetime excited states, such as strontium, may use single photon transitions to build the interferometer, essential in very long baseline gradiometry [44]. The Raman transition is shown in figure 2.3, where population is driven between states $|g\rangle$ to $|e\rangle$, via a two photon interaction, detuned from an upper state $|i\rangle$. The transition occurs simultaneously,

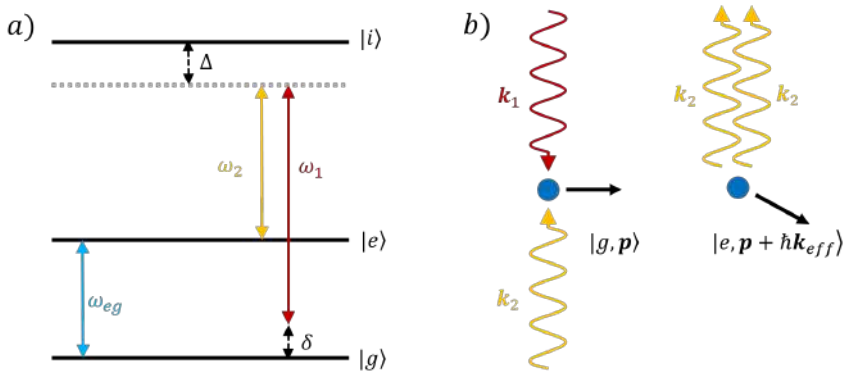

Figure 2.3 a) Two long lived states $|g\rangle$ and $|e\rangle$, separated by a frequency ω_{eg}, can be coupled by a two photon stimulated Raman transition. Diagram is not to scale and $\omega_1, \omega_2 >> \omega_{eg}, \delta, \Delta$. b) The atom simultaneously absorbs a photon from one beam and has emission stimulated from the virtual level by the other. The atom gains momentum equal to the difference between the wavevectors $\mathbf{k}_{\text{eff}} = \mathbf{k}_1 - \mathbf{k}_2$ which is maximal when operated in a counter-propagating configuration as shown. The frequencies do not excite transitions from other states but will contribute to higher order terms in the solutions to the state evolution.

but may be considered as the absorption of one frequency and the stimulated emission of the other. Transition rates are determined by the coupling strength of the equivalent resonant single photon transitions and the detuning from such states. The coupling between the two internal states via a Raman transition will be presented.

For the two-photon transition to be resonant, the frequency difference between the photons must equal that of the ground state splitting, ω_{eg}, plus energy shifts due the atoms kinetic energy. For an atom with momentum $\mathbf{p}$, a transition from state $|g\rangle$ to state $|e\rangle$ will impart photon momentum equal to $\hbar$ times the difference of the wavevectors. This difference is defined as $\mathbf{k}_{\text{eff}} = \mathbf{k}_1 - \mathbf{k}_2$. A detuning from the resonance condition can be written as

$$\delta = \omega_1 - \omega_2 - \omega_{eg} + \frac{|\mathbf{p}|^2}{2m\hbar} - \frac{|\mathbf{p} + \hbar\mathbf{k}_{\text{eff}}|^2}{2m\hbar}, \tag{2.4}$$

and will be referred to as the two-photon detuning.

To avoid single photon scattering, both frequencies ω_1 and ω_2 are detuned by a frequency Δ, from their respective ground state to the intermediate state $|i\rangle$. In order to neglect spontaneous emission, $\Delta >> \delta$. Defining the frequency difference between state $|g\rangle$ and $|i\rangle$ as ω_{ig}, the single photon detuning is given by

$$\Delta = \omega_{ig} - \omega_1 + \frac{|\mathbf{p}|^2}{2m\hbar} - \frac{|\mathbf{p} + \hbar\mathbf{k}_1|^2}{2m\hbar}, \tag{2.5}$$

but could equivalently be defined relative to ω_2. It also useful to define the energy of each state, for an atom with initial momentum $\mathbf{p}$

$$\hbar\omega_g = \frac{|\mathbf{p}|^2}{2m}, \tag{2.6a}$$

$$\hbar\omega_e = \frac{|\mathbf{p} + \hbar\mathbf{k}_{\text{eff}}|^2}{2m} + \hbar\omega_{eg}, \tag{2.6b}$$

$$\hbar\omega_i = \frac{|\mathbf{p} + \hbar\mathbf{k}_1|^2}{2m} + \hbar\omega_{ig}. \tag{2.6c}$$

Treating the system in a semi-classical approach, the atom is treated as a quantum object in a classical light field. In this basis, the internal energy and external momentum state can be represented by the tensor product of the two Hilbert spaces, such that for an atom with internal state $|x\rangle$ and momentum $|\mathbf{p}_x\rangle$, the state can be written as $|x, \mathbf{p}_x\rangle = |x\rangle \otimes |\mathbf{p}_x\rangle$.

The wavefunction of the atom is a linear combination of all states, with time varying coefficients, whose square give the probability of the atom being in the state at a certain time. The wavefunction for the three level system for a single momentum $\mathbf{p}$ is thus

$$|\Psi_p(t)\rangle = C_{g,\mathbf{p}}(t)\,|g,\mathbf{p}\rangle + C_{e,\mathbf{p}+\hbar\mathbf{k}_{\text{eff}}}(t)\,|e,\mathbf{p}+\hbar\mathbf{k}_{\text{eff}}\rangle + C_{i,\mathbf{p}+\hbar\mathbf{k}_1}(t)\,|i,\mathbf{p}+\hbar\mathbf{k}_1\rangle\,, \tag{2.7}$$

where the coefficients $C_{x,\mathbf{p}_x}$ evolve with a fast oscillation, at the frequency defined in equations (2.6), and a slow varying coefficient such that

$$C_{x,\mathbf{p}_x}(t) = c_{x,\mathbf{p}_x}(t)e^{-i\omega_x t}. \tag{2.8}$$

To determine the state after an evolution in the optical field, one needs to solve the time dependent Schrödinger equation, (2.9), for the wavefunction in equation (2.7).

$$i\hbar\frac{d}{dt}\,|\Psi(t)\rangle = \hat{H}\,|\Psi(t)\rangle\,. \tag{2.9}$$

2.2.2 Evolution of a three level atom in a light field

The solution of the Schrödinger equation requires the Hamiltonian of the system. Here, the Hamiltonian can be considered as consisting of three parts: $\hat{H}_P$, associated with the kinetic energy of the atom, via the momentum operator; $\hat{H}_A$, acting on the internal energy levels of the atom; $\hat{H}_{int}$, the interaction between the atom and the light field, via the electric dipole operator. Written out, these are

$$\hat{H}_P = \frac{\hat{\mathbf{p}}^2}{2m}, \tag{2.10}$$

$$\hat{H}_A = \hbar\omega_e\,|e\rangle\langle e| + \hbar\omega_g\,|g\rangle\langle g| + \hbar\omega_i\,|i\rangle\langle i|\,, \tag{2.11}$$

$$\hat{H}_{int} = -\hat{\mathbf{d}}\cdot\mathbf{E}, \tag{2.12}$$

with the total Hamiltonian of the system therefore

$$\hat{H} = \hat{H}_P + \hat{H}_A + \hat{H}_{int}. \tag{2.13}$$

For the two photon transition, the light field is given by the sum of the two waves

$$\mathbf{E} = \mathbf{E}_1 \cos(\mathbf{k}_1 \cdot \mathbf{z} - \omega_1 t + \phi_1) + \mathbf{E}_2 \cos(\mathbf{k}_2 \cdot \mathbf{z} - \omega_2 t + \phi_2). \tag{2.14}$$

Substitution of equations (2.7), (2.13), and (2.14) into the Schrödinger equation, (2.9), yield the three level coupled equations

$$\frac{dc_{g,\mathbf{p}}(t)}{dt} = -ic_{i,\mathbf{p}+\hbar\mathbf{k}_1}\Omega_1 e^{-i(\Delta t+\phi_1)}, \tag{2.15a}$$

$$\frac{dc_{e,\mathbf{p}+\hbar\mathbf{k}_{\text{eff}}}(t)}{dt} = -ic_{i,\mathbf{p}+\hbar\mathbf{k}_1}\Omega_2 e^{-i((\Delta+\delta)t+\phi_2)}, \tag{2.15b}$$

$$\frac{dc_{i,\mathbf{p}+\hbar\mathbf{k}_1}(t)}{dt} = -ic_{g,\mathbf{p}}\Omega_1^* e^{i(\Delta t+\phi_1)} - ic_{e,\mathbf{p}+\hbar\mathbf{k}_{\text{eff}}}\Omega_2^* e^{i((\Delta+\delta)t+\phi_2)}, \tag{2.15c}$$

where the definition of the Rabi frequency is introduced

$$\Omega_1 \equiv -\frac{\langle g|\,\hat{\mathbf{d}}\cdot\mathbf{E}_1\,|i\rangle}{2\hbar}, \tag{2.16a}$$

$$\Omega_2 \equiv -\frac{\langle e|\,\hat{\mathbf{d}}\cdot\mathbf{E}_2\,|i\rangle}{2\hbar}. \tag{2.16b}$$

The Rabi frequency represents the coupling between each ground state and the intermediate state, through the interaction with the light field. It gives

the rate of population transfer between states, determined by both the properties of the light and the coupling strength between the two atomic states. This coupling may be calculated through the reduction of the dipole operator, into a matrix of Clebsch-Gordan coefficients, which determine the allowed transitions and relative coupling between states [47].

For spontaneous emission to be suppressed, $|\Delta| >> |\delta|, |\Omega_1|, |\Omega_2|$. Application of this condition finds that the time dependent oscillation of the coefficients $c_{g,\mathbf{p}}$ and $c_{e,\mathbf{p}+\hbar\mathbf{k}_{\text{eff}}}$, is much less than the terms oscillating at frequency Δ. Integration of equation (2.15c) can be performed using adiabatic elimination, where the the time dependence of $c_{g,\mathbf{p}}$ and $c_{e,\mathbf{p}+\hbar\mathbf{k}_{\text{eff}}}$ is neglected and oscillations at frequency Δ approximately average out. The results from this integration remove the effect of the intermediate state $|i\rangle$, reducing the system to a coupled two-level system in an external, near-resonant light field [9].

$$\frac{dc_{g,\mathbf{p}}(t)}{dt} = i\frac{|\Omega_1|^2}{\Delta}c_{g,\mathbf{p}}(t) + i\frac{\Omega_1\Omega_2^*}{\Delta}c_{e,\mathbf{p}+\hbar\mathbf{k}_{\text{eff}}}(t)e^{i(\delta t+\phi_2-\phi_1)}, \tag{2.17a}$$

$$\frac{dc_{e,\mathbf{p}+\hbar\mathbf{k}_{\text{eff}}}(t)}{dt} = i\frac{\Omega_1^*\Omega_2}{\Delta}c_{g,\mathbf{p}}(t)e^{-i(\delta t+\phi_2-\phi_1)} + i\frac{|\Omega_2|^2}{\Delta}c_{e,\mathbf{p}+\hbar\mathbf{k}_{\text{eff}}}(t). \tag{2.17b}$$

The atom light interaction causes the energy of each state to change in what is called a light shift, or the a.c. Stark effect [52]. The light shifts cause a change in Rabi frequency leading, to the introduction of Ω_{AC}

$$\Omega_{AC} \equiv \frac{|\Omega_2|^2}{4\Delta} + \frac{|\Omega_1|^2}{4\Delta}. \tag{2.18}$$

Additionally, the difference term δ_{AC} gives a frequency detuning due to the light shifts

$$\delta_{AC} \equiv \frac{|\Omega_2|^2}{4\Delta} - \frac{|\Omega_1|^2}{4\Delta}. \tag{2.19}$$

Note that it is possible to remove the effects due to the light shift by setting the Rabi frequencies of the two to be equal. Practically this can be achieved by fixing the ratio of intensities between the two frequencies [11].

The effective Rabi frequency is also defined, which is the product of the two transition rates and a phase term $\phi_{\text{eff}} = \phi_1 - \phi_2$, obtained from the light interaction, such that

$$\Omega_{\text{eff}} \equiv \frac{\Omega_2^* \Omega_1}{2\Delta} e^{i(\phi_1 - \phi_2)}. \tag{2.20}$$

The solution for the coupled rate equations (2.17), requires the basis to be changed to a rotating frame, in which the Hamiltonian is time independent, so that the typical two level solutions can be applied. In rotation back, a phase term rotating at frequency $\Omega_{AC} \pm \delta$ is introduced. This derivation is done elsewhere [9, 63], and introduces a generalised Rabi frequency, Ω_R, which accounts for light shifts and frequency detuning

$$\Omega_R \equiv \sqrt{\Omega_{\text{eff}}^2 + (\delta - \delta_{AC})^2}. \tag{2.21}$$

The solutions for the coefficients $c_{g,\mathbf{p}}$ and $c_{e,\mathbf{p}+\hbar\mathbf{k}_{\text{eff}}}$ at a time $(t_0 + \tau)$ are

$$\begin{aligned} c_{g,\mathbf{p}}(t_0+\tau) = e^{i(\Omega_{AC}+\delta)\tau/2} \Bigg\{ \left[\cos\left(\frac{\Omega_R \tau}{2}\right) + i \frac{\delta_{AC} - \delta}{\Omega_R} \sin\left(\frac{\Omega_R \tau}{2}\right) \right] c_{g,\mathbf{p}}(t_0) \\ + \left[i \frac{\Omega_{\text{eff}}}{\Omega_R} \sin\left(\frac{\Omega_R \tau}{2}\right) e^{i(\delta t_0 + \phi_{\text{eff}})} \right] c_{e,\mathbf{p}+\hbar\mathbf{k}_{\text{eff}}}(t_0) \Bigg\}, \end{aligned} \tag{2.22}$$

and

$$c_{e,\mathbf{p}+\hbar\mathbf{k}_{\text{eff}}}(t_0+\tau) = e^{i(\Omega_{AC}-\delta)\tau/2}\Bigg\{\left[i\frac{\Omega_{\text{eff}}}{\Omega_R}\sin\left(\frac{\Omega_R\tau}{2}\right)e^{-i(\delta t_0+\phi_{\text{eff}})}\right]c_{g,\mathbf{p}}(t_0) + \left[\cos\left(\frac{\Omega_R\tau}{2}\right) - i\frac{\delta_{AC}-\delta}{\Omega_R}\sin\left(\frac{\Omega_R\tau}{2}\right)\right]c_{e,\mathbf{p}+\hbar\mathbf{k}_{\text{eff}}}(t_0)\Bigg\}. \tag{2.23}$$

Thus through reduction of the three level atom, into an effective two level system, it is possible to describe the time evolution of a state $|\psi\rangle$, in a light field $\mathbf{E}$, driving stimulated Raman transitions.

2.2.3 Rabi oscillations

The coefficients describing the time variance of the states are not directly observable. Instead, the modulus squared of each coefficient gives the probability of finding the atom in that state. The initial conditions are set such that at time t_0, the wavefunction of the atom is in state $|g\rangle$, i.e. $c_{g,\mathbf{p}}(t_0) = 1$ and $c_{e,\mathbf{p}+\hbar\mathbf{k}_{\text{eff}}}(t_0) = 0$. The population of each state at a time $t_0+\tau$ becomes

$$\begin{aligned} P_g(\mathbf{p}, t_0+\tau) &= |c_{g,\mathbf{p}}(t_0+\tau)|^2 \\ &= \cos^2\left(\frac{\Omega_R\tau}{2}\right) + \left|\frac{\delta_{AC}-\delta}{\Omega_R}\right|^2 \sin^2\left(\frac{\Omega_R\tau}{2}\right), \end{aligned} \tag{2.24}$$

and

$$\begin{aligned} P_e(\mathbf{p}+\hbar\mathbf{k}_{\text{eff}}, t_0+\tau) &= |c_{e,\mathbf{p}+\hbar\mathbf{k}_{\text{eff}}}(t_0+\tau)|^2 \\ &= \left|\frac{\Omega_{\text{eff}}}{\Omega_R}\right|^2 \sin^2\left(\frac{\Omega_R\tau}{2}\right) = \left|\frac{\Omega_{\text{eff}}}{\Omega_R}\right|^2 \frac{1-\cos(\Omega_R\tau)}{2}, \end{aligned} \tag{2.25}$$

which shows the population oscillating between the two states, at a frequency

Ω_R. This is the Rabi oscillations expected for two level atoms in a driving field [52]. Due to the choice of long lived states, the population does not radiatively decay, as expected in the traditional two level atom picture, and therefore populations can reach total inversion. As mentioned, the detuning due to light shifts can be nulled; however, more generally the detunings can be grouped such that $\delta_{tot} = \delta_{AC} - \delta$. The effect of changing δ_{tot} on the population probability is explored in figure 2.4. Taking equation (2.25), it can be seen that since the amplitude and frequency of the oscillation are dependent on Ω_R, then as the detuning difference, $\delta_{AC} - \delta$, tends to zero, the maximum probability tends to 1, and the frequency tends to Ω_{eff}.

To relate the Rabi frequency to measurable quantities, equation (2.16) may be written to depend upon the intensity of the light, instead of the electric field, such that

$$|\Omega| = \sqrt{\frac{I}{I_{\text{sat}}}}\frac{\Gamma}{\sqrt{2}}. \tag{2.26}$$

The Rabi frequency has been written in terms of the transition linewidth, Γ, and a saturation intensity, I_{sat}, which is calculated from the coupling strength of the transition and the polarisation [47]. From this, the effective Rabi frequency may be written as

$$|\Omega_{\text{eff}}| = \frac{\sqrt{I_1 I_2}\Gamma^2}{4I_{\text{sat}}\Delta}. \tag{2.27}$$

Thus, the effective Rabi frequency can be increased by increasing the intensity of each Raman frequency, or by decreasing the single photon detuning Δ.

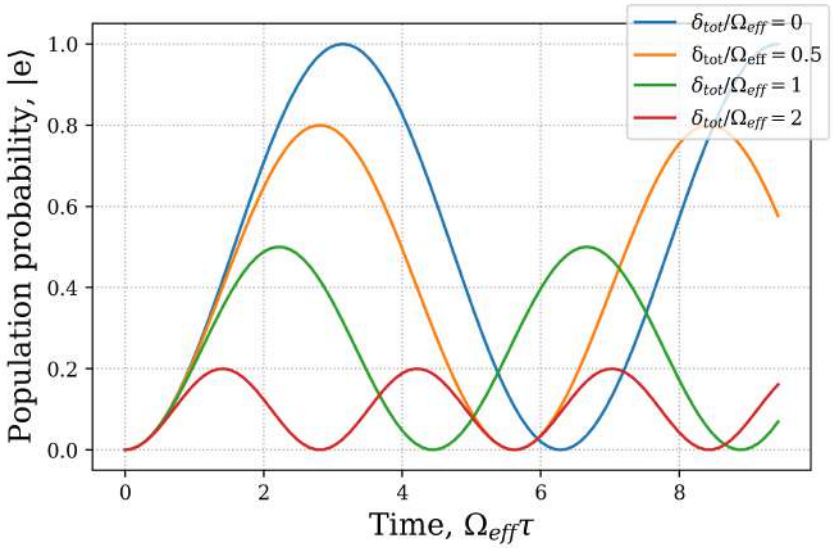

Figure 2.4 Rabi oscillations between two states when the atom starts in state $|g\rangle$. By increasing the total detuning δ_{tot} in equation (2.25), the population probability decreases and the oscillation frequency increases. Since the maximum transfer means a cleaner quantum state, it is important to minimise δ_{tot}.

2.2.4 Atom optics

With predictable behaviour from the atom light interaction, careful choice of light duration will result in the ability to change the state of the atom in a controlled way. The probability to find the atom in the upper state is maximum for a pulse duration $\tau = \pi/\Omega_R$, which for $\delta_{tot} = 0$ corresponds to total population inversion. When the pulse time causes the maximum population inversion, it will be referred to as a π pulse, of length τ_π. Similarly, a pulse of length $\tau = \pi/2\Omega_R$ transfers half the population to the excited state, leaving the atom in an equal mixture of the two states. This will be referred to as a $\pi/2$ pulse, of length $\tau_{\pi/2}$. In the case $|\delta_{tot}| > 0$, a pulse time $\tau = \pi/\Omega_R$ leaves the atom in a mixed state; subsequent atom light interactions compound the effect of the mixed state reducing the resulting fidelity of a pulse scheme.

As well as population transfer, the momentum state is changed due to the coupling of the internal energy and external momentum state. This means that an atom undergoing a π pulse has complete state inversion and the possible trajectories changed. This forms an atom optic equivalent of a mirror, where the efficiency of transfer is analogous to the reflectivity. A $\pi/2$ pulse transfers half of the population, creating two trajectories from the different momentum states. This forms the atom optic equivalent of a beam splitter, where a deviation from 50 % transfer is equivalent to a change in splitting ratio. Using these atom optic beam splitters and mirrors, an atom interferometer can be constructed using combinations of pulse schemes separated in time.

2.2.5 Raman beam configuration

For moving atoms, such as those falling in a gravitational field, an additional resonance condition is introduced due to the Doppler shift. The dependence on velocity comes from the shift $\mathbf{v} \cdot \mathbf{k}_{\text{eff}}$, where the scalar product of the velocity vector, $\mathbf{v}$, and the effective wavevector, $\mathbf{k}_{\text{eff}}$, determine the magnitude and sign of the effect. This shift can be added into the resonance condition of equation (2.4) [63]. Providing the other resonance conditions are met, such that $\delta_{tot} = 0$, then the shift comes into the generalised Rabi frequency as

$$\Omega_R = \sqrt{\Omega_{\text{eff}}^2 + (\mathbf{v} \cdot \mathbf{k}_{\text{eff}})^2}. \tag{2.28}$$

For Raman beams aligned to the z axis, there is a resonant velocity, v_R, for which the shift is cancelled. The included Doppler shift can now be written as $(v_z - v_R)|\mathbf{k}_{\text{eff}}|$. There are two atom velocities relevant to this Doppler shift: the centre of mass motion of the atomic cloud; and the

thermal velocities of the atoms, discussed further in section 2.2.6.

For the atom cloud's centre of mass motion, the configuration of the Raman beams, relative to the velocity of the atom cloud, determines the size of the shift. The Raman beams are always assumed to be co-axial; however, the two frequencies could be delivered from the same direction (co-propagating), or opposite directions (counter-propagating). In the case where the atom motion is perpendicular to the Raman beams, the scalar product of the velocity and the wavevector is zero - the resonant frequency of the Raman transitions are insensitive to the centre of mass velocity of the atom cloud. This may be utilised in cold atom accelerometers and rotation sensors operating in the xy plane [64], where the free-fall of the atoms due to gravity does not shift the resonant frequency of the Raman transitions delivered from beams in the horizontal direction.

As introduced in equation (2.4), the effective wavevector $\mathbf{k}_{\text{eff}} = \mathbf{k}_1 - \mathbf{k}_2$ determines the magnitude of momentum transferred to the atom. In a co-propagating case, the vectors of $\mathbf{k}_1$ and $\mathbf{k}_2$ point in the same direction. The momentum transferred to the atoms is thus $\hbar|\mathbf{k}_{\text{eff}}| = \hbar|\mathbf{k}_1| - \hbar|\mathbf{k}_2|$. In this case, $\mathbf{v} \cdot \mathbf{k}_{\text{eff}}$ is small, making the Doppler shift negligible.

In the counter-propagating case, the vectors $\mathbf{k}_1$ and $\mathbf{k}_2$ point in opposite directions, meaning the transferred momentum is largest, as $\hbar|\mathbf{k}_{\text{eff}}| = \hbar|\mathbf{k}_1| + \hbar|\mathbf{k}_2|$. Whilst this larger momentum transfer is desirable for increasing the interferometer area, the Doppler shift introduces a non-negligible shift to the Raman transition's resonant frequency. The lasers are typically compensated with a linear frequency ramp to stay on resonance with the atom, this is explored later in section 2.4.3.

2.2.6 Velocity selection

Operation in a counter-propagating beam configuration will introduce an additional constraint when considering a cloud of atoms at finite temperature. A cloud of atoms will contain a distribution of thermal velocities causing different Doppler shifts, depending on the velocity. Typically, this means that the resonance condition of the Raman transition cannot be satisfied by all atoms simultaneously. The length of a Raman pulse determines its Fourier linewidth in frequency space, with a long Raman pulse resonant with a narrow velocity class and vice versa.

For an atom starting in state $|g\rangle$, application of a π pulse is resonant only with a select velocity class. It can be shown that the excited state population probability from equation (2.25), for a π pulse can be written as [63]

$$P_e(v_z) = \left[\frac{\pi}{2}\,\mathrm{sinc}\left(\frac{\pi}{2}\sqrt{1+\left(\frac{(v_z - v_R)|\mathbf{k}_{\mathrm{eff}}|}{\Omega_{\mathrm{eff}}}\right)^2}\right)\right]^2 . \tag{2.29}$$

In the case of an atomic ensemble, there is a velocity spread determined by the Maxwell-Boltzmann distribution. Centred around an average velocity, $\overline{v_z}$, the width of the distribution is dependent upon the ensemble temperature, such that $\sigma_v = k_B T/m$. The total distribution can be written as [63]

$$g(v_z) = \frac{1}{\sqrt{2\pi}\sigma_v}\exp\left(-\frac{(v_z - \overline{v_z})^2}{2\sigma_v^2}\right). \tag{2.30}$$

The probability of finding an atom in each state after a π pulse is proportional to the product of the functions $P_e(v_z)$ and $g(v_z)$. To demonstrate the velocity selection of a counter-propagating Raman transition, an ensemble

of 5 µK is modelled in a resonant light field, with an effective Rabi frequency of 50 kHz. This is demonstrated in figure 2.5, where it is seen that only a narrow velocity class from state $|g\rangle$ is transferred by the π pulse.* The total number of atoms transferred by a pulse would be given by the integration of the functions over the velocity. In order to maximise the number of transferred atoms, the ensemble temperature should be reduced and the effective Rabi frequency increased. A trade off between atom number and interferometer contrast will be required, since an ideal interferometer would only address a narrow velocity class, to avoid a broad range of resonant frequencies occurring at the same time. In that case, multiple Rabi oscillations, or interferometer sequences, would occur simultaneously, with the final state output being an average of all paths.

2.3 Zeeman shift

Hyperfine ground-states have degenerate magnetic sub-level structure under zero magnetic field. These magnetic sub-levels known as m_F states can be separated by a magnetic field with an energy shift given by the Landé g_F factor for the low-field regime [47]. This is known as Zeeman splitting, with the change in energy of the magnetic sub-levels given by

$$\Delta E_{|F,m_F\rangle} = \mu_B g_F m_F B_z. \tag{2.31}$$

For the two ground states of ^{87}Rb, this splitting of each m_F state is given by the m_F number times the factor $-0.70\,\mathrm{MHz\,G^{-1}}$, and $0.70\,\mathrm{MHz\,G^{-1}}$ for the $F = 1$ and $F = 2$ states respectively [47]. The effect of the splitting is shown

*It should be noted that while the Raman pulse results in a narrow velocity class, no cooling has taken place. Raman transitions may be used in the cooling scheme of Raman cooling or Raman sideband cooling [65, 66]; however, that is outside the scope of this thesis.

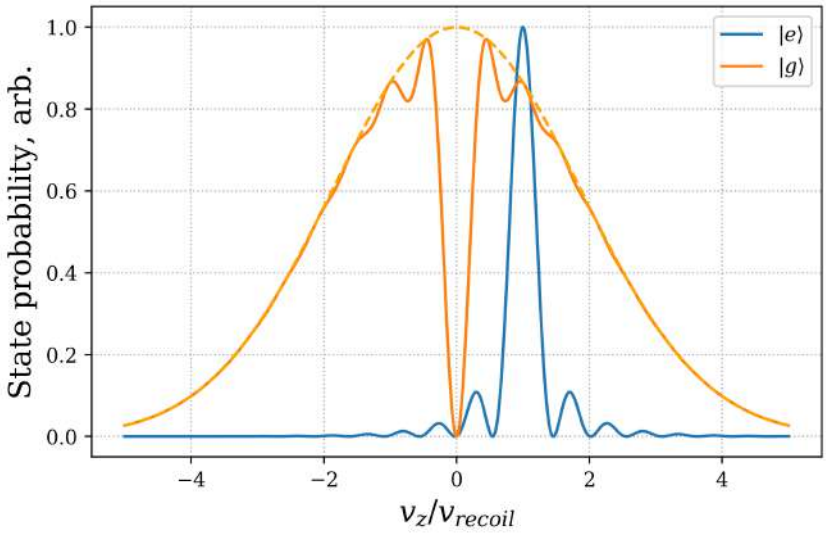

Figure 2.5 Resonant π pulse transfers only a subset of the velocity distribution from state $|g\rangle$ to $|e\rangle$. Dotted line shows the ensemble distribution before transfer. The model is representative of a 5 µK atomic ensemble with an effective Rabi frequency of 50 kHz. The velocity is scaled to the recoil velocity of the two-photon transition, where $v_{recoil} = \hbar|\mathbf{k}_{\text{eff}}|/m$. To transfer a larger fraction of the atoms, the temperature of the ensemble should be reduced, or the effective Rabi frequency increased.

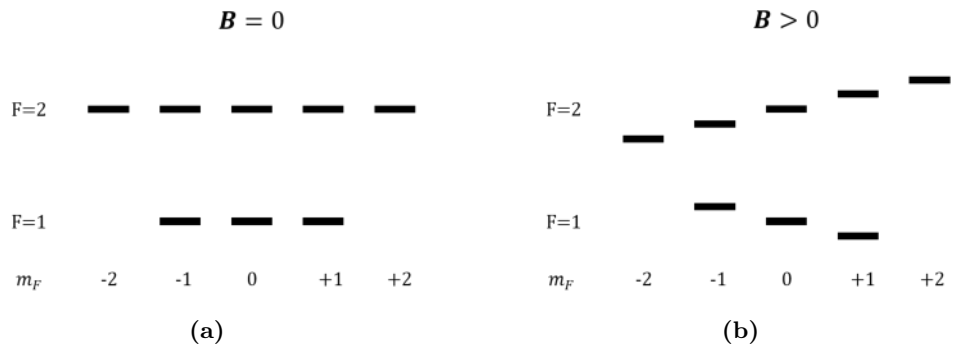

Figure 2.6 (a) With zero applied magnetic field, the m_F sub-levels have degenerate energies. (b) With an applied external magnetic field, the degeneracy in energy is lifted, with the levels splitting by $\mu_B g_F m_F B_z$ in the low-field regime.

in figure 2.6, where the energy of the states is displaced in the presence of a magnetic field.

A Raman transition requires a quantisation axis to define the light polarisation relative to the atoms, dictating the transitions that can be driven, which is achieved with a magnetic field aligned with the Raman beam axis. This bias field may also be used to lift the degeneracy of the magnetic sub-levels, selecting a Raman transition to be resonant only with the transition $|F = 1, m_F = 0\rangle \leftrightarrow |F = 2, m_F = 0\rangle$. States that are not overlapped in energy will not contribute to the interferometer, meaning that the interferometer is insensitive to external magnetic fields to first order. In order for the states to be sufficiently separated, the energy spacing of the states must be larger than the velocity width of the transitions.

To second order, the transition frequency between $|F = 1, m_F = 0\rangle \leftrightarrow |F = 2, m_F = 0\rangle$, is shifted by the square of the field strength, with a shift given by the quadratic Zeeman effect. The value of the shift is derived from the Breit-Rabi formula, quoted here as

$$\Delta\omega_{|F=1\rangle\leftrightarrow|F=2\rangle} = \frac{(g_J - g_I)^2 \mu_B}{2\hbar\Delta E_{|F=1\rangle\leftrightarrow|F=2\rangle}} B^2, \tag{2.32}$$

and may be calculated as $2\pi{\cdot}575.15\,\mathrm{Hz\,G^{-2}}$ [47]. Therefore, the high bias fields required to separate the states will induce shifts to the resonant frequency of the Raman transition. Furthermore, spatial gradients in the field profile will induce position dependent shifts to the resonant frequency, and accelerations to the atoms, proportional to $\mathbf{B}\cdot\mathbf{B}'$, inducing phase shifts to the interferometer [67].

2.4 Atom interferometry

2.4.1 Mach-Zehnder sequence

Using the atom optics introduced in section 2.2.4, an interferometer can be built, analogous to the optical Mach-Zehnder sequence [68]. A typical atom optic Mach-Zehnder sequence uses pulse lengths of $\pi/2-\pi-\pi/2$, where each pulse is separated by a time T. Since the energy and momentum states are coupled in the Raman transitions, the first $\pi/2$ pulse causes the wavefunction to follow two trajectories. The states are allowed to evolve for a time T, after which a π pulse inverts the population, and therefore the momentum states. The two trajectories intercept after an additional time T, where a second $\pi/2$ pulse is used to mix the states together, closing the interferometer. The output of the interferometer is encoded in the population of the two states, and a phase shift between the two arms causes the state populations to oscillate, causing a fringe pattern. At the output of the interferometer, the probability of finding an atom in the excited state is given by

$$P_e = P_0 - \frac{C}{2}\cos(\delta\phi^{tot}), \tag{2.33}$$

where there is an oscillation of amplitude C, centred around an offset value P_0.

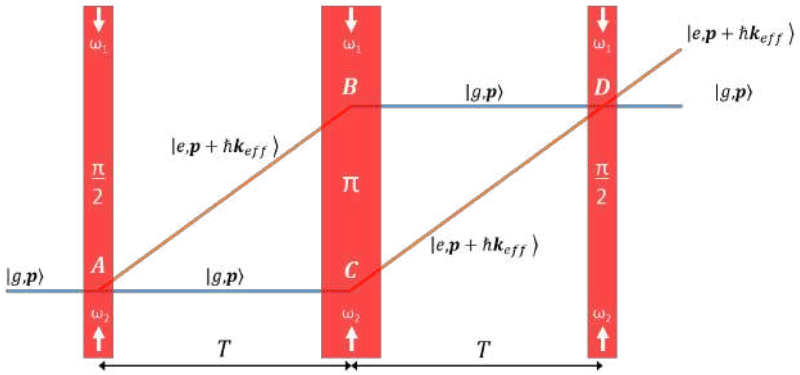

Figure 2.7 Space-time diagram of the interferometry sequence which consists of $\pi/2 - \pi - \pi/2$, separated by an evolution time T. Due to the coupling of the energy and momentum states, the interferometer encloses an area of $\hbar|\mathbf{k}_{\text{eff}}|T^2/m$.

This sequence can be visualised for atoms without a gravitational acceleration in figure 2.7, where the trajectories shown are the classical paths. These classical paths can be calculated from a path integral approach where any path is possible; however, only the paths close to the classically expected trajectories survive the integration [69].

The classical path is determined by the extremal classical action, which can be calculated by integration of the Lagrangian of the system [69]. Through construction of two paths, using the classical actions via different midpoints, the phase shift between the two arms can be calculated. The total phase is simply the sum of all contributions, where a phase contribution from mechanism x is denoted by $\delta\phi^x = \phi^x_{path1} - \phi^x_{path2}$. The contributions for the phase come from: the propagation of the atom, $\delta\phi^{prop}$, the phase from the atom's internal evolution, $\delta\phi^{int}$, and the phase due to the interaction of the laser, $\delta\phi^{laser}$. The total phase of the interferometer is therefore

$$\delta\phi^{tot} = \delta\phi^{prop} + \delta\phi^{int} + \delta\phi^{laser}. \tag{2.34}$$

Since the atom spends the same amount of time in each energy state

Internal state	Momentum	Phase shift
$\|g\rangle \to \|e\rangle$	$\mathbf{p} \to \mathbf{p} + \hbar\mathbf{k}_{\text{eff}}$	$+\big(k_{\text{eff}} z(t) - \omega_{\text{eff}} t - \phi_{\text{eff}}\big)$
$\|e\rangle \to \|g\rangle$	$\mathbf{p} + \hbar\mathbf{k}_{\text{eff}} \to \mathbf{p}$	$-\big(k_{\text{eff}} z(t) - \omega_{\text{eff}} t - \phi_{\text{eff}}\big)$
$\|g\rangle \to \|g\rangle$	$\mathbf{p} \to \mathbf{p}$	0
$\|e\rangle \to \|e\rangle$	$\mathbf{p} + \hbar\mathbf{k}_{\text{eff}} \to \mathbf{p} + \hbar\mathbf{k}_{\text{eff}}$	0

Table 2.1 For transitions between ground and excited states, there is an associated phase shift due to the interaction with the light. The phase shifts depend on the wavevector, k_{eff}, frequency, ω_{eff}, and phase, ϕ_{eff}, of the light field. Additional phase shifts due to the light shifts have been neglected since the contribution can be nulled. Table reproduced from [70]

for both paths, the phase contribution due to internal evolution is zero. It is shown in the appendix B.1, that the contribution of the phase due to propagation is also zero, even in a gravitational field, leaving only the phase due to the atom light interaction.

2.4.2 Laser phase

To calculate the phase contribution from the atom light interaction, consider the four possible interactions the atom can have with the light in each of the two states. With the two level case, the atom will acquire a phase shift when the interaction causes it to change state. Table 2.1 summarises the phase shifts due to the interaction with the light. There are additional phase shifts due to light shifts; however, since these can be nulled, they have been neglected here. A full treatment of phase shifts, including light shifts, can be found in the literature [71].

As with the propagation phase (appendix B.1), the phase difference between two paths can be constructed, this time for the laser interactions. Path ABD, follows the top path of figure 2.7, changing state as follows $|g\rangle \to |e\rangle \to |g\rangle \to |e\rangle$. Path ACD follows the lower path, with the state evolving as $|g\rangle \to |g\rangle \to |e\rangle \to |e\rangle$ in the $\pi/2 - \pi - \pi/2$ sequence. The total

accumulated phase can be calculated for each path from table 2.1 and is

$$\begin{aligned}\phi^{ABD} &= \left[k_{\text{eff}} z_A - \phi_I\right] - \left[k_{\text{eff}} z_B - \omega_{\text{eff}} T - \phi_{II}\right] + \left[k_{\text{eff}} z_D - \omega_{\text{eff}} 2T - \phi_{III}\right] \\ &= k_{\text{eff}}\left(z_A - z_B + z_D\right) - \omega_{\text{eff}} T - \phi_I + \phi_{II} - \phi_{III},\end{aligned} \tag{2.35}$$

$$\begin{aligned}\phi^{ACD} &= 0 + \left[k_{\text{eff}} z_C - \omega_{\text{eff}} T - \phi_{II}\right] + 0 \\ &= k_{\text{eff}} z_C - \omega_{\text{eff}} T - \phi_{II}.\end{aligned} \tag{2.36}$$

The total phase shift from the laser interaction is the difference between the two paths, such that

$$\begin{aligned}\delta\phi^{laser} &= \phi^{ACD} - \phi^{ABD} \\ &= k_{\text{eff}}\left(z_B - z_A + z_C - z_D\right) + \phi_I - 2\phi_{II} + \phi_{III} \\ &= k_{\text{eff}} g T^2 + \phi_I - 2\phi_{II} + \phi_{III},\end{aligned} \tag{2.37}$$

where the construction of an additional parallelogram $A_0B_0C_0D_0$, is used to remove the spatially dependent terms using equation (B.6). This uses the curved paths an atom takes due to the presence of gravity, introducing the term gT^2 from application of the equations of motion (see appendix B.1 for the formalism).

Therefore, the total phase shift in the Mach-Zehnder atom interferometry sequence comes from the interaction with the laser. The phase shift is dependent on acceleration, which for the vertically aligned beams means the interferometer is sensitive to gravitational acceleration.

2.4.3 Interferometer phase

As demonstrated, the interaction with the light provides the phase shift to the interferometer, and the total phase shift is equivalent to

$$\Delta\phi = \phi_{P1} - 2\phi_{P2} + \phi_{P3}, \tag{2.38}$$

where

$$\phi_{Pn} = k_{\text{eff}} z(t) - \omega_{n,eff} t - \phi_{n,eff} \tag{2.39}$$

is the phase contribution from the nth pulse.

For the falling cloud, the resonant frequency will change with time due to the Doppler shift. Adding a linear frequency ramp, α, into the frequency of one of the beams will keep the transition resonant. This frequency ramp is referred to as the chirp rate. The frequency of ω_1 has a time dependence introduced, becoming $\omega_1(t) = \omega_1 + 2\pi\alpha t$. This addition makes the interferometer phase of equation (2.38) become

$$\Delta\phi = (k_{\text{eff}} g - 2\pi\alpha)T^2 + \phi_I - 2\phi_{II} + \phi_{III}. \tag{2.40}$$

With no systematic effects, the phase components from the light cancel, $\phi_I - 2\phi_{II} + \phi_{III} = 0$, leaving the relation

$$\Delta\phi = (k_{\text{eff}} g - 2\pi\alpha)T^2. \tag{2.41}$$

For a value of $k_{\text{eff}} g = 2\pi\alpha$, the phase shift is cancelled for all values of T. A chirp rate scan for different interferometer T times allow for the local value of g to be determined, by finding the point where the fringes have a common minima. A larger interferometer time allows the value of α to be found with

greater precision, since the interferometer area is larger. This is how a cold atom gravimeter measures the value of gravity.

Deriving the interferometer phase in the semi-classical approach does not result in the optimal phase dependence for the interferometer [72]. Although yet to be experimentally demonstrated, it has been theoretically shown that the sensitivity of such an interferometer may see its sensitivity doubled with appropriate modifications [72]. In this work, the interferometer phase shall be assumed to be that of 2.40.

2.4.4 Gravity gradients

The gradient of the gravitational field of Earth is non zero at all points, due to the $1/r^2$ dependence of the field strength. For the field of the Earth, this may be small over length scales of a laboratory system; however, due to variations in local gravity, this effect can become relevant. The two points to consider here are the effect of the gravity gradient on each atom interferometer, and the gradient between two atom interferometers.

Gravity gradient effect on the Mach Zehnder sequence

In the derivation of interferometer phase, the Lagrangian was modelled with a uniform gravitational field. With the inclusion of a linear gradient γ, the Lagrangian becomes [31]

$$L\Big(z, v\Big) = \frac{1}{2}mv(t)^2 - mgz(t) + \frac{1}{2}m\gamma z^2. \tag{2.42}$$

By solving the Euler-Lagrange equation, and following the classical path approach as before, the final gravity value has terms collected in orders of gradient [31]

$$\begin{aligned} g =& g_0 + \gamma\left[\frac{7}{12}g_0T^2 - \overline{v_0}T - z_0\right] + \\ & \gamma^2T^2\left[\frac{31}{360}g_0T^2 - \frac{1}{4}\overline{v_0}T - \frac{7}{12}z_0 - \frac{1}{2}v_{recoil}\big(T - T_{det}\big)\right] + ... \end{aligned} \tag{2.43}$$

Introduction of variables $g_0, \overline{v_0}, z_0$ are made representing the gravitational field, average path velocity, and position, respectively, for an initial time t_0. Evaluating the effect of each order for an interferometer of $T = 100\,\mathrm{ms}$, $g_0 = 9.81\,\mathrm{m\,s^{-2}}$, $\gamma = 3086\,\mathrm{E\,m^{-1}}$ (the gravity gradient of Earth), $\overline{v_0} \approx 0.1\,\mathrm{m\,s^{-1}}$, $z_0 = 0\,\mathrm{m}$, $v_{recoil} = 0.0117\,\mathrm{m\,s^{-1}}$ (two-photon recoil). This puts first order corrections $g^{(1)} \approx 10^{-8}g_0 = 10ng$, which is comparable to a gravimeter signal size. Second order corrections are $g^{(2)} \approx 10^{-16}g_0$, far below the resolution of any current atom interferometer, and can be neglected.

Each interferometer has a dependence on the gravity gradient of Earth, but for manifestations in the gradiometer, this may provide a systematic offset. More important is the dependence on parameters such as initial starting position and cloud velocity, which will manifest as uncommon noise sources in a differential measurement between two interferometers and should be well controlled.

Differential measurement of gravity gradient

When performing a measurement of the gravitational field using a Mach-Zehnder sequence, the measurement of g is a time weighted average of the gravitational field between the initial $\pi/2$ pulse and the final one. Despite this, a gravity gradiometer can be constructed from two such interferometers, separated vertically by a baseline b; a system of this operation is shown in figure 2.8. The field gradient between the two interferometers causes a

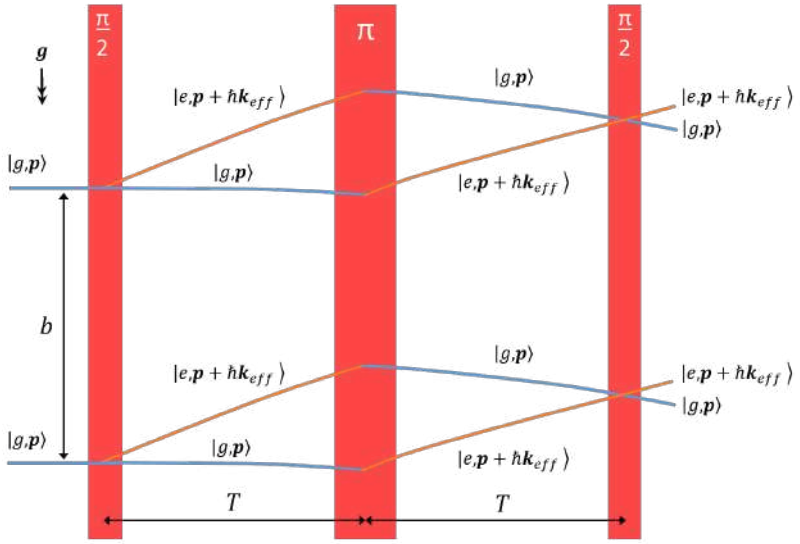

Figure 2.8 Gravity gradiometer using two Mach-Zehnder interferometers separated by a baseline b. Each interferometer measures the time weighted average value of g between the first and last $\pi/2$ pulse. A gradient is taken with the two averaged values of g over the separation. An atom interferometer allows for the measurements to be taken simultaneously.

difference in the output phase for each sensor. To calculate the value of the gravity gradient between the two sensors, a subtraction of the phase can be converted by the equation

$$\nabla g = \frac{\Delta\phi_1 - \Delta\phi_2}{k_{\text{eff}} T^2 b}. \tag{2.44}$$

Consequently, the sensitivity of a gravity gradiometer has a dependence on the separation b, as well as k_{eff} and T. It is this value of gravity gradient that will be used as the measurement parameter.

2.4.5 Sensitivity limit of the atom interferometer

Outside of squeezing, the sensitivity of an atom interferometer is fundamentally limited by the quantum projection noise. Also called atom shot noise, the finite number of atoms in each interferometer cycle creates an

uncertainty in the output phase determined by the statistical fluctuations in atom number. For an ensemble of N uncorrelated atoms, the uncertainty in the number of atoms is given by a Poisson distribution, resulting in $\delta\Delta\phi = 1/\sqrt{N}$.

For a single interferometer, the phase is equal to $\Delta\phi = 2k_{\text{eff}}gT^2$, and the uncertainty in the phase is limited by the quantum projection noise, such that

$$\frac{\delta\Delta\phi}{\Delta\phi} = \frac{\delta g}{g} = \frac{1}{k_{\text{eff}}T^2\sqrt{N}}. \tag{2.45}$$

Extending to a gravity gradiometer, the value of $\delta\nabla g$ may be calculated by propagating the uncertainties associated with $\Delta\phi_1$ and $\Delta\phi_2$ in equation (2.44), by adding the errors in quadrature. The resulting uncertainty in the gradient $\delta\nabla g$, depends on the number of atoms in each interferometer, giving

$$\delta\nabla g = \frac{\nabla g}{2k_{\text{eff}}T^2 b}\sqrt{\frac{N_2 + N_1}{N_1 N_2}}. \tag{2.46}$$

For similar atom numbers in each interferometer, the approximation $N_1 \approx N_2 = N$ can be made. The resulting atom shot noise limited sensitivity of the gravity gradiometer, the uncertainty over the measured value, is thus given by

$$\frac{\delta\nabla g}{\nabla g} = \frac{\sqrt{2}}{2k_{\text{eff}}T^2 b\sqrt{N}}, \tag{2.47}$$

for a single shot. To calculate the sensitivity, in units of $\text{E}/\sqrt{\text{Hz}}$, the measurement time for a single shot is included. Taking a repetition rate of 0.7 Hz, and an atom number of 10^5, the atom shot noise limited sensitivity of the interferometer at $T = 85\,\text{ms}$ is $23\,\text{E}/\sqrt{\text{Hz}}$. In order to be limited by

atom number fluctuations, technical noise must be lower than the quantum projection noise.

2.5 Phase components in a gradient measurement

The gravity gradient is calculated from the differential phase $\Delta\Phi$, between the two interferometers. In a system with no phase noise, extraction of the differential phase can be achieved most simply by plotting each interferometer fringe and determining the phase shift between the two. In the presence of phase noise, the interferometer fringes are completely scrambled, leaving no way to determine the phase of each fringe.

Denoting the two phases as being composed of common mode (CM), and non-common mode (NCM), phase components and phase noises, the phase shift between the two interferometers is given by

$$\begin{aligned}\Delta\Phi = \Delta\phi_1 - \Delta\phi_2 &= \phi_1^{CM} + \delta\phi_1^{CM} + \phi_1^{NCM} + \delta\phi_1^{NCM} \\ &\quad - \left(\phi_2^{CM} + \delta\phi_2^{CM} + \phi_2^{NCM} + \delta\phi_2^{NCM}\right) \\ &= \phi_1^{NCM} + \delta\phi_1^{NCM} - \phi_2^{NCM} - \delta\phi_2^{NCM} \\ &= \phi_{\nabla g} + \phi_{systematic} + \delta\phi_{noise}.\end{aligned} \tag{2.48}$$

The common phase components, including the noise terms $\delta\phi^{CM}$, have been cancelled, leaving only the non-common terms. The phase term, $\phi_1^{NCM} - \phi_2^{NCM}$, has been grouped into terms responsible for the gravity gradient $\phi_{\nabla g}$, and systematic phase offsets $\phi_{systematic}$. Phase noise terms, $\delta\phi_1^{NCM} - \delta\phi_2^{NCM}$, have been grouped into a single uncommon noise term $\delta\phi_{noise}$.

Common phase noise terms can come from within the sensor or external effects. Typical terms will come from the effect of vibration and microseismic noise, and Raman laser phase noise. Provided the noise source af-

fects the phase noise of each interferometer in the same way, the gradient measurement will remove the noise. Most noise sources will have an element of common and non-common noise; it is important to match the conditions of each interferometer carefully, such that the noise sources are made to be common. Practically this means rigid fixing of the two interferometers with a common, aligned beam axis.

Systematic phase shifts may be caused by differences in the two interferometers, such as initial atom cloud velocities, or gradients in the local magnetic field. Provided they remain constant for all measurements, systematic phase shifts do not need to be suppressed to achieve the required sensitivity for civil engineering applications, where accuracy is not being targeted. However, systematic phase shifts which have a dependence on external environmental effects should be removed, lest they skew the measurement of gravity gradient.

Phase noise in the gradiometer will couple in with different timescales and through different effects. Since the output of the interferometer is encoded in the state populations, anything that causes fluctuations to the measured populations will result in phase noise. Additionally, in order to resolve interference fringe of each interferometer, the contrast, the percentage of each atom transferred, should be maximised. Therefore, it is important that:

1. The atoms are prepared in a clean state, ideally with a consistent number of atoms.

2. The Raman transitions are driven repeatably and with a high fidelity for each interferometer, transferring the correct fraction of atoms for pulses within the interferometer sequence, and between subsequent

measurements.

3. The state population detection scheme is independent of fluctuations in atom number and provides a proportional and accurate detection of the population of atoms in each state.

During the construction and optimisation of the instrument, sub-systems and experiment steps are evaluated against these criteria to ensure phase noise and time dependent phase shifts are not introduced to the measurement of gravity gradient.

2.6 Extracting gravity gradient using ellipse phase

In order to remove the common mode phase components and extract the gradiometer phase shift, the output fringes from the two interferometers are plotted parametrically to create an ellipse, demonstrated in figure 2.9. For a fixed phase difference between the two interferometers, an ellipse can be formed, even in the presence of common phase noise between the two interferometers. The common phase noise can be any value, including greater than 2π, without any effect on the output ellipse. The phase of the ellipse can be extracted to determine the phase shift between the two interferometers, without the need to reconstruct the interferometer fringes, through an ellipse specific least squares fitting approach [73].

The extraction of the differential phase may be performed by fitting to the ellipse using other fitting techniques for a modest improvement to accuracy [74], or through cancelling the phase with an applied phase shift [75]. Alternative proposals include: reconstruction of the fringes through correlation between the interferometers and a classical accelerometer to remove large phase offsets [76]; and phase shear, where the entire interferometer

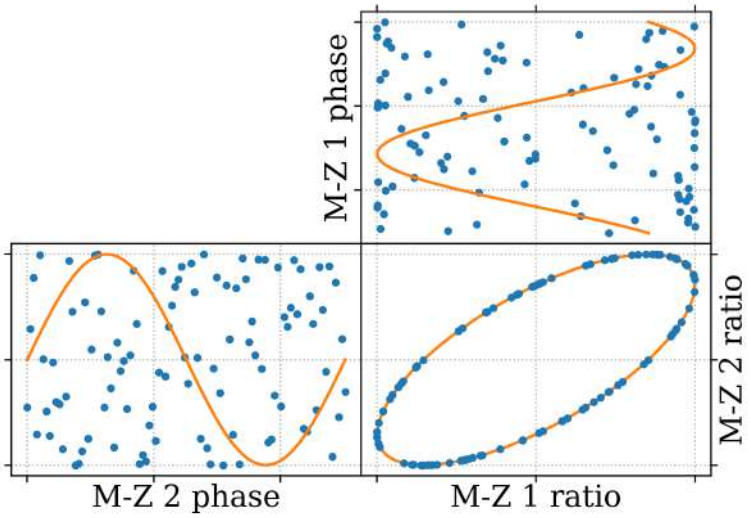

Figure 2.9 For two Mach-Zehnder interferometers operating simultaneously, the fringes, shown with orange lines, may be plotted parametrically to produce an ellipse. The differential phase shift between the two interferometers is given by the phase of the ellipse. In the case of common phase noise, as shown with the blue points, the scrambled interferometer fringes recover the same ellipse phase, with the noise removed.

fringe is imprinted on the spatial extent of the atom cloud [77].

Until performance of the instrument is limited by the error introduced by the fitting algorithm, ellipse specific least squares fitting will be used. The alternative methods may be considered in future upgrades to the system. The details of the ellipse specific fitting are covered later in section 4.3.

2.7 Coriolis force

As a direct result of the Sagnac effect, any horizontal velocity of the atomic clouds will result in a phase dependence of the interferometer to the orientation relative to the East-West axis [31, 78].

For a single atom interferometer, the phase shift is proportional to the cross product of the velocity in the East-West axis, $\mathbf{v}_{EW}$, and the momentum imparted upon the atom due to a Raman pulse, $\mathbf{k}_{\text{eff}}$. The phase shift on the interferometer is given by

$$\Delta\phi_{cor} = 2\boldsymbol{\Omega} \cdot (\mathbf{v}_{EW} \times \mathbf{k}_{\text{eff}})T^2, \tag{2.49}$$

which can also be expressed in terms of g and the recoil velocity, $\mathbf{v}_{rec}$,

$$\Delta g_{cor} = 2\boldsymbol{\Omega} \cdot (\mathbf{v}_{EW} \times \mathbf{v}_{rec}). \tag{2.50}$$

The bias to the interferometer phase is scaled by $\boldsymbol{\Omega} = \Omega\cos(\theta_{lat})$ - the Earth's rotation angular velocity, which is largest at the equator, where $\theta_{lat} = 0$. Rewriting the vectors in terms of latitude, angle of orientation, θ_{or}, and horizontal velocity of the atoms, v_{hor}, equation (2.50) becomes [78]

$$\Delta g_{cor} = 2\Omega\cos(\theta_{lat})v_{hor}\cos(\theta_{or}). \tag{2.51}$$

In the gradiometer, a simple extension to this formula can be made by taking the gradient of two sensors with horizontal velocity $\mathbf{v}_{hor1}$ and $\mathbf{v}_{hor2}$. Using equation (2.44), the Coriolis effect introduces a shift to the gravity gradient, equal to

$$\nabla g_{cor} = \frac{2\Omega \cos(\theta_{lat}) \cos(\theta_{or})}{b} |\mathbf{v}_{hor1} - \mathbf{v}_{hor2}|. \tag{2.52}$$

In the case of gravity gradient, any contribution due to Coriolis is cancelled when the relative velocity of the two clouds is zero, or aligned with the East-West direction. Consequently, rotation of the sensor about the vertical axis will reveal if there is a differential velocity and the direction it points in. Setting the differential velocity as Δv_{hor}, for different values of the differential velocity, the expected shift can be large for moderate velocities, as seen in figure 2.10. At the latitude of the University of Birmingham (52.45 deg), changes in the ellipse phase can be 10's of Eötvös per rotational degree.

2.8 Initial demonstrator

The instrument described within this work takes lessons learnt from the development of a previous gravity gradiometer, discussed in detail in reference [10]. The previous instrument was used to evaluate the core design and robustness of a sensor package suitable for measuring gravity gradients in a civil engineering environment, focusing on assessing the challenges of sensitive operation in the field.

The instrument consists of a sensor and control package, separating the atomic part (vacuum system, atom preparation, interrogation, and detection), from the supporting hardware (lasers, electronics, and computer con-

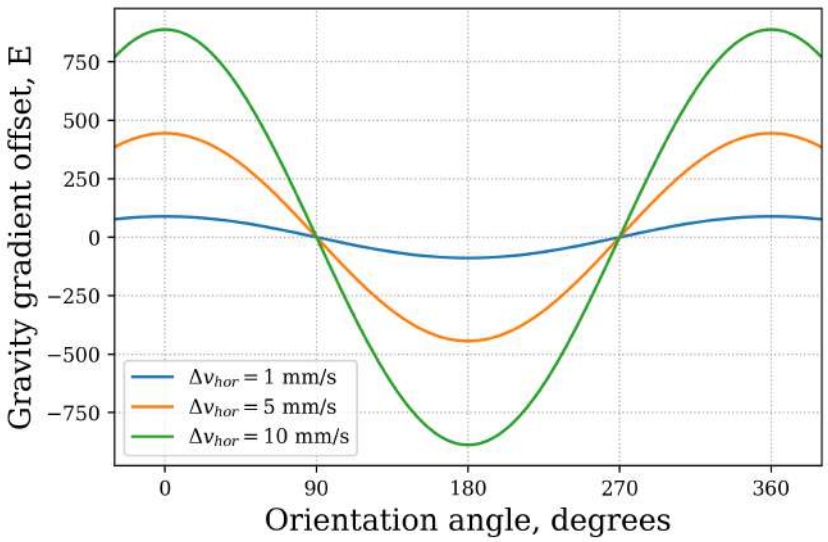

Figure 2.10 Modelled gravity gradient shift due to Coriolis force for different differential horizontal velocities. The sensitivity to angle can become large for moderate differential velocities. In a measurement environment, good survey practice will reduce errors introduced due to changing instrument alignment.

trol). The two are joined by an umbilical to route electrical cables and optical fibres, in order to operate with a static control package and person portable sensor package.

Assessment of the performance of the instrument was done during an overnight measurement campaign in a cave. A cave environment provides a stable temperature and low background magnetic fields, although humidity can be exceptionally high. The instrument ran in a gradiometer configuration, with interferometer T time of 50 ms and baseline 0.75 m, measuring gravity gradient in a steady state for 10 hours to determine the stability of the phase measurement. The data was analysed by fitting ellipses to every 25 shots, extracting the phase value for each ellipse, using the method described later in section 4.3. With the phase data, an Allan deviation was taken (appendix B.2), shown in figure 2.11, where the instrument shows an average short term sensitivity of $(2857 \pm 76)\,\mathrm{E}/\sqrt{\mathrm{Hz}}$, at a 0.5 Hz measure-

ment bandwidth. The averaged value of gravity gradient for the instrument, shown in blue, quickly diverges from the value expected when averaging the Poisson distributed phase values, $1/\sqrt{N}$ shown in orange, suggesting the instrument is limited by instability, i.e. drifts in the phase value. Comparing the resolution of the sensor, estimated from the Allan deviation at 175 E, with some of the expected signal sizes from civil engineering, outlined in figure 1.2, targets on the scale of large sinkholes could be addressed with such an instrument.

Using equation 2.47, the atom shot noise limited sensitivity, for the operating parameters above, is calculated $209\,\mathrm{E}/\sqrt{\mathrm{Hz}}$ for an estimated 10^5 atoms participating in each interferometer. Simulation of the data suggests that the short term sensitivity of the instrument is limited to $1100\,\mathrm{E}/\sqrt{\mathrm{Hz}}$ due to noise associated with detection, notably Johnson noise and shot noise due to background light [79]. Addressing the technical noise of the instrument was less feasible than the building of a new sensor package, which may be designed to overcome the limitations identified in the initial demonstrator. Many limitations occurred due to uncommon phase noise between each interferometer, many of which were caused by the coupling to environmental effects, increasing instability and short term noise in the instrument.

2.9 Design considerations for a new sensor package

In order to achieve a lower sensitivity and be able to resolve smaller civil engineering features, a new sensor package is designed for use with the existing control package. The instrument will again be split into a control/sensor package architecture, keeping the laser system, RF chain, electronics, and

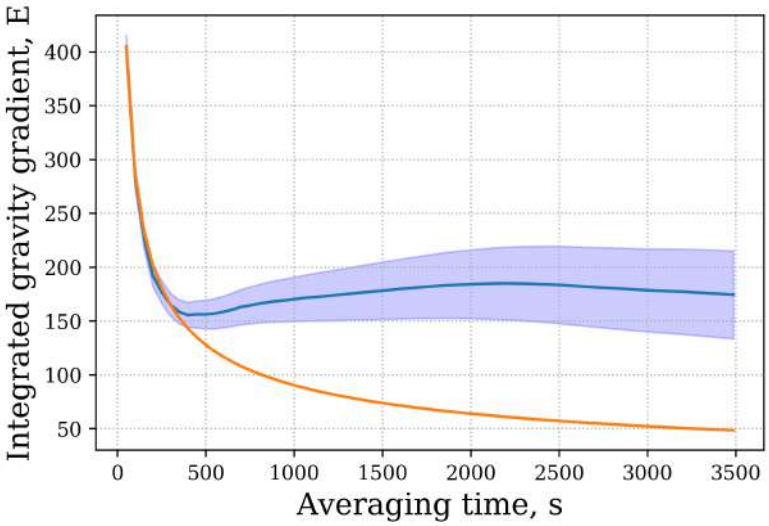

Figure 2.11 Allan deviation (blue line), taken using the first iteration instrument. Shaded area shows the error of the Allan deviation (see appendix B.2). An overnight measurement was made in a cave to ensure stability for temperature and background fields. Average short term sensitivity is measured as $(2857 \pm 76)\,\mathrm{E}/\sqrt{\mathrm{Hz}}$, for a 0.5 Hz measurement bandwidth, with the orange line $1/\sqrt{\tau}$ averaging. The lowest average gradient value is 155 E in 398 s before instability dominates.

computer control separate from the atom cloud generation, interrogation, and detection. The two are linked with optical fibres and electrical wires, bundled into an "umbilical" to deliver light and electrical power to the sensor package by a single cable assembly.

The considerations for a new sensor package are outlined here, with particular emphasis on shielding the atoms from environmental effects to reduce uncommon phase noise between the interferometers. For the instrument to operate out of the laboratory, there are additional constraints that require a level of system integration not usually required in cold atom systems. Such constraints tend to trade off ease of optimisation of the system for a compact form factor and mechanical robustness, in order to increase the overall portability of the instrument.

In order to determine the desirable sensitivity of the instrument, some simple modelling of gravity anomalies is performed. From this, the parameter space can be explored in order to set the baseline and interferometer T time.

2.9.1 Target modelling

A target can be modelled by approximating its shape to a simple well understood gravity anomaly. In this case, a tunnel or pipe is modelled as an infinitely long cylinder with a density contrast set to be that of air in soil. Whilst this model is a simplification to a real anomaly, it provides an order of magnitude calculation to determine signal sizes. More complex models can be made to include: soil density variations, terrain corrections, and time dependent environmental noise [17].

For this model, the gravitational acceleration from the anomaly is calculated at two positions, separated vertically by a variable baseline b. Using

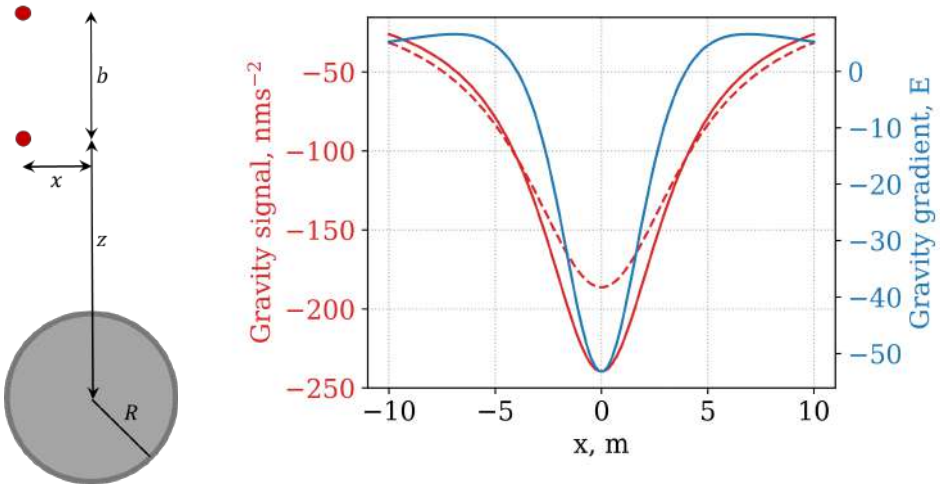

(a) **(b)**

Figure 2.12 a) An infinite cylinder anomaly measured by two sensors. b) Expected gravity signal for top (dashed), and bottom (solid), sensor when z =3.5 m and a 1 m baseline, where the radius is 1 m and a density contrast of soil to air ($-2000\,\mathrm{kg\,m^{-3}}$). The gradient, shown in blue, is calculated from the gravity signals and the baseline. The signal shape from the gradient is narrower than the gravity signals in this case, demonstrating how a gravity gradient measurement also complements a single height gravity map, by providing additional information.

the values of gravitational field at the two points and the baseline between them, the gravity gradient is calculated. The equation

$$g_z(x) = 2\pi G R^2 \Delta\rho \frac{z}{x^2 + z^2}, \tag{2.53}$$

can be used to calculate the gravity signal from an infinite cylinder of radius R, at a distance $\sqrt{x^2 + z^2}$ [80]. The density contrast between the cylinder and surrounding material is modelled as dense soil/clay to air ($-2000\,\mathrm{kg\,m^{-3}}$) [80]. The layout, gravity, and gravity gradient, as calculated for each sensor position, is depicted in figure 2.12.

2.9.2 Determining baseline and noise limits

Comparing the gravity gradient to the quantum projection noise limit (section 2.4.5), gives a signal to noise ratio, for an atom shot noise limited system. Changing the baseline and the interferometer T time in the calculation of the gravity gradient for the modelled anomaly, allows exploration of the parameter space the instrument can take. For the anomaly modelled in section 2.9.1, the expected signal to noise ratio is shown in figure 2.13. In order to provide sensible limits for the scanned parameters, the maximum modelled interferometer T time is taken to be half the evolution time of a comparable lab system, quoted as operating with background vapour loaded atomic clouds, cooled to a few µK [10, 11, 67, 81]. This evolution time is taken as 500 ms.

The interferometry region dimensions are evaluated for both an atom drop and launch configuration, setting the maximum interferometer T time in each case. These limits are set to a maximum T time of 230 ms in a launch configuration, and a maximum T time of 100 ms in a drop configuration. To determine the baseline of such an instrument, the physical limitations of

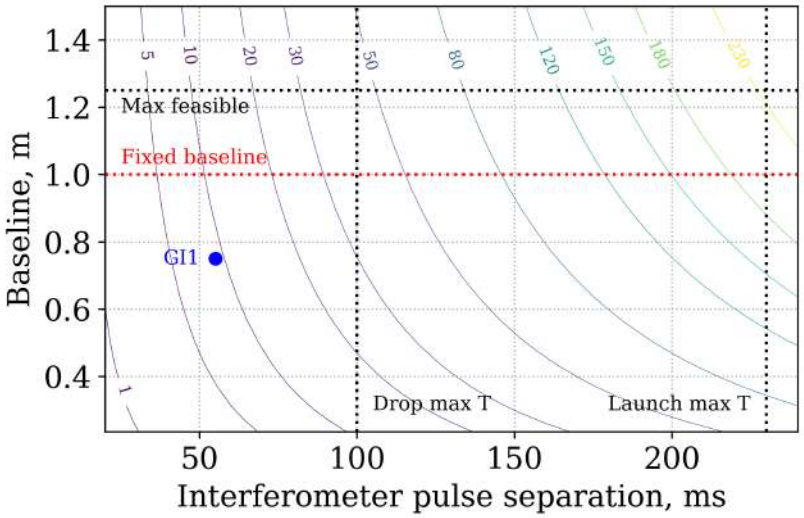

Figure 2.13 For the modelled anomaly in figure 2.12, contour plot showing the signal to noise ratio for an atom shot noise limited gradiometer, when the baseline and interferometer T time is changed. To enable the instrument to leave the lab, a maximum baseline is shown as well as the baseline chosen for the sensor. In drop or launch configuration, vertical lines show the maximum evolution time limited T time. For completeness, the operating T time and baseline of the first iteration sensor is included (GI1).

such a system must be considered. In this case, the instrument is to operate in the field environment so must be able to leave the laboratory, fit into the lift, and into a suitable transport vehicle. From the modelling shown in figure 2.13, the baseline of the instrument is set to be 1 m to allow space for optical access and magnetic shielding.

Operating the instrument in a drop configuration, at a reduced interferometer T time of 85 ms, a signal to noise ratio of roughly 25 is expected for the modelled anomaly. If technical noise, or noise introduced by operation in the field, increases the noise to 10 times the atom shot noise limit, then it should still be possible to detect a target with a signal to noise ratio of 2 or better. Therefore, in order to measure an interesting feature without significant averaging, the technical noise should be within a factor of 10 of the atom shot noise limit.

2.9.3 Environmental phase shifts

For the instrument to operate successfully in the field, there must be sufficient protection from environmental conditions, especially where these may produce uncommon phase shifts. One main environmental challenge when using atoms comes from the effect of external magnetic fields. These can introduce phase shifts into the interferometer, and change the trapping centre of the MOT, adding cloud position instability and additional phase shifts. To mitigate these, a magnetic environment which shields the atoms from external fields is essential.

Since in principle the gradiometer is less sensitive to tilt variations, wind noise will not affect the gradiometer in the same way. Uncommon phase shifts will still be introduced into the gradient measurement if the beam path between the two interferometers is air, due to two identified issues:

phase shifts due to air pressure changes, and wavefront aberration from the required vacuum windows. Whilst the gradiometer is in principle unaffected by changes in air pressure, the refractive index of air is changed, introducing a phase shift to the interferometer [82]. Additional noise due to turbulent air, caused by wind or sensor motion, will therefore contribute to uncommon phase noise in the gradient measurement. To mitigate this, continuous vacuum should link the two interferometry regions, which also removes two uncommon optical interfaces, reducing phase shifts from Raman beam aberration [83].

Detection of the atom population is performed optically by measurements of fluorescence. Therefore, changes in ambient light level will contribute to noise in the detection system. It will be necessary to completely isolate the system from any external light sources.

2.9.4 Cold atom requirements

MOT geometry

Fluctuations in polarisation or intensity in the trapping beams of a magneto-optical trap will result in fluctuations in atom number and cloud position stability. One mitigation for this is to move to a trap design where the trapping beams are derived from a single input beam. Fluctuations in laser intensity or polarisation become common between trapping beam pairs, reducing the effects of trap centre motion and atom number fluctuations [10]. A particularly robust geometry is that of the prism MOT [84, 85], which consists of four mirror coated prisms redirecting the light by 45° towards the trap centre, creating two of the trapping beam pairs. The third trapping beam pair is formed through either reflection or a separate beam delivered from the bottom. The prism MOT scheme is shown in figure 2.14. Also

shown on each beam is the $\sigma^{\pm}$ transitions required for the MOT, chosen by the circular polarisation of the input light.

The prism MOT trades off optical efficiency for overall stability. Other novel trap designs exist, including metasurface beam splitters [86], or gratings [87], which may offer better use of optical power when creating the MOT. Geometries such as a grating MOT restrict the optical access around the cloud, reducing the potential Raman beam size used to couple the two atom clouds.

Two prism MOTs may be placed back to back to create two trapping centres, with the through beam forming the rear beam of the other trap. In this configuration, trapping beams may be delivered from the top and bottom to create two trapping centres in order to generate two atom clouds. Having a common vertical trapping beam pair also helps to couple any vertical centre of mass motion of the atom clouds, reducing changes in the gradiometer baseline. This idea is shown in figure 2.15.

Atom cloud temperatures

Thermal atoms are sufficient to realise the performance of this instrument, which can be achieved using an optical molasses, a sub-Doppler cooling scheme [52, 53]. In an atom interferometer, cold atoms are required to reduce the velocity distribution of the atomic cloud, in order to reduce dephasing in the interferometer, and to allow the Raman transitions to address a significant fraction of the atoms for a particular velocity class. Due to a finite Raman beam size, the Raman beam profile has a Gaussian intensity profile, creating a different Rabi frequency for atoms in different parts of the beam. For an atomic ensemble within the Raman beam, the spatial overlap creates inhomogeneous Rabi oscillations when the the sum of all

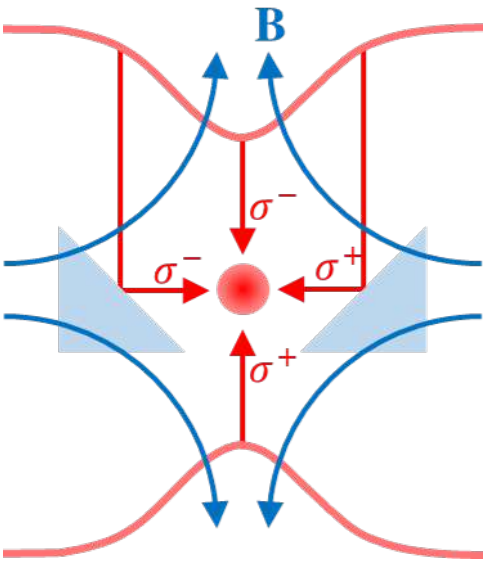

Figure 2.14 Operation of a prism MOT. Blue arrows denote magnetic field lines produced by a pair of quadrupole coils. When illuminated by a Gaussian beam from both sides, the reflections create the correct circular polarisation required by the atoms to scatter preferentially. An additional pair of prisms (not shown) form the trapping beams in the plane of the page.

atoms is taken [88]. To properly model the average Rabi oscillation form, a multi-frequency model is required to account for the different Rabi frequencies present. In a simplified case, a damping coefficient can be used to approximate this behaviour for a single Rabi frequency.

For sequential Raman pulses, the inhomogeneous transfer effect is compounded and produces a reduction in interferometer contrast. However, with an expanding cloud, an atom sees a changing Rabi frequency for each interferometer pulse, creating further dephasing of the interferometer. This can be overcome using a composite pulse scheme [89], but typically the effect is reduced through use of a large beam size compared to the atomic cloud size. With a fixed beam size, it will be important to keep the atomic ensemble temperature low, to reduce cloud expansion over the evolution time. With an optical molasses, ensemble temperatures can typically reach a few µK,

when the background magnetic environment is nulled [52, 53]. Atoms of single µK will only move a few mm from their starting position over the length of the interferometer, which will be the targeted ensemble temperatures.

2.9.5 Sensor package design overview

Taking the identified issues with the previous sensor into account, the design of the new sensor package has the following design objectives: an increased interferometer T time, a uniform bias field, continuous vacuum between interferometers, enclosed magnetic shielding, and low detection noise. The sensor package design will take inspiration from the original sensor, where robustness has been demonstrated, and will use a pair of prism MOTs sharing a common Raman beam. The constraints of size and weight have been relaxed for this sensor, provided it is still portable and usable in the field environment. For operation without specialised infrastructure, the complete instrument should be operable using a 230 V, 13 A standard power outlet, and be portable with only a few operators, without the need for additional lifting equipment. The outline of the sensor head requirements is included in figure 2.15, and the prism MOT geometry is shown.

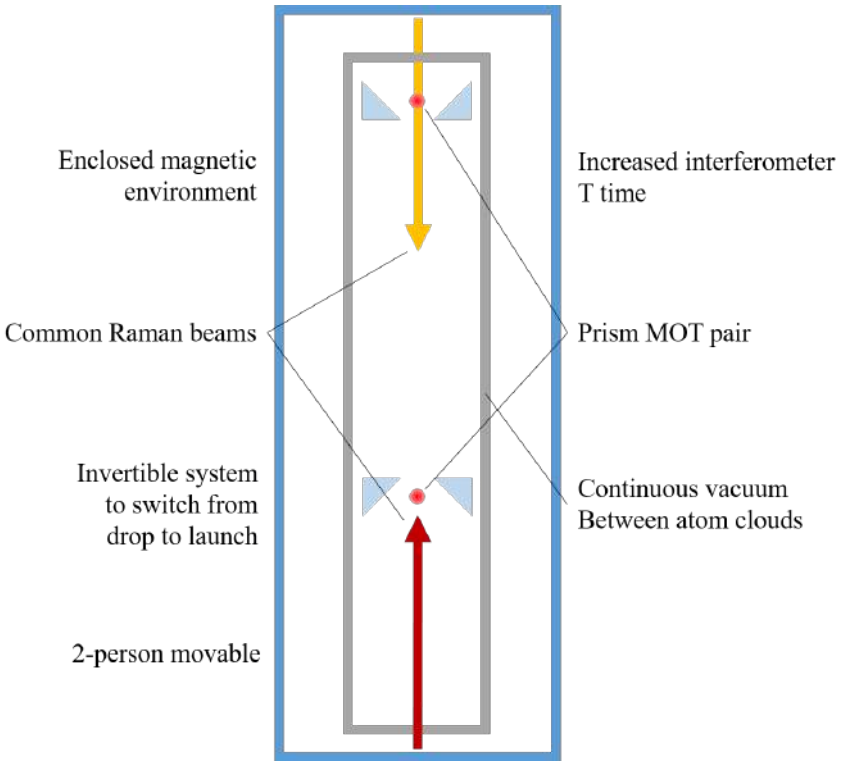

Figure 2.15 Requirements of the new sensor package to maximise the sensitivity and usability of the instrument. Back to back prism MOT geometry shown, where they share a common through beam pair.

Chapter 3

Sensor package

Here the design, manufacture, and assembly of the new sensor package is presented, meeting the coarse dimensions set by the modelling of the parameter space in section 2.9. The sensor package has been designed to address the issues identified in the initial demonstrator, as well as include new features to allow for upgrades at a later date. To maximise the sensitivity of the instrument, the sensor head is designed to be invertible, allowing for a maximised interferometer T time in either drop or launch configuration. As discussed in section 2.9, the core design is based on two vertically separated prism MOTs who share the same Raman beam axis.

The sensor package is where the atoms are prepared, interrogated, and detected. To enable this, the required sub-systems are: a cold atom source, vacuum environment, magnetic environment, laser beam delivery, detection optics and electronics. These will all be designed to reduce uncommon phase shifts and phase noise between the two interferometers.

3.1 Vacuum chamber

Preparation, interrogation, and detection of atoms is all achieved in an ultra-high vacuum (UHV) environment. This is required in all cold atom experiments to reduce heating and decoherence due to collisions with background atoms. The vacuum system may be considered as two halves consisting of a science chamber and an interfacing module, connected by continuous vacuum. The science chamber is where atoms are prepared, interrogated, and detected; it contains the prisms for the magneto-optical trap as well as all the optical access. The interfacing modules contain the pumps, valves, and feedthroughs as well as the nipples used for continuous vacuum between each half. To facilitate switching the system from an atom drop to launch configuration, the vacuum system is largely symmetric in its design, with the naming convention "Top" and "Bottom" used to define each science chamber relative to the direction of gravity.

To facilitate upgrades to the system, as well as minimise maintenance, the vacuum is designed to contain as few components as possible and no moving parts. All beam delivery is external to the vacuum, using viewports for optical access.

3.1.1 Vacuum pressure requirement

To calculate the required pressure, the inelastic collision rate r_c, can be calculated to determine the collision of atoms in the interferometer with background atoms [9]. This is determined by the average velocity, $\langle v \rangle$, of the background atoms, and the mean free path of an atom, l. Expressing the mean free path in terms of the number density n, and the atomic collision cross section σ_c, the collision rate is given by

$$r_c = \frac{\langle v \rangle}{l} = \sqrt{2} n \sigma_c \langle v \rangle. \tag{3.1}$$

The average velocity of thermal atoms is given by the Maxwell-Boltzmann distribution, $\langle v \rangle = \sqrt{8 k_B T / \pi m}$, expressed in terms of the mass m, and temperature T, of the atoms. The number density can be related to pressure by the ideal gas law such that $n = p/k_B T$. Assuming that the collisions are limited to Rb-Rb, the collision rate can be written as

$$r_c = P \sigma_c \sqrt{\frac{16}{\pi m_{Rb} k_b T}} \tag{3.2}$$

In order to use a background vapour loaded prism MOT, the system must be designed such that collisions in the vacuum chamber are dominated by inelastic scattering with background Rb atoms. An approximate atomic collision cross section for background atoms would be $3 \cdot 10^{-18}\,\mathrm{m}^2$ [90]. For cloud evolution times on the order of half a second, the collision rate should be on the order of $2\,\mathrm{s}^{-1}$, putting an upper bound on the pressure P of $1 \cdot 10^{-7}\,\mathrm{mbar}$, calculated using equation (3.2). In order for Rb-Rb collisions to dominate, the partial pressure of other background gases should be a few orders of magnitude lower, requiring ultra-high vacuum pressures for the system before Rb is introduced.

3.1.2 Science Chambers

At a minimum, the science chamber needs to contain: the optics and atom source to produce a prism MOT, a distance for the atoms to drop or be launched, and optical access for beam delivery and detection. The prism MOT determines the minimum diameter of the science chamber, with 15 mm prisms used to accept a 2" cooling beam. The internal diameter of the

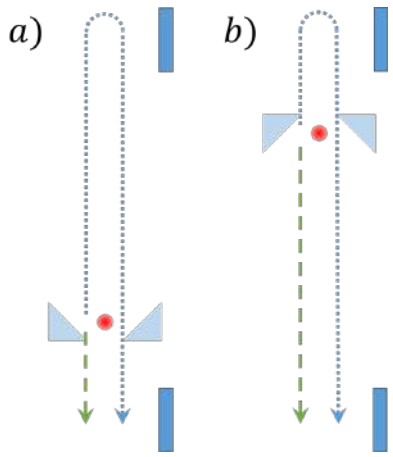

Figure 3.1 Extremes of cloud trajectories for a) launch, and b) drop configuration for one of the science chambers. Dotted blue lines denote an atom launch, dashed green lines for an atom drop.

vacuum chamber must therefore be a minimum of 45 mm to fit the trap.

Since the instrument will be operated using the existing control package, an atom launch will not initially be possible. The vacuum chambers have been designed with this in mind to allow for hardware testing of an atom launch without needing to invert the sensor. Consequently the trapping region is located off centre, with a detection region above and below it. Each detection region is formed of four windows, in a square pattern in the xy plane. Having two distinct detection regions allows for four modes of operation, which is demonstrated in figure 3.1. Initial launch tests will likely see case 3.1b) where a small atom launch can be performed with the sensor in a drop configuration. With four windows at each level, there is optical access for multiple diagnostic equipment and detection optics. Around the trapping region are an additional set of windows to allow for diagnostic photodiodes and cameras.

The final design of the science chambers can be seen in figure 3.2. Manufacture of the science chambers was done in-house by the Physics Workshop

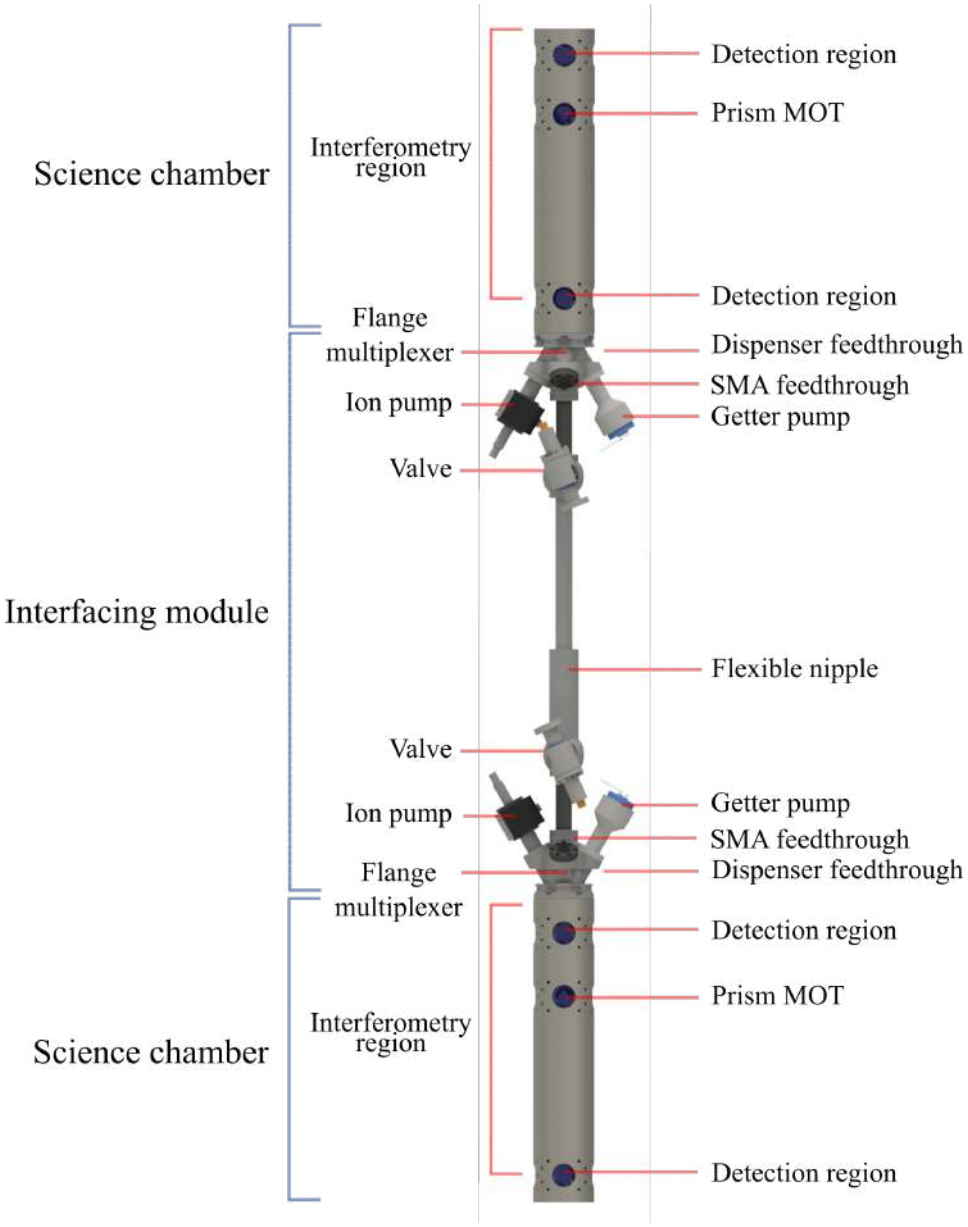

Figure 3.2 Vacuum chamber in drop configuration. Each custom machined science chamber connects to an interfacing module, attaching pumps and electrical feedthroughs. Separating the two halves is continuous vacuum constructed from solid and flexible nipples.

staff, from a titanium billet, using a CNC lathe and 5-axis machine. Titanium is used due to its low outgassing rate and low magnetic susceptibility [91]. Additionally, titanium has a thermal expansion well matched to N-BK7, the material used in the side viewports, which will reduce thermally induced stresses on the glass and bonds [92]. The total length of each science chamber is 345 mm, with an internal bore depth of 335 mm. The internal diameter of the chamber is 50 mm, with an outer diameter of 70 mm. To achieve the deep internal bore of the chamber, and get a smooth finish, an anti-vibration boring bar was used during machining. Holes for viewports are recessed into the chamber wall, with a clear aperture of 20 mm. Parts manufactured by the Physics Workshop are cleaned using the procedure outlined in appendix C.1.

For each science chamber, optical access for cooling and Raman light delivery is achieved through a 60 mm diameter, 12.7 mm thick, $\lambda/60$, UV fused silica window, anti-reflection coated at 780 nm.* Each science chamber also has twelve 5 mm thick, N-BK7, side windows, with a broadband anti-reflection coating over the range 650−1050 nm.† The windows are attached using indium sealing, described in appendix C.2.

On the opposite end of each science chamber, a knife edge compatible with a 2.75" ConFlat® connector is machined in. This allows for interfacing with standard vacuum parts and form the connection to the interfacing module.

3.1.3 Prism mount

In each science chamber, four 15 mm mirror coated prisms form the optics for the horizontal trapping beams of the MOT. To fix the relative alignment

*Manx precision Optics, custom
†Thorlabs, WG11050-B

of the prisms, they are bonded to a titanium mount using a UHV compatible epoxy.* The suitability of the UHV epoxy for this application is tested in appendix C.3, where it is estimated that the relative misalignment of the prisms as a retro-reflector is (80 ± 20) mdeg. To set the alignment of the prisms to the chamber, the mount slides in via grooves along the internal bore of the science chamber [10]. The mount rests on four 0.5 mm wide flat ledges, and is secured using a small amount of a room-temperature curable HV epoxy.†‡

The prism mount also contains fixing points for two dispensers. Each 17 mm dispenser contains 6.8 mg of Rb,§ and is positioned on the underside of the mount, pointing upwards, and slightly towards the chamber wall. They are mounted in this way in an attempt to increase the Rb partial pressure around the trapping region. Securing each dispenser is a pair of customised barrel connectors which are screwed onto Kapton® coated wire. The wire is run through ceramic beads and bushes which are epoxied to the mount, providing further electrical insulation. Neither the dispenser or wire touch the prism mount directly, reducing the risk of electrical shorting due to a scratched wire coating. Additionally, each dispenser leg is bent to act as a spring, allowing for any movement due to thermal stresses. The Kapton® coated wires are twisted in pairs and run along the chamber walls to keep them from the beam path. It was not possible to twist the wires from the same dispenser together to create twisted pairs to reduce magnetic fields.

*Epoxy Technology, EPO-TEK® H21D

†Agilent, Torr Seal®

‡The high curing temperature of EPO-TEK® H21D would require the science chamber to be stood in the vacuum oven to cure. Unfortunately, the vacuum oven lacks the height and would require the chambers to be laid flat while curing, risking misalignment of the prism mounts to the flat ledges.

§SAES, RB/NF/6.4/17 FT10+10

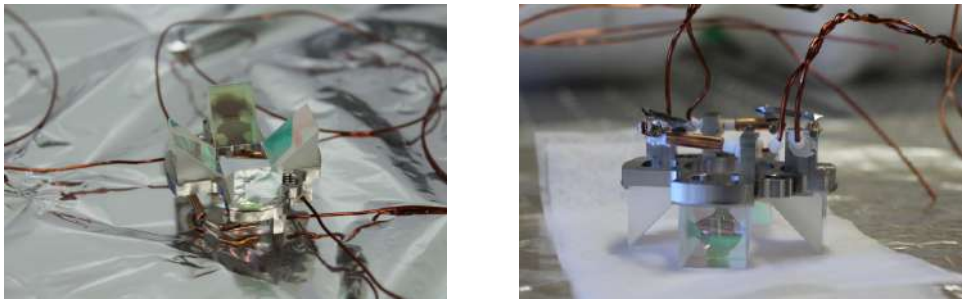

(a) **(b)**

Figure 3.3 Titanium prism mount with 15 mm right angle mirrors bonded to surface. Dispensers are mounted below the prism mount using cut barrel connectors and Kapton® coated wire. Ceramic beads and bushes insulate dispensers from mount. Legs of dispensers are bent to act as a spring to remove mechanical and thermal stresses.

3.1.4 Interfacing module

Having designed the science chambers to fit the prism MOT and drop specifications, connecting the pumps and electrical feedthroughs requires an interfacing module. This module will aid in setting the relative alignment of the two science chambers, as well as providing continuous vacuum between the two, for the Raman beams to travel. With the maximum Raman beam waist set by the distance between the prisms, any vacuum sections must be greater than 15 mm in diameter in order to facilitate the maximum Raman beam size. Using standard, ConFlat® compatible parts, the minimum CF size would be DN16 flanges, which have an internal clear aperture of 16 mm. Since the Raman beam is required to be well collimated and aligned to both clouds, there will be sufficient room for the Raman beam to pass.

The interfacing module is built from standard parts and can be considered as two symmetric halves joined with a bellow. Each half is built around a flange multiplexer which allows for the connection from one DN25 port to five DN16 ports,* The DN25 port attaches to the bottom of the science

*Kimball Physics, MCF275-FlgMplxr-Cr1A5

chamber using an annealed copper gasket, which are used throughout. For the smaller ports, one is on axis and forms the beam path between the two science chambers, as well as providing access to the vacuum valves.* The other four ports form a square pattern, facing outwards at 30 deg to the axis, defined by the two other flanges. Attached to these splayed ports are the following: a $2\,\mathrm{l\,s^{-1}}$ ion pump configured for noble gases;† a housed, $15\,\mathrm{l\,s^{-1}}$, non-evapourable, passive getter;‡ a four-pin electrical feedthrough;§ an SMA feedthrough.¶ Both assembled sections can be seen in figure 3.4 without the nipple and bellow to join them. Again, the a render of the complete vacuum system is seen in figure 3.2.

On the inside of the chamber, attached to the SMA feedthrough, is a 6.8 GHz RF antenna, made from half of a rigid SMA cable. The jacket of the cable has been removed to expose 11 mm of straight wire and positioned to point orthogonal to the chamber axis. The antenna is intended to be used to facilitate population transfer between the two ground states, which may be used in state, and velocity selection using an upgraded control package.

The four-pin electrical feedthrough attaches to the wires connected to the dispensers, via barrel connectors insulated by a Macor® housing. The internal wires are positioned such that they do not impose on the Raman beam path, and out of the detection window aperture.

To complete the interfacing module, a custom length nipple and a flexible bellow are used to link the two halves together. This sets the baseline of the system to 1 m, with some fine tuning achieved by compressing or expanding the bellow to fit into the mounting structure, discussed in section 3.2. Having

*VAT, 54124-GE02
†Gamma Vacuum, 3S-DI-1H-5K-N-N
‡SAES, CF-16-MK2-172-2X16-10
§Lesker, IFDCG012012
¶Lesker, IFDCG012012

Figure 3.4 Flange multiplexers from Kimball Physics join vacuum peripheries to science chambers. The two pieces are linked with a flexible and rigid nipple (not shown).

a flexible bellow in the middle also allows for small angles from connectors to be corrected and removes the stress on the ConFlat® joins when mounting in the rigid structure.

3.1.5 Bakeout for UHV

To reach the required vacuum pressure calculated in section 3.1, the system must undergo a bakeout procedure. Vacuum is achieved through use of a turbo pump, with dry backing pump, taking the system from air pressure to high vacuum. In the high vacuum regime, much of the pressure comes from gas loads from within the bulk materials. During manufacture, gases become trapped inside the materials used to make the vacuum chamber. Heating the chamber increases the rate of diffusion of these gases, resulting in lower outgassing rates once cooled down. The turbo pump is used throughout to remove the gases and is capable of dealing with a large gas load compared to the ion pumps. Heating of the chamber is achieved through wrapping with heating tapes and aluminium foil to provide even heat and avoid hot and cold spots.

To perform the bakeout, the ion pump magnets are removed, compression flanges remain attached to the windows and significant layers of aluminium foil are wrapped tightly to the chamber body. Temperature probes are attached along the chamber body into the mounting holes and between the CF flanges, to best gauge the temperature of the bulk material. Six individually controlled heating tapes are wrapped around the foil covered chamber, with more foil added over the tapes to help spread the heat. The bakeout should avoid temperature gradients for efficiency of outgassing and also to avoid thermal stresses across the windows and CF flanges. As well as spatial temperature gradients, the temporal gradients were limited to a

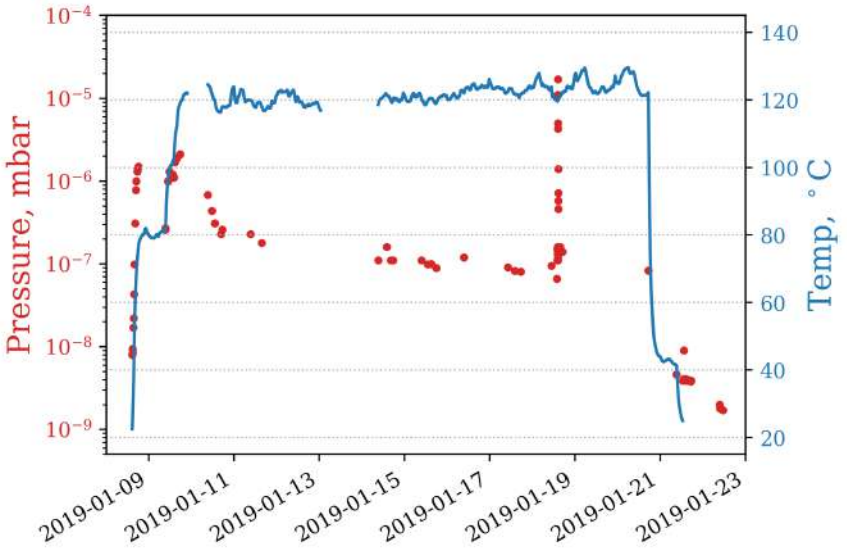

Figure 3.5 Pressure as measured by the turbo pump, and the maximum recorded temperature of the bakeout process. Pressure initially follows the temperature as materials outgas. Once the outgassing rate is reduced, the pressure falls for a fixed temperature. Spike in pressure is seen when the dispensers are degassed.

heating rate of $1\,^{\circ}\mathrm{C}\,\mathrm{min}^{-1}$.

The maximum temperature of the bakeout is limited by the use of indium sealed windows, where the melting point of indium is 156.6 °C. To avoid leak channels forming due to the softening of the indium, the bakeout temperature was limited to a maximum of 130 °C. In practice, the temperature was kept to around 120 °C to allow for drift in temperature throughout the day.

During bakeout, it was useful to de-gas the dispensers by applying a current below their activation current. This allows for trapped gases and surface contaminants to be removed, and means that in principle a dispenser can be activated without needing to reconnect a turbo pump at a later date. The pressure spike can be seen in figure 3.5, as well as the pressure temperature data of the bakeout procedure.

Towards the end of the bakeout, the passive getters were activated by

applying 5 A to the pins for 1 hour. Once the chamber was cool, the magnets were attached to the ion pumps and the valves closed. UHV was confirmed using the ion pump driving current as an indication of the pressure in the system. With the valves closed, the pressure became lower, showing that the turbo pump is no longer required and can be removed. UHV is maintained in the system using the ion pumps and the passive getters, where both ion pumps measure a pressure of $1.2 \cdot 10^{-10}$ mbar when the dispensers are off.

3.2 Mounting structure

The vacuum chamber is mounted in a rigid support structure consisting of four 1" diameter aluminium bars and mounting points to the chamber. The mounting structure sets the baseline and relative alignment of the two science chambers. Plates attach to the bolt threads used for the large window compression flanges and have an external SM3 thread to interface with the telescope lens tube. A central interlocking plate reduces torsion in the structure keeping the vertical bars straight and thus the chambers aligned. Additional bolt holes around the central plate allow for mounting other parts. The total length of the mounting structure is 1.77 m and can be seen rendered in figure 3.6.

As well as supporting the vacuum chamber, the mounting structure is used to interface with the bias coils, base-plate, and magnetic shields. The lower half of the bias coil is supported by four steps which come out of the uprights, as well as being held in the centre with removable supports (not shown in figure 3.6). The upper half of the coil supported from the central supports and clamped with the top-most plate. This forces it to be concentric around the vacuum system, and thus the Raman beam axis.

A base plate bolts to the bottom struts from beneath. The base plate

is large to provide stability for the system, since the centre of mass of the system is near the middle of the vacuum chamber. Built into the base plate are three height adjustable feet which can be individually levelled from the top using a spanner or ratchet. This is possible due to a modification to the top of the screw thread of the articulated foot, where a 10 mm hexagon head has been machined. The base plate is large enough to accommodate a cylindrical box, which is used to contain optical and electrical distribution.

3.3 MOT coils

In the magneto-optical trap, a quadrupole field is required. This is achieved using a pair of coils with current flowing in opposite directions. To determine the distance between the two coils, a simple calculation of two thin current loops of radius r and separation s is performed. As seen in figure 3.7, the field has a maximum gradient when the ratio $s/r = 1$. However, at this point, the field is not linear. A linear field is achieved at the coil centre when the second order derivative of the field is zero. This corresponds to a separation of $\sqrt{3}$ times the coil radius. To ensure a linear trapping gradient, the coils have been designed to have $s/r = 1.7$.

For structural integrity, the coil former is machined from aluminium. A radial slit is made through the former to cut the flow of eddy currents in the metal, which would otherwise increase the time for the coils to be switched. This gap has been filled with plastic to maintain former rigidity. Thermal paste is applied during winding to aid flow of heat to the former which acts as a heat sink. Each coil contains 92 turns of 1 mm Kapton®coated wire (184 turns per former), held in place by epoxy on the outer layers. Each coil's axial field is characterised off the system and shown in figure 3.8 where the central gradient is measured as $(5.1 \pm 0.1)\,\mathrm{G\,cm^{-1}\,A^{-1}}$ and

Figure 3.6 Anodised aluminium support structure holds the vacuum system rigidly and sets the baseline and alignment between the two science chambers

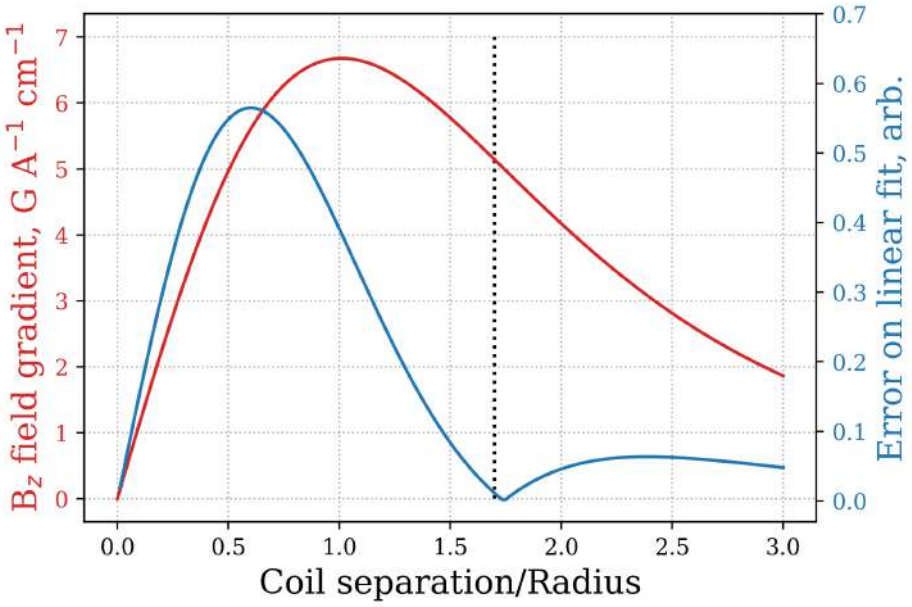

Figure 3.7 Changing the ratio between separation between two current loops with opposite flowing current. The maximum in axial field gradient occurs when the separation is equal to the coil radius, however the axial field between the two coils is most linear when the ratio is $\sqrt{3}$. Dotted line shows the ratio used in the experiment coils.

$(5.0 \pm 0.1)\,\mathrm{G\,cm^{-1}\,A^{-1}}$ for the two coil pairs. The MOT coil formers are bolted directly to the vacuum chamber with the magnetic zero point aligned with the middle of the prisms. During operation, the MOT coil current is 2.5 A, generating a trapping field gradient of $12.5\,\mathrm{G\,cm^{-1}}$ in the axial direction.

3.4 Compensation fields

The magnetic environment is prepared through use of bulk mu-metal shielding, discussed later in section 3.8, with compensation coils used for fine adjustment. This allows for any residual field to be cancelled, which would allow for colder temperatures in sub-Doppler cooling. As well as this, internal coils can be used to shim the position of the cloud, which may be

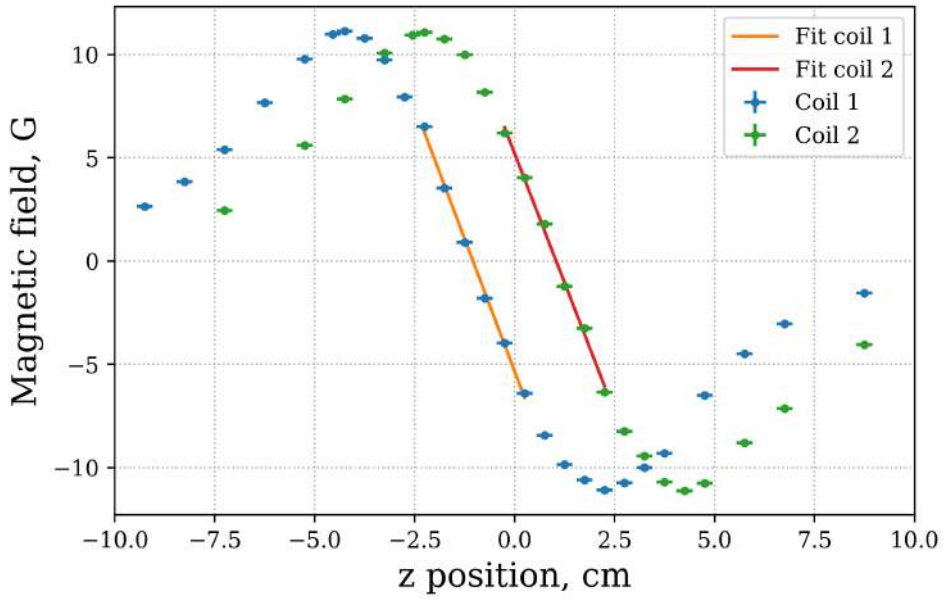

Figure 3.8 Measured field profile of the axial quadrupole trapping coils using a magnetometer. Offset in distance for clarity. Linear fit gives a gradient of $(5.1 \pm 0.1)\,\mathrm{G\,cm^{-1}\,A^{-1}}$ for Coil 1 and $(5.0 \pm 0.1)\,\mathrm{G\,cm^{-1}\,A^{-1}}$ for Coil 2.

used to compensate a misalignment between the magnetic zero point and the optical trapping centre, compensate for Raman beam misalignment, or offsets between the clouds. Due to space constraints, a pair of coils in the Helmholtz geometry cannot be used.* To maximise the region of field uniformity, a rectangular coil pair is used. The characterisation of this pair is shown in figure 3.9, where it can be seen that the field has an extended length of uniformity in the vertical direction.

*For a uniform field between two coils, the separation is given by the Helmholtz spacing. For circular coils this is the coil radius and for square coils, the separation is 0.5445 times the length of a side. At this spacing, the second derivative of the field is zero at the centre [93].

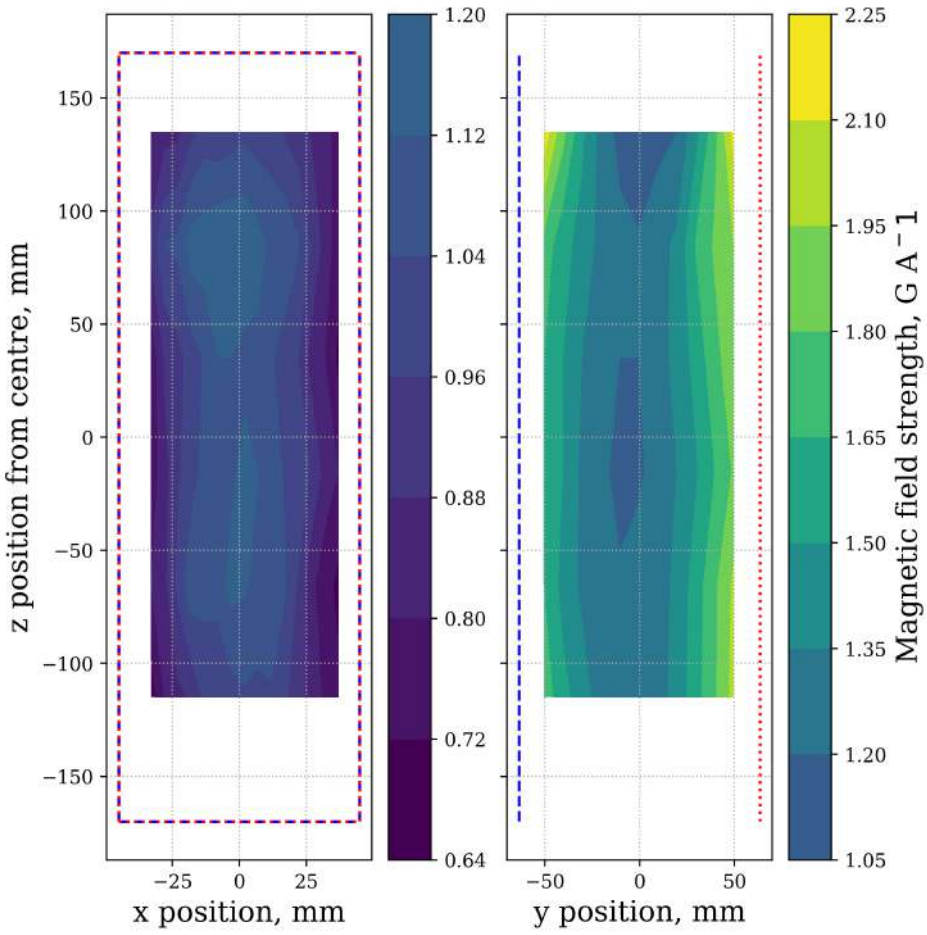

Figure 3.9 Measurement of B_y for a pair of compensation coils. Dotted and dashed lines show the relative orientation of the two coils as seen from the plane of the measurement. The origin is at the centre of the coil pair. Measurements of $B_y(x)$ are in the plane $y = 0$ and vice-versa. Data taken by J. Winch.

3.5 Bias coil

The bias coil is used to define the polarisation of the Raman beam, and can be used to split the m_F states, as discussed in section 2.3. To reduce phase shifts in the gradiometer, a uniform field is desired across the entire interferometry region. In the design of a coil, two extremes are considered: a coil tight to the chamber, or a large diameter coil enclosing most of the sensor package. In the case of the coil close to the chamber, the windings must not get in the way of the detection optics and mounting structure, since this may result in a discontinuous coil, creating additional field components in the xy plane. A tighter coil allows for lower power consumption and, in general, a lower inductance, decreasing the switching speed. If switching speed is not important, then a large coil allows for easier installation and removal, as well as more uniformity around the xy plane, since the coil can better approximate a solenoid. In this case, a large external coil is used due to the constraints of mounting optics.

To produce a uniform field, along an axis one can use an arrangement of multiple coils, whose total field is uniform, or a solenoid. In a constrained environment, the coil system must meet the following conditions: to be rigidly fixable to the sensor package, such that the field is aligned with the Raman beam; to have a uniform field over both interferometry regions; operable from a single current supply, to ensure current noise is common between each interferometer. A solenoid is often used to create a uniform magnetic field; however, their limitations in a constrained environment make them unsuitable. Take a solenoid whose length is much larger than its diameter; the resulting field inside does not become uniform until a length roughly 3 times the diameter of the coil. When the coil's length becomes smaller, so

too does the region of uniformity, where uniformity is considered to be a field variation of less 0.5 %. Thus, in the geometry of interest, the current density must be tailored to produce a uniform field in the region of interest. A multiple coil system would allow for the tuning of the magnetic field in situ by a potential divider circuit, allowing for single supply operation; however, ensuring a rigid structure adds complexity to the design.

Due to these constraints, the best option for this sensor package is a large coil which hugs the inside of the magnetic shield, forming a continuous magnetic field along the entire length of the sensor head, maximising the space for optical access. An optimised current density is required to achieve this and must prioritise a uniform field in the interferometry regions. To optimise current density, variable pitch spacing is employed. This means that for a constant input current, the field profile at any given point is determined by the pitch spacing of all the adjacent wire loops.

To calculate the winding pattern required, the length of the coil is split into regions of current density. Each region is required to have an integer number of turns, making the pitch spacing in that region simply determined by the size of the region and the number of required turns. The coil dimensions were set to be 1.6 m in length and 0.25 m in diameter, and has been split into 101 regions of current density, 15.8 mm long. For practicality of installation and removal, the coil must be split into two halves, which will cause a region in the centre with a single turn. If the turn spacing is small compared to the radius of the coil, then the helical coil can be be considered as consisting of many current loops, separated by the pitch spacing of subsequent turns. Treating the coil this way, the magnetic field along the axis is given by the summation of all current loops at every point. For a coil of radius a, N turns, and each turn, n, having a position z_n, then the magnetic

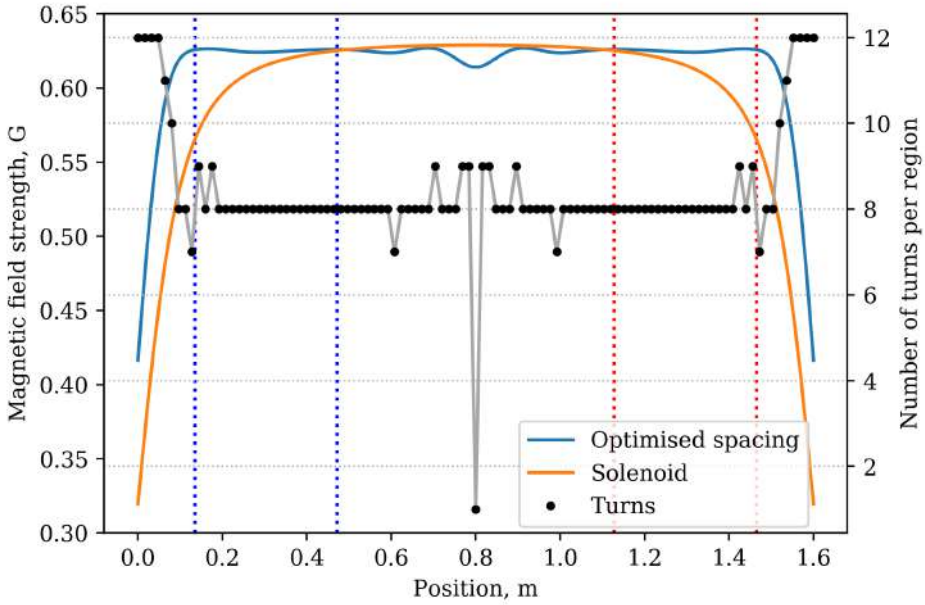

Figure 3.10 Biot-Savart calculations of the optimised field geometry. Comparison between a solenoid shows that the field is more uniform between the interferometry regions, shown by the pairs of dotted lines. The coil is split into 101 regions, each 15.8 mm long.

field at each position, for an input current I, is given by a summation of the Biot-Savart law

$$B(z) = \sum_{n=1}^{N} \frac{\mu_0 I a^2}{2\Big((z - z_n)^2 + a^2\Big)^{3/2}}. \tag{3.3}$$

To optimise the field, the number of turns in each current density region is changed, starting from a solenoid of uniform turn density. It is required that the field be uniform in the interferometry regions. Achieving this requires the maximum current density at the edges of the coil to ramp up the magnetic field. In total, there are 849 turns on the coil with pitch spacing varying from 1.32 mm to 2.26 mm. The optimised field is shown in figure 3.10 alongside the number of turns per region, demonstrating the turn density changes along the coil.

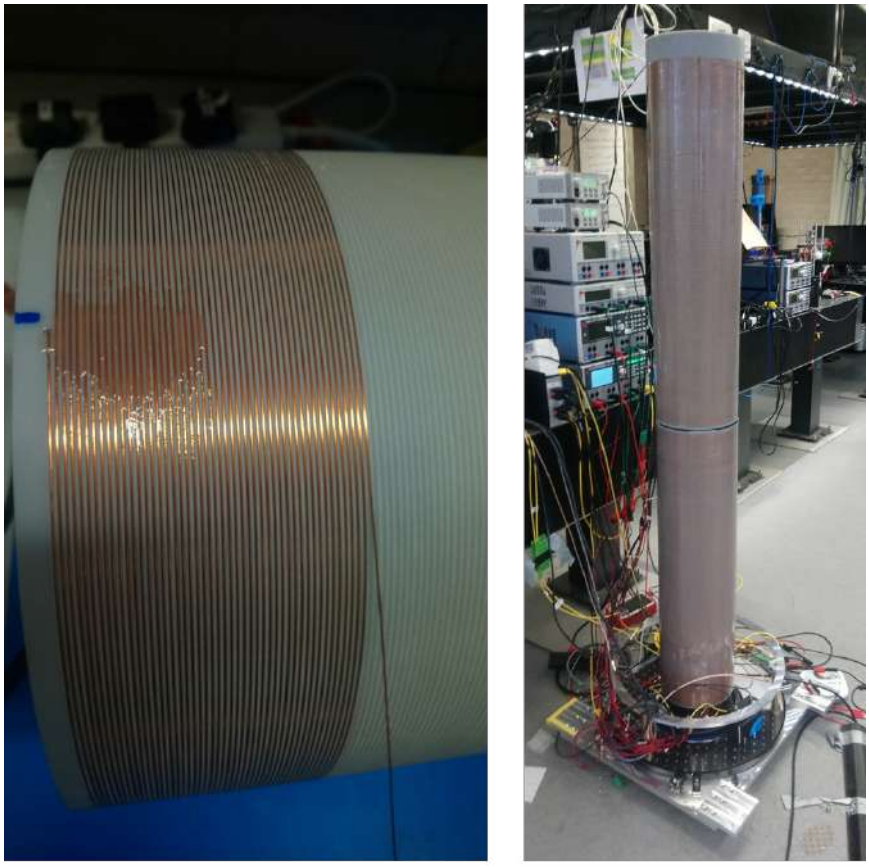

Figure 3.11 a) Winding the bias coil with 0.5mm wire. b) Early fitting of the wound bias coil on the sensor head. It can be seen that the bias coil completely encloses the vacuum system.

The coil is wound onto a 3D printed former with grooves for each turn, made of a resin material to prevent additional eddy currents. Using a series of rollers, the wire was easily added to the former and an epoxy layer stops the wires from being scratched or displaced. The winding of the former, and the fitting of it to the sensor package is shown in figure 3.11.

Measurements of the field profile were taken with the axial probe of a Gauss meter, along the length of the coil by attaching the probe to taut string and pulling through the centre. The measurements are shown in

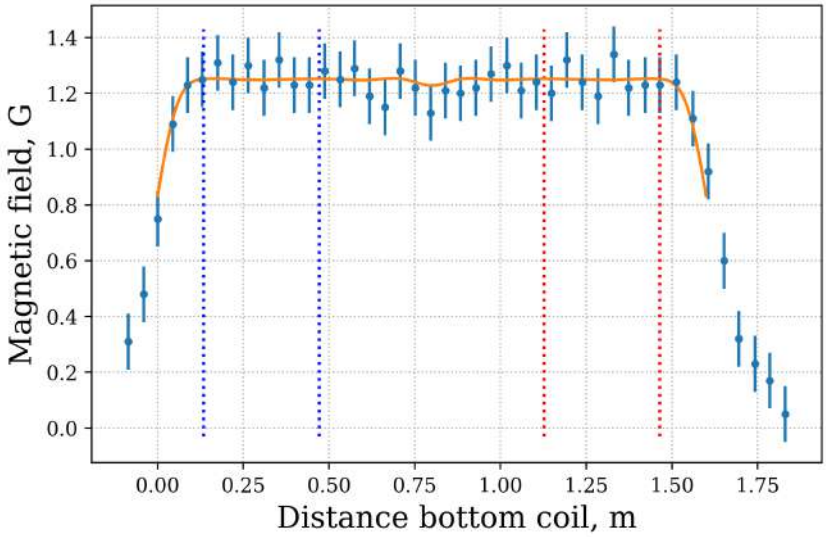

Figure 3.12 Measurements of the bias coil field profile with a axial probe Gauss meter, compared to the modelled field. Large error from poor method obscures the flat field profile; however, the sharp rise at coil edge is observed. The interferometry regions are again shown with pairs of dotted lines, showing a uniform field to within error.

figure 3.12, where despite large measurement noise obscuring the flat central profile, the sharp rise at the coil edges is measured.

3.6 Beam delivery

To have a high enough intensity for the Raman laser, the Raman and cooling beams must be different sizes. The cooling beam size is determined by the size of the prism MOT, with a minimum beam diameter of 47 mm required to cover the prisms completely. The Raman beam must pass through the centre of the prisms and through the vacuum bellow without clipping, setting the maximum beam diameter of 15 mm. In the system, the cooling and Raman beam are delivered co-axially, which adds complexity to the telescope design. Previous designs focused the Raman and cooling beam close together, with one reflected from a D shaped mirror, which produces

near co-axial beams but the focusing introduced wavefront aberration [10]. To overcome this, the design used here requires the two beams to be cross linear polarised, combined together on the output port of a polarising beam splitter (PBS). The cooling and Raman beam delivery is shown in figure 3.13, where alignable fibre coupled collimators are used to direct the beams onto the reflection inputs of two PBS. With the arrangement of the beam splitters such that the polarisation's are crossed, both beams may be delivered co-axially, with the fraction of the output power determined by the alignment of the half-wave plate (HWP), with the PBS.

A variable retarder plate (VRP), is aligned with its axis at 45 degrees to each linear polarisation. The VRP is used to switch between circular and linear polarisation for the MOT/molasses, and Raman phases respectively. Each VRP is individually temperature stabilised at 40 °C, to reduce polarisation fluctuations.

The beams are expanded using a Galilean telescope consisting of a diverging lens, focal length −25 mm, and a converging lens with focal length 100 mm. With the telescope collimated, the relative beam sizes are set by the ratio of the two collimator outputs. The Raman collimator has an output beam diameter of 3.1 mm when used with the 6 µm mode field diameter input fibre.* The cooling collimator is a fixed focus, parabolic mirror collimator, with an output beam size of 12 mm.† The ratio of the collimators therefore set the output beam diameter to be 12.4 mm for the Raman beam and 48 mm for the cooling beam.

To secure the optics in place, a custom telescope body, manufactured from aluminium, rigidly fixes the VRP, PBS, and HWP in place. The lenses are secured in lens tubes, which are screwed into the telescope body. The

*Raman collimator: *Thorlabs* PAF2A-18B
†Cooling collimator: *Thorlabs* FC12APC-PO1

alignment process for each beam is performed by a lockable 2-axis tip tilt stage, where the fibre collimators are mounted to the telescope body. This allows each beam to be aligned independently. The collimation is determined by the relative separation of the two telescope lenses. Each telescope can be thus be aligned and collimated, such that the cooling and Raman port from each telescope are coaxial; however, the relative alignment of each telescope is more difficult due to no access to the beam path and limited diagnostics. With each telescope on the system, the atomic signal can be used as a metric for alignment quality. Since the telescopes are inside the magnetic shielding, it becomes difficult to optimise the beam alignment using the atoms. This is where the high level of integration of the sensor package begins to hinder the optimisation process. This application would be ideal for piezo actuated telescopes, where the alignment could be controlled externally, whilst maintaining a fully integrated sensor package. A telescope made of customised parts would need to be manufactured in order to fit this functionality, but presents technical challenges with robustness. The telescopes are screwed to the mounting structure with the 3" lens tube and further braced against the mounting structure. This rigidly fixes the telescopes to the system and avoids telescope motion, and thus beam misalignment, when the sensor package is moved.

3.7 Detection and diagnostics

Detection of the atoms is done by capturing the fluorescence with a photodiode. The cooling light can be sent through the main telescopes, to measure the state populations and readout the interferometer, or diverted with an optical switch to a side port light sheet, used for optimisation and diagnostics. The hardware used is briefly mentioned here, with further information

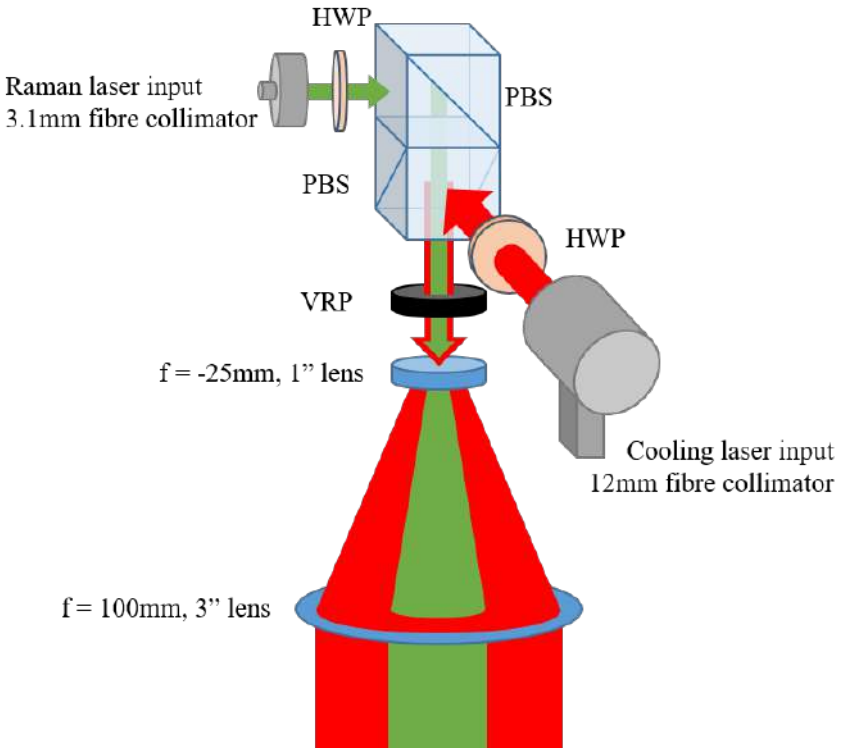

Figure 3.13 Exploded view schematic of the telescope used in the system. The ratio of the output beam sizes is determined by the ratio of the beam sizes at the collimators. A careful choice of the collimators has produced an output cooling beam of $1/e^2$ waist 48 mm, and a Raman beam waist of 12.4 mm.

about the detection scheme used in the experiment discussed later, in section 5.2.

3.7.1 Detection photodiodes

A fast photodiode* is mounted onto a custom amplification circuit† to provide a large signal to noise ratio. The output of the photodiode is fed directly into a 10^7 gain inverting amplifier to provide the convert photocurrent to voltage. Second stage amplification is used to perform active low-pass filtering, at inverting unity gain, with a cutoff frequency of 28.2 kHz. The amplifiers are supplied using 2 stage power line filtering. The low-pass filtering stage limits the photodiode detection bandwidth to 40 kHz.

These photodiodes are mounted rigidly to the chamber using 3D printed mounts. A double lens system is used, where a lens collimates the light from the atoms (if treated as a point source), with a second lens focusing the collected light onto the photodiode. The detection solid angle from this system is estimated as 0.16 sr, covering only a small fraction of the surface the atoms radiate over.

The detection diodes are fed into a data acquisition device (DAQ), via SMA cables. Ferrite cores on each end of the cable reduce electrical pickup noise and the cables are kept as short as possible. The data is digitised by the high resolution, low noise DAQ before being sent over the umbilical for data processing.‡

Hamamatsu S1336-44BK

†Designed by F. Hayati

‡DAQ: *National Instruments* USB-6366

3.7.2 Light sheet detection

Orthogonal to the detection photodiode, a horizontal light sheet is installed into each chamber. A fibre collimator, with output beam waist of 8 mm, is masked to a thin sheet of a few hundred µm, using a 3D printed slit and some blackout foil.

The light sheet is used as a diagnostic tool to measure atom cloud temperatures and arrival time, as well as giving an indication of the number of atoms. Much as for the light delivered from the vertical beams, the fluorescence from the atoms is measured by the detection photodiodes. Light sheet detection is not for measuring relative state populations in this sensor; however, the slit can easily be removed to allow for horizontal detection schemes and adapted to form double light sheets, or just open port pulsed detection. The light sheet can also be balanced by retro-reflecting through the introduction of a mirror with a quarter-wave plate (QWP), attached to the front. The QWP stops intensity interference patterns forming when retro-reflecting, creating a beam of equal intensity from both directions. In principle, this helps to increase the signal through more power and reduce the number of atoms pushed out of the detection region.

3.8 Magnetic shielding

Part of operation in a field environment is mitigation of external effects. For a portable system, the background fields will change with measurement location and may also have a temporal and frequency dependence. An identified issue with the previous sensor was susceptibility to external magnetic fields; here, the sensor head is designed to be completely enclosed by bulk magnetic shielding, in order to reduce uncommon phase shifts in the gra-

diometer. There are two standard approaches to magnetic shielding: bulk shielding, using a high permeability ferromagnetic material, such as mu-metal, to guide fields around the shield; or active field cancellation, where several coil pairs are used in a feedback system to constantly cancel fields.

Bulk shielding, whilst increasing the sensor weight, provides additional benefits over an active shielding approach: a reduction in complexity, where a 3-axis feedback system is not required; improvements to general robustness and protection against other environmental effects, such as ambient light levels and fast temperature fluctuations; magnetic shielding that is always providing the same level of protection, independent of sensor position or orientation. A bulk magnetic shield instead increases the requirements on integration, with all instrument components required to fit within the confines of the shield. Interfacing and wiring must be concentrated into small openings in the shield to avoid flux leakage into the sensor. The initial demonstrator could not opt for a fully enclosed shield due a strict sensor weight limit [10]; however, the limits have been relaxed for this sensor package in order to prioritise environmental isolation.

Bulk magnetic shielding is used in multiple places in the instrument to remove the effect of external and internal magnetic fields. Any tweaking to the magnetic field profile where the atoms are positioned, may be achieved using the compensation coils described in 3.4.

3.8.1 Main shield

The "main shield" is used to fully enclose the sensor package, to approximate an ideal mu-metal can. The construction is three cylindrical layers of 1 mm mu-metal,* with external height 1.8 m and outer diameter 264 mm. The

*ASTM A753 Alloy 4

shield exists as three parts: top, bottom, and base cap; to form a contiguous enclosed shield that can be removed section by section.

The base cap sits on the base plate and has four 28 mm holes to accommodate the uprights of the mounting structure. Lips on the edges of the three layers receive the bottom shield. Once attached to the base plate and mounting structure, the base cap is not removable.

The bottom shield consists of three concentric cylinders which form a push fit join with the base cap. Along the bottom are two semi-circular holes of 38 mm diameter, through all three layers. These holes are the only opening of the magnetic shield and are used to route all cabling into and out of the sensor head. Incorporation of these mouse holes allows for the bottom shield to be removed without the need to remove any wiring or unplug components.

Finally the top shield has three concentric cylinders to match the bottom shield, and each layer is capped at the top. The join between the layers of the bottom and top shield are vertically offset by 111 mm between each layer. This helps to reduce field leakage through the joint.

During shield characterisation from the manufacturer,* the shielding factor was measured to be between 440 and 530 at the interferometry regions, measuring less than 1.5 mG of residual Earth field. Additional testing with AC fields up to 500 Hz showed no significant interference.

The shields are powder coated to protect them on the outside in a glossy blue, but left exposed where the parts join, to ensure a good connection. On the inside layer only, a matt black powder coating is used to absorb any stray light.

In total, the shield weighs approximately 45 kg, but in its two halves

Magnetic Shields Limited

may be fitted or removed from the system within 5 minutes by two operators, without the need to disassemble anything. The ease of removal is vital to enable optimisation of internally adjusted components, such as the telescopes, since there is no additional access.

3.8.2 Additional magnetic shielding

Base shield

A cover is needed for the base of of the sensor head since it houses some optical and electrical distribution. To provide additional protection, this cover is made from two layers of mu-metal and fits around the main shield. An opening for cable and optical fibre interfacing is placed at 135° from the mouse-hole openings in the main shield to reduce further field leakage.

Ion pump shield

Since the entire sensor head is shrouded in magnetic shielding, fields produced internally are confined inside the shield. A large source of constant magnetic field is the strong rare-earth metal magnet on the ion pump. A single layer, 1.5 mm mu-metal shield, is fit to each ion pump to help confine the field. Tests showed that the functionality of the ion pump is unaffected, and measurements of the field close to the ion pump gave a reduction of magnetic field from 180 G to below 1 G with the field falling off quickly with distance. The ion pumps were designed to be far away from the MOT region, so the residual field should not affect the cold atom generation. The field may also be compensated with the compensation coils if the optical molasses temperature is affected. Since the ion pumps are always on, the presence of the ion pump field may distort the bias field, creating a systematic phase offset in the gradiometer.

3.9 Integrated sensor head

For the characterisation of the instrument in this thesis, the sensor head is operated in a drop configuration. The assembled sensor head is shown without the main and base magnetic shields, or the bias coil, in figure 3.14a. In the figure, it can be seen that all available space around the vacuum system and mounting structure has been utilised with wiring, optical delivery, detection assemblies, and diagnostic equipment. The bias coils and magnetic shields fit tightly over all the cabling, which is all routed through the two yellow mouse hole covers at the base, seen in figure 3.14a. Once in the base box, the wires are connected to an industrial connector to interface with the umbilical,* and the optical fibres connected to bulk head adapters.

The fully integrated instrument is shown in figure 3.14b, where the instrument is in a mobile form, ready to be transported to a measurement site. Here, the main magnetic shields can be seen, as well as the base box shield. The sensor package has also been mounted into a trolley to facilitate transport, without the need to lift the sensor. The sensor package is estimated to weigh 80 kg, requiring several operators to lift, or just one to push. The control package is on castor wheels and is manoeuvrable with one operator. On level ground, the entire instrument may be moved with two people, but three is recommended.

* *Harting* Han-Modular®

(a) **(b)**

Figure 3.14 a) Assembled sensor head in drop configuration, shown with shields and bias coils removed. Base box contains electrical and optical interfacing between the umbilical and sensor package. The cabling is routed through two small holes in the shields and along the mounting structure to each required component. b) Sensor head and control package fully integrated, shown in mobile form on the transportation trolley.

Chapter 4

Control package

The control package consists of the laser system, RF driving electronics, sensor head electronics, and a computer control system. Each will be briefly be discussed, including the details pertinent to operation with the new sensor package.

As discussed in section 2.9, an existing control package is used with the new sensor head in order to achieve application relevant sensitivities. The existing control package has proven to be field robust and meets most of the technical requirements to drive the new sensor package.

The control package, seen in figure 3.14b, consists of sub-systems, packaged into standard 19" rack mounted boxes, fixed into a portable flight case. All necessary inputs and outputs of the control package are fed into the umbilical to reach the sensor package.

4.1 Laser and RF system

The laser system is based around frequency doubling the 1560 nm output from C-band telecom laser components. This approach allows for a me-

chanically robust, reliable, all fibre, laser system at high power and narrow linewidth, using off-the-shelf components [12, 94].

The laser system may be considered as two arms: one generating the cooling and repumper frequencies, the other generating the two Raman frequencies. In this scheme, the cooling laser is locked to an atomic transition, and the Raman laser offset locked to the cooling seed frequency. The architecture of the laser system is shown in figure 4.1, where the RF frequencies applied to the electro-optic devices is also shown.

Each arm of the laser system follows the same architecture. A 1560 nm fibre seed,* of output power 40 mW, is split into two paths using a fibre splitter. Each path goes through an electro-optic modulator (EOM),† in order to add various frequency sidebands to the seed laser frequency. The output of each EOM is amplified in an erbium doped fibre amplifier (EDFA),‡ up to a power of 2 W at 1560 nm. The output from each amplifier is frequency doubled in a periodically poled, lithium niobate, ridged wafer waveguide (PPLN),§ outputting light at 780 nm at powers of up to 1 W. From here, the architecture varies, with light being delivered to the sensor package telescopes passing through a fibre acousto-optic modulator (AOM),¶ to facilitate fast switching of the light.

4.1.1 Cooling and spectroscopy laser arm

In order to lock the cooling seed laser, one of the laser paths is sent to a ^{85}Rb vapour cell for atomic locking. Frequency sidebands are added to the EOM at a tunable frequency of around 1.2 GHz, with the higher frequency

NKT Photonics Koheras BASIK
†*iXblue* MPZ-LN-10
‡*Photonic Solutions* OEM BoostiK
§*NEL* WH-0780-000
¶*Gooch & Housego* FibreQ T-M150-0.5C2W-3-F2P

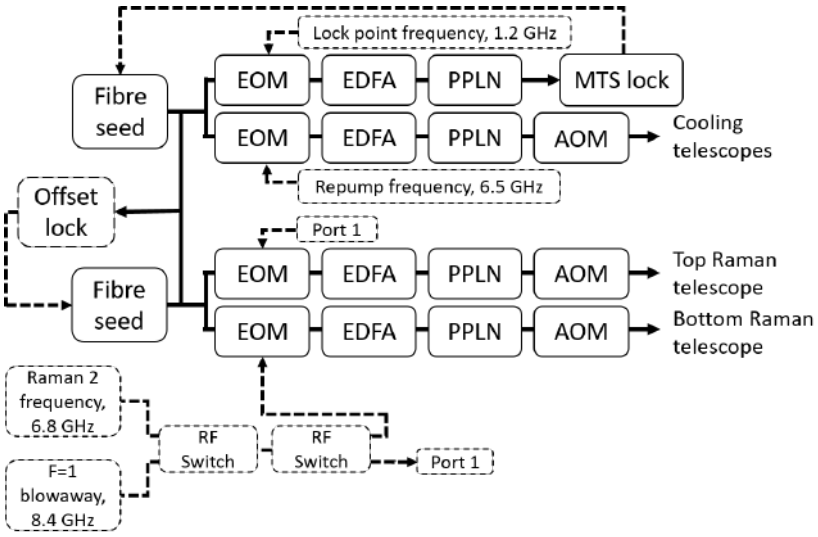

Figure 4.1 Diagram of the all fibre laser system (represented by solid lines), and key parts of the RF chain (represented with dashed lines). System uses a two seed architecture, frequency stabilised to each other with a side of filter offset lock [95]. The cooling laser frequency is frequency sideband locked to an atomic resonance using modulation transfer spectroscopy [96]. Raman laser frequencies are generated using the fibre seed frequency and an EOM generated frequency sideband. The frequency sideband can be applied to either arm of the Raman laser through the use of an RF switch. Through an RF switch, the Raman frequency sideband can also be used to create a blowaway resonant with atoms in the $|F = 1\rangle$ ground state (discussed further in 5.5)

sideband locked to the $|F = 3\rangle \leftrightarrow |F' = 4\rangle$ transition of ^{85}Rb, through modulation transfer spectroscopy [96]. By locking the sideband to an atomic reference, the seed laser frequency may be tuned, by changing the RF frequency applied to the EOM. Since the carrier frequency is not required for locking, the EOM is driven at a high modulation index to have the most optical power in the frequency sidebands. The total power sent to the modulation transfer setup is on the order of 500 mW, which gives the largest signal to noise from the locking signal.

The second path from the cooling seed laser is used to deliver light to each MOT telescope. The frequency of the cooling light is determined by the RF frequency applied to the lock EOM sideband, and can be ramped by 30 MHz ms^{-1}, over a range of 200 MHz, limited by the RF generation. The repump frequency is generated through a frequency sideband of roughly 6.5 GHz, modulated onto the EOM in the cooling chain [97]. The repump frequency can then be turned on and off by switching the RF source.

To calculate the power requirement for the cooling laser, each MOT must deliver an intensity to the atoms of order the saturation intensity ($\approx$1.7 mW for circularly polarised light [47]). For the beam sizes of the telescopes shown in section 3.6, each telescope must have a minimum input power of 30 mW to reach the saturation intensity. Typically, intensities of a few times the saturation intensity are used in the magneto-optical trap to increase the number of trapped atoms [98], requiring a telescope input of 100 mW. The total power from the laser should therefore be at least 200 mW, and to account for losses in fibre delivery and telescopes, an output power from the laser of 300 mW is desirable. With a maximum insertion losses of 3 dB from the AOM, output powers from the PPLN should be around 600 mW, achievable with use of a 2 W EDFA in each path. When measured, the

output from the AOM is around 500 mW, with a maximum central intensity from each cooling telescope of 7.5 mW cm^{-2}.

4.1.2 Raman laser arm

The Raman laser seed is locked to the cooling laser seed by beating the two sources together, and locking the beat frequency to an RF frequency, through a side of filter lock [95]. The Raman laser frequency can therefore be moved by changing the cooling laser frequency. During a measurement, the Raman seed frequency is left at a constant value to maintain the value of the single photon detuning, Δ. The value of Δ is 1.9 GHz, detuned below the $|F' = 3\rangle$ state.

The two experiment outputs of the Raman seed, are used to create the two arms of the Raman laser, both comprised of the same component chain as the cooling laser: EOM, EDFA, PPLN, AOM. Each AOM output is sent down the umbilical, into one of the Raman telescope inputs. In order to generate the second Raman frequency, a frequency sideband is applied to one of the EOMs, with a modulation frequency equal to the hyperfine splitting of the ^{87}Rb ground state - 6.8 GHz. By tuning the RF frequency, the value of the two photon detuning, δ, may be nulled. Applying a linear chirp to the RF frequency allows for compensation of the Doppler shift, as discussed in section 2.4.3.

Each Raman arm has a power of up to 200 mW into the telescopes, with a maximum central intensity, measured from the Raman telescope, of 170 mW cm^{-2}. By feeding back on the AOM, the Raman laser is intensity stabilised to within a few percent, with the intensity pickup being provided by the detection diodes. With stabilised intensities and balanced powers, the Raman intensity from each arm is around 100 mW cm^{-2}, which using

equation (2.27), should provide Rabi frequencies of up to 340 kHz, allowing for short π times and a large velocity class to be selected.

Raman beam delivery

To improve contrast and reduce bias in the interferometer, parasitic transitions are removed through appropriate choice of the Raman beam delivery. Parasitic transitions are introduced when a modulated Raman beam is reflected, and the reflected sidebands form additional Raman transitions [99]. As well as phase shifts in the final interferometer, these additional transitions create a spatial dependence on the Rabi frequency, with spatially varying oscillations seen on the length scales of the modulation frequency [10, 99, 100].

Parasitic transitions are overcome in this system by delivering the modulated Raman light from one direction only, allowing only one pair of Raman frequencies to be resonant with a transition. This approach is not achievable in a high sensitivity cold atom gravimeter, without an additional phase lock between the two beams, since uncommon phase shifts between the two laser arms are introduced by splitting the beam delivery. In the gradiometer, this laser phase noise is common to each interferometer and is subtracted as a common phase noise in the gradient measurement.

In order to enable the removal of systematic effects, such as external magnetic fields, it can be useful to alternate the direction of the photon momentum kicks in the interferometer [101]. In this beam configuration, this can be readily achieved by alternating which EOM the 6.8 GHz modulation is applied to, changing the direction the Raman sideband comes from, and thus the direction of the wavevector $\mathbf{k}_{\text{eff}}$. The flipping of the wavevector, changes the direction of the photon momentum kick, transferred to the atom

in the interferometer pulses, reversing the shifts due to inertial effects on the interferometer. Systematic effects, such as magnetic fields, still contribute with the same sign. Using area reversal techniques, where the direction of the wavevector is alternated each measurement, systematic effects may be suppressed in the interferometer when the subtraction of the interleaved measurements is taken. The hardware is implemented in the system; however, area reversal will not be demonstrated in this work.

The inputs to the Raman EOMs can also be used to add a frequency sideband to generate light resonant with atoms in the $|F = 1\rangle$ hyperfine ground state. In the MOT, molasses, and detection phase of the sequence, this role is fulfilled by the repump frequency sideband added to the cooling laser. Consequently, it is not possible to have the repump sideband without sufficiently extinguishing the cooling light. This sideband applied to the Raman seed frequency can be used to blow-away atoms in the $|F = 1\rangle$ ground state and will be used for sub-level state selection, described in section 5.5.

4.2 Timing and electronic control

Precision timing is required in the instrument in order to perform sequences repeatably and without changes to the length of light pulses. This is particularly crucial for the Raman pulses, where the pulse length determines the state population transferred by the atom optics. Consequently, the experiment sequence is performed on an FPGA development board,* operating on a custom python interface. The system is capable of performing outputs at 100 ns resolution, with 40 digital outputs and six analogue outputs. The digital outputs are used to trigger RF switches to change the laser frequency and power, and to trigger IGBTs to switch the power in the coils

National Instruments sbRIO-9627

and dispensers.

MOT and compensation coils, as well as dispensers, are powered using stable, low noise benchtop power supplies.* For the bias coil, a low noise, stable power supply is used to reduce any additional phase noise in the interferometer.†

The control package, and by extension the sensor head, are powered by a single 240 V, 13 A, mains power cable. To reduce noise coupling into the system from the mains supply, and to allow for continuous operation whilst moving between mains outlets, an uninterruptible power supply (UPS), is used.‡ The system may run on the UPS for 17 minutes before the battery is close to depletion, which was found to be enough time to set up further mains connections. The instrument can also be run using a generator as the mains power source, opening the possibility for measurements in more isolated areas.

4.3 Data analysis using ellipse-specific fitting

Extraction of gradiometer phase is done using an ellipse-specific fitting method [102], which has been successfully applied to gravity gradiometry [73]. The analysis is performed on data saved onto the experiment PC and can be done in tandem with the experiment sequence running.

Since the outputs of the two interferometers can be considered sinusoidal, with a phase shift between them, a conic section constraint, with the discriminant of an ellipse, is fit using a least-squares optimisation. The algebraic form of a conic section is the function

Rohde & Schwarz HMP2030
†*Keysight Technologies* B2962A
‡*Power Solve* Xi 1500

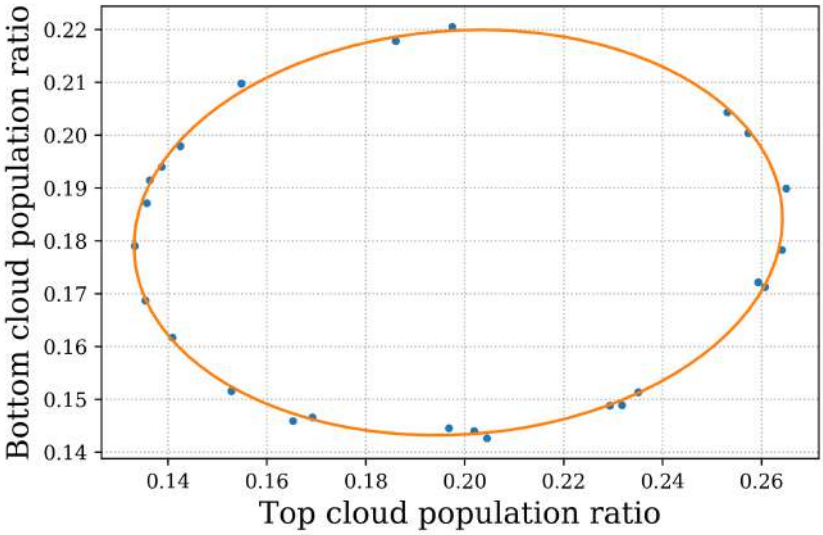

Figure 4.2 25 points of synthetic data, with normally distributed amplitude noise, is fit with an ellipse to determine the gradiometer phase.

$$F(a,x) = \mathbf{a.x} = Ax^2 + Bxy + Cy^2 + Dx + Ey + F = 0, \tag{4.1}$$

where a and x are the matrices $[A, B, C, D, E, F]$, and $[x^2, xy, y^2, x, y, 1]^T$, respectively. The discriminant $B^2 - 4AC < 0$, constrains the solution to an ellipse and from the fit parameters, the ellipse phase $\Delta\Phi$ can be extracted.

$$\Delta\Phi = \cos^{-1}\left(-\frac{B}{2\sqrt{AC}}\right) \tag{4.2}$$

To avoid constraining the fit, the ellipse parameters are free to change. This means that there is a minimum number of points required to reproduce the phase shift between the two fringes. With noise on the measurements, the minimum number of points is increased. To calculate this, synthetic data is produced, with a similar noise profile to that of an experiment data set. Noise is added to the population ratio as amplitude variations, using the later introduced equation (6.4), independently for each cloud. Drift in the

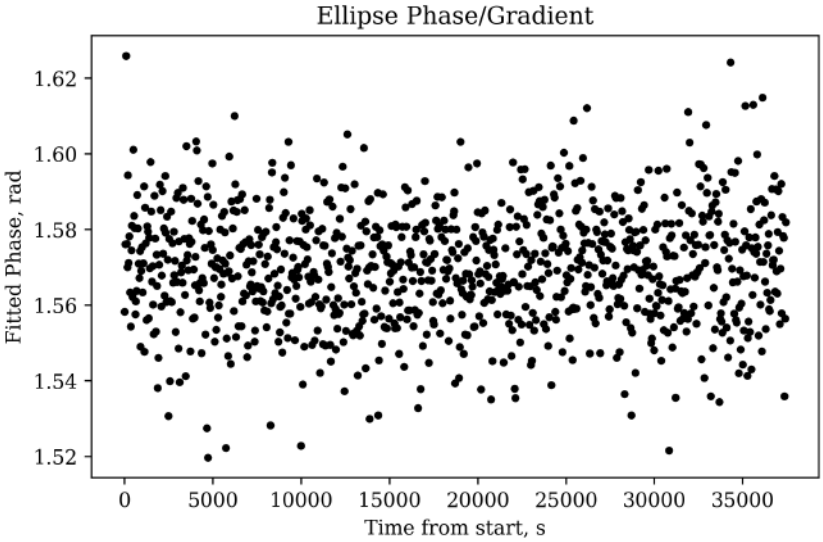

Figure 4.3 Extracted phase values for ellipses fit to 25 shots of synthetic data. The mean and standard error of the phase shift are taken for the total data set.

phase or uncommon phase noise is deliberately avoided so that the applied phase shift is well known. An increasing number of shots is fitted to each ellipse and the phase value extracted for a total of 25000 shots. One set of phase data is shown in figure 4.3, where each ellipse is fit with 25 shots.

Taking the mean and standard error for the data, fit with different numbers of ellipses, is shown in figure 4.4. With increasing number of shots, the mean phase tends towards a constant value, and the standard error decreases. Since the number of ellipses fit is reduced, the standard error only sees a relatively small improvement. Either through the fitting or the position and noise on the ellipse, there is a bias on the value of the extracted phase. Calculating an Allan deviation of the data set, discussed further in appendix B.2, the short term sensitivity can be evaluated. Figure 4.5, shows the change in sensitivity with the number of fitted shots, where it is seen that the short term sensitivity is overestimated when the number of shots fit to the data is too low.

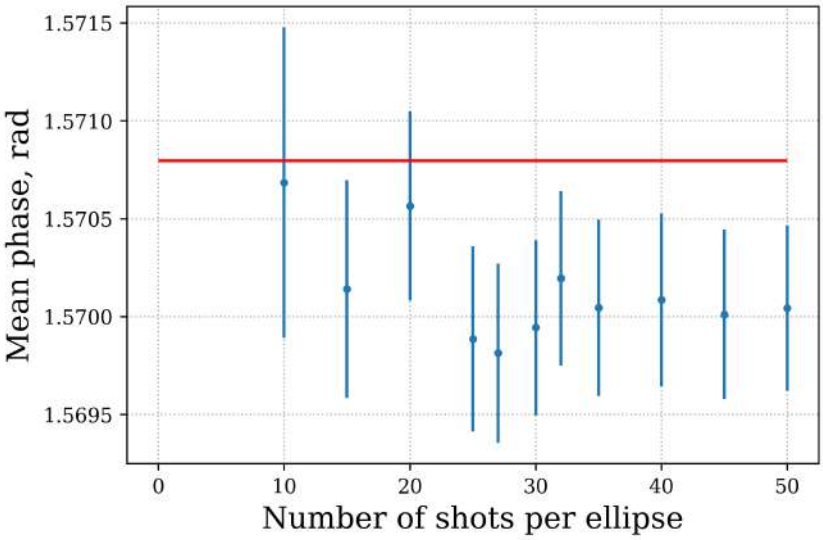

Figure 4.4 Average and standard error of the calculated phase when the number of shots per ellipse is changed. A bias is seen on the fitted phase when compared to the applied phase shift (red line). The standard error of the fitted phase decreases with more shots per ellipse.

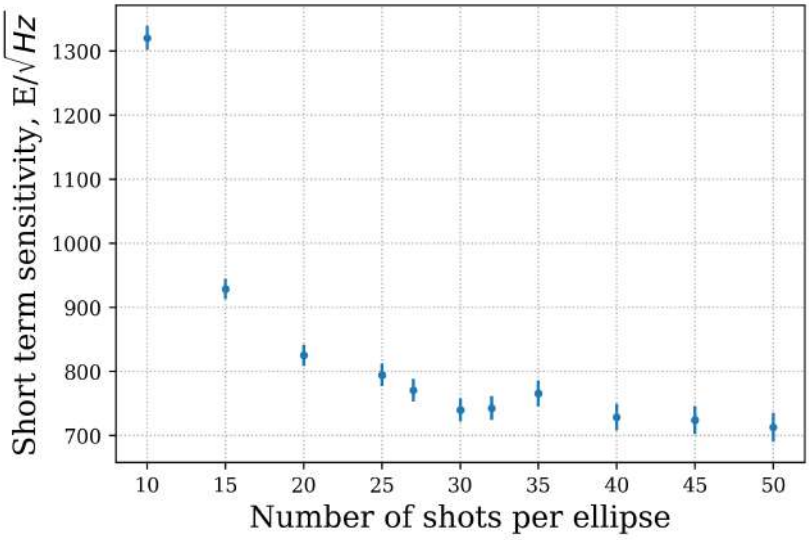

Figure 4.5 Number of shots used in each ellipse fit changes the value of short term sensitivity for data with normally distributed amplitude noise. Since fitting to more shots requires more measurement time, the error is approximately the same.

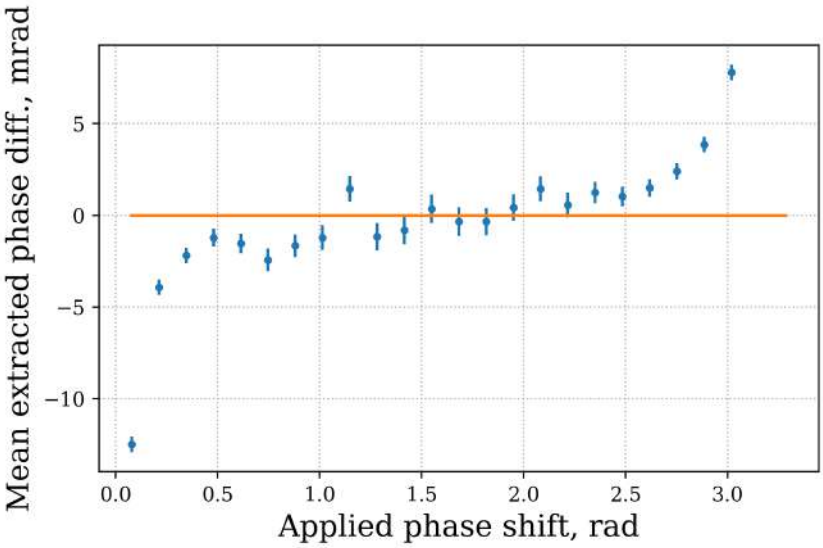

Figure 4.6 Bias seen in the fitted phase, compared to an applied phase shift. Bias is strongest when the phase is a multiple of π, causing the ellipse to tend towards a line. For data with applied amplitude noise only, the phase error is lowest when the ellipse tends towards a line.

The fitting bias seen in figure 4.4 is also dependent upon the applied phase. For ellipse phases corresponding to multiples of π, the ellipse forms a line. In the presence of amplitude noise, it is not possible to determine which side of the ellipse the point is on, thus the phase of the ellipse is incorrectly calculated. This is shown in figure 4.6, where the difference between the applied phase shift and the calculated ellipse phase is shown for ellipse phases between 0 and π. Despite a bias being seen, the uncertainty on the ellipse phase is lowest for the extreme points, possibly due to the modelling containing amplitude noise only. To provide context of noise sizes, a phase uncertainty of 1 mrad at an interferometer T time of 85 ms, gives a phase error of 9 E.

Chapter 5

System characterisation

With the sensor package constructed and connected to the control package, the full system may be characterised. This will focus on the atomic characterisation, and measurements pertinent to its operation as a gravity gradiometer. The characterisation loosely follows the steps of the experiment sequence used in the final gradiometer measurement. The sequence contains the following steps, with an indication of the length of time per step:

1. Trapping atoms in a magneto-optical trap (1 second).
2. Cooling trapped atoms using an optical molasses, a sub-Doppler technique (10 millisecond).
3. Internal energy state preparation (10 millisecond).
4. Interferometry sequence (200 millisecond).
5. Detection of atoms in each internal state (10 millisecond).

Within each of these steps, the hardware has been meticulously checked for consistent operation; much of the early work was to ensure that the

hardware was performing the desired sequence without bugs. The latter forms most of the work presented in this section, where parameters have been optimised for overall instrument sensitivity. The figures of merit for a sensitive instrument are related to total atom flux, i.e. the number of atoms partaking in the interferometry sequence per second, requiring: the most number of cold atoms, high fidelity in the interferometer pulses, and effective state detection.

Due to the high level of system integration, some instrument parameters cannot be easily measured, so must be optimised indirectly. In the case of the MOT characterisation, the presented data sets are from an earlier level of system integration but are representative of the current system. All other characterisation data is taken for the current fully integrated system.

5.1 Cold atom preparation

5.1.1 MOT characterisation

The magneto-optical trap (MOT), provides the trapping and first stage cooling of the atoms in the experiment. Since the sensitivity of the interferometer is dependent on the number of atoms, it is important to optimise the system for a total atom flux, the initial stage of which requires the optimisation of the MOT. Maximising the number of atoms in the trap requires the following optimisations: cooling and repump laser frequency, cooling laser intensity, cooling beam polarisation, MOT trapping field strength, and the rubidium and background vacuum pressures. Not all of these characterisations are presented.

MOT dynamics

Describing the dynamics of the MOT can lead to further characterisation of the vacuum pressure, and so the following theory is introduced. The dynamics of the MOT can be well approximated by the rate equation [103, 104]:

$$\frac{dN}{dt} = \alpha P_{Rb} - \gamma N(t) - \beta P_{Rb} N(t), \tag{5.1}$$

which describes the loading rate of atoms into the trap, dN/dt, by loading and loss terms dependent on the MOT parameters. The atom loading rate is described by the MOT trapping cross-section, α, and the partial pressure of Rb, P_{Rb}. Loss terms in the system are dominated by collisions with background Rb atoms through the constant β, and collisions with non Rb atoms grouped together as the rate term γ.

The solution to the rate equation has an exponential form, with the number of atoms, N, at a time, t, given by

$$N(t) = N_{eq}\big(1 - \exp(-t/\tau)\big), \tag{5.2}$$

with an exponential time constant, τ. Fitting to MOT loading curves allows the values of the equilibrium atom number, N_{eq}, and τ to be found. The terms α, β, γ, can also be calculated by varying the background Rb pressure, achieved by changing the dispenser currents.

MOT frequency, intensity, and polarisation

For the magneto-optical trap, the cooling light needs to be red-detuned, as discussed in section 2.1.1. There is an optimum detuning for the MOT, determined in part by the capture velocity [52]. Thus, in order to max-

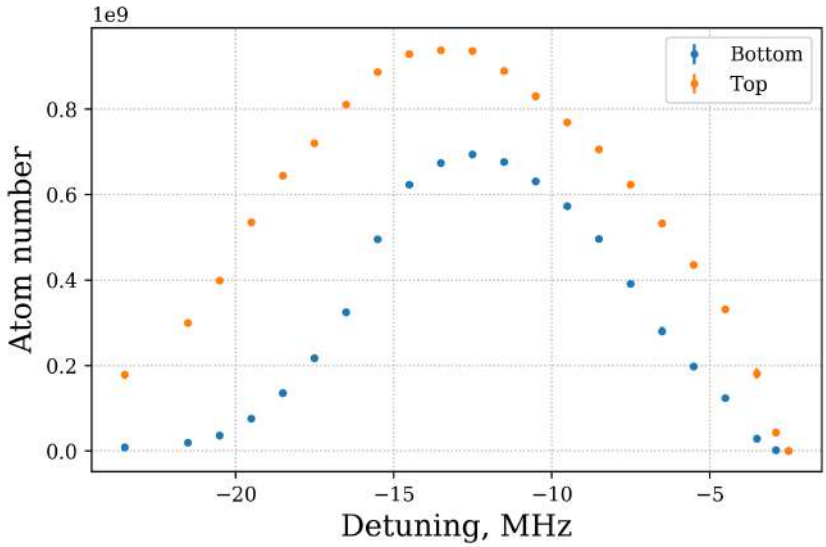

Figure 5.1 Changing the frequency detuning of the cooling light changes the estimated number of atoms in the magneto-optical trap. The optimum frequency for both MOTs is at −13 MHz. There is a slight offset between the two clouds, and a difference in the shape of the response, pointing to differences in beam alignment.

imise the atom number in the MOT, the cooling frequency is scanned whilst keeping the repump light on resonance. Figure 5.1 shows the number of atoms vs the cooling frequency detuning from the $|F = 2\rangle \leftrightarrow |F' = 3\rangle$ transition. The atom number peaks for the MOT cloud in each chamber, at a detuning close to twice the transition linewidth, and is set to −13 MHz from resonance. There is a slight offset between the two clouds, which may be explained by differences in alignment of the cooling telescopes.

In the optimisation of the intensity, it was found that more cooling laser power always gave the larger atom number, so the lasers were set to output maximum power, resulting in a cooling central intensity of around $7.5\,\mathrm{mW\,cm^{-2}}$ from each cooling telescope.

Cooling beam polarisation was optimised by setting the variable retarder plates to give circular light, measured off of the system using a polarimeter. Then as an assembled system, the VRP voltages are tweaked to maximise

the atom number in both chambers.

MOT loading

With the intensity, frequency, and polarisation set, the MOT coils are set to 2.5 A. The free parameter to increase the atom number and loading rate, is the amount of rubidium and background gas in the vacuum chamber. In this system, dispensers are used to provide rubidium vapour in a reasonably deterministic way. Measurements taken close in time will yield similar results; however, as the rubidium is depleted from the dispenser, the output of the dispenser, for the same driving current, may yield different results.

In order to optimise the dispenser current, the current was increased and the system left to equilibrate for a few minutes between measurements. Several MOT loading curves were taken at each dispenser current, and fit using equation 5.2, to extract the equilibrium atom number, N_{eq}, and the loading time τ.

Figure 5.2 shows a comparison between the two chambers, at the maximum dispenser current tested. Here, the loading time is similar; however, the atom number differs by nearly a factor of three. This difference may be explained by alignment of the detection optics, or a difference in the background partial pressures. That they are the same order of magnitude is sufficient for gradiometry to be performed.

Figure 5.3 shows the equilibrium atom number versus the dispenser current for both chambers, where the MOT is in a steady state. Higher dispenser currents in the bottom chamber shows a saturation of atom number not seen in the top chamber. Despite there being more Rb in the chamber, the equilibrium atom number is not increasing. This can be understood through putting the measurable quantities, N_{eq} and τ, in terms of the load-

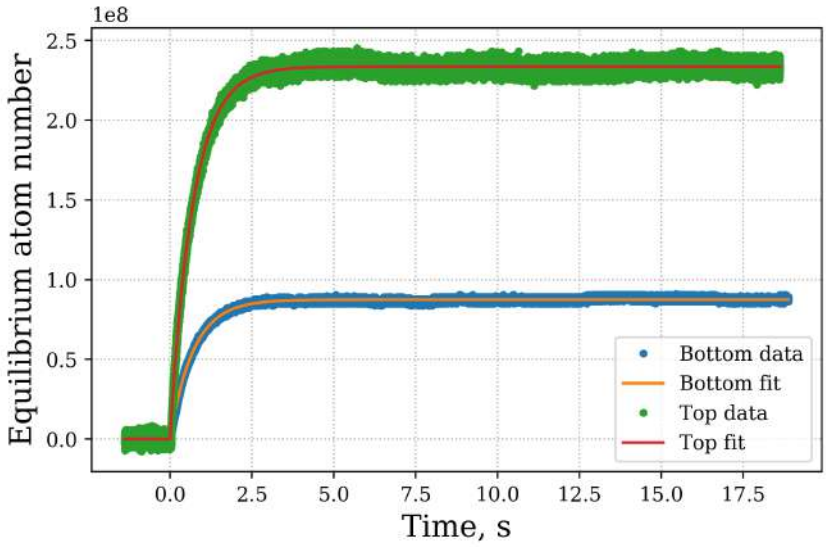

Figure 5.2 Comparison of the MOT loading curves for both chambers at their respective highest dispenser current. Loading time is similar but a difference in signal size of nearly a factor of 3 is seen. The MOT coil current is set to 2.5 A.

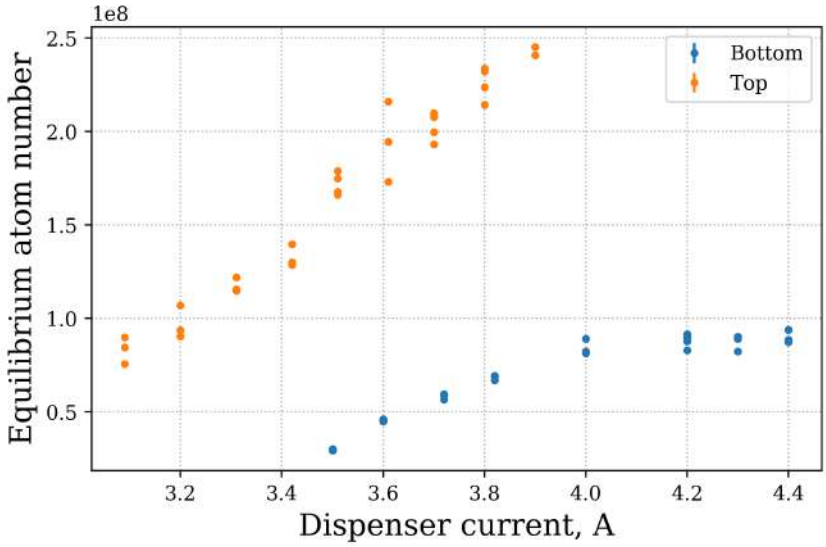

Figure 5.3 Estimated equilibrium atom number from fluorescence versus dispenser current. Saturation of equilibrium atom number is seen in the bottom chamber at high dispenser currents but not in the top chamber.

ing and loss rate terms, such that

$$N_{eq} = \alpha P_{Rb} \tau \tag{5.3}$$

and

$$\tau = \frac{1}{\beta P_{Rb} + \gamma}. \tag{5.4}$$

Equation (5.3) therefore shows that it is possible for the atom number to remain constant, if the loading rate of atoms into the trap is equal to the collisional loss of atoms from the trap. Plotting instead for the loading time, τ, versus the dispenser current, shown in figure 5.4, also begins to show high dispenser currents reducing the change in loading time. Finally, the equilibrium atom number and the loading time can be plotted against each other, seen in figure 5.5. This graph shows a linear relationship between the two, independent of the dispenser current, and can be used to calculate the partial pressures of the rubidium and other gas load from the loading rate constants α and γ [103]. The gradient of the line is given by the combination of equations (5.3) and (5.4)

$$N_{eq} = \frac{\alpha}{\beta}(1 - \gamma\tau). \tag{5.5}$$

To calculate the non-rubidium background pressure from the rate constant γ, the conversion $P = \gamma \cdot 1.53 \cdot 10^{-8}\,\mathrm{s\,mbar}$ may be used [103, 104]. Using the fit parameters from the lines in figure 5.5, the non-rubidium partial pressure is calculated to be $(3.0 \pm 0.1) \cdot 10^{-9}$ mbar and $(2.15 \pm 0.09) \cdot 10^{-9}$ mbar for the top and bottom chambers respectively. This value is higher than the ion pump readings, which is expected due to the conductance from the pump to the trapping region. The linearity of the graph also suggests that the background gas load is negligibly changed by the dispenser currents tested.

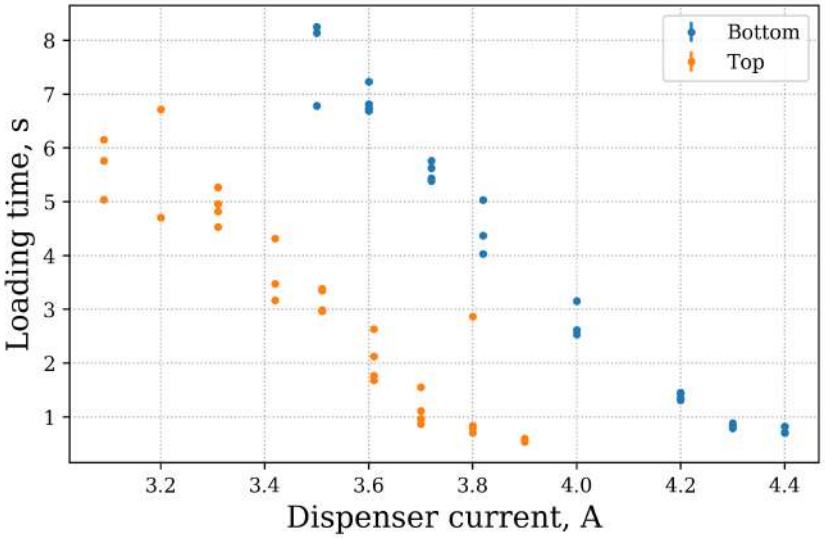

Figure 5.4 For different dispenser currents, the loading time constant is taken from the fit to the loading curve. An offset between the two chambers is seen, where the bottom dispensers must be driven at a higher current to achieve the same loading rate. The loading time constant is tending towards a maximum value at high dispenser currents.

From the same fit, the value of $1/\gamma$ gives the upper limit of the trap lifetime as $(9.2 \pm 0.5)\,\mathrm{s}$ and $(12.7 \pm 0.5)\,\mathrm{s}$, for the top and bottom chamber respectively. As determined by the calculations in section 3.1.1, the background pressure of the vacuum system meets the requirements to allow for atom evolution times of at least 500 ms.

Finally, the maximum possible atom number achievable in the system is calculated from the y intercept of figure 5.5. For the top and bottom chamber this is $(2.45 \pm 0.07) \cdot 10^8$ and $(9.9 \pm 0.2) \cdot 10^7$ atoms, respectively, which is 1000 times the modelled atom number used in determining the instrument requirements (section 2.4.5 onwards). Atom loss is expected in the molasses sequence, state selection, and interferometer contrast, so detection of the atoms after the interferometry sequence will reveal the number of atoms partaking the interferometer. This is measured in the appendix D.3.

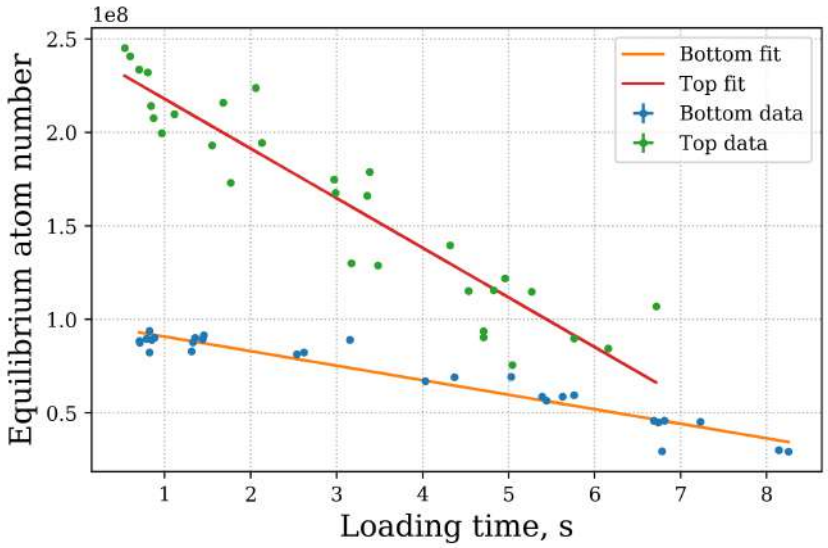

Figure 5.5 Equilibrium atom number vs loading time. Linear fit to equation (5.5) allows the rate constants α, β, γ to be calculated, which in turn may be used to evaluate the vacuum quality and MOT performance [103].

5.1.2 Optical molasses

In order cool the atoms to below the temperature achieved after the MOT phase, sub-Doppler cooling processes, typically referred to as an optical molasses, are used to lower the ensemble temperature to the recoil limit. To create an effective optical molasses, magnetic fields are turned off, the light level reduced, and the optical frequency further red detuned from the MOT phase. As discussed in section 3.8, bulk magnetic shielding is used as the primary method for achieving low background magnetic fields in the sensor package, with additional compensation coils being used to reduce any remaining field.

The molasses sequence was optimised for the coldest temperature, whilst maximising atom number, and achieved through an iterative optimisation process. The resulting sequence is shown in figure 5.6, and outlined here. At the end of the MOT phase, the coils and MOT light are turned off for

3.25 ms to allow the eddy currents from the coils switching to decay.* The light is turned back on at a reduced power, giving a central intensity of around 1.3 mW cm^{-2}, and the frequency ramped from the MOT frequency by 50 MHz in 1.5 ms. In the second stage, the power is reduced further and the frequency ramped by an additional 47 MHz in 11.5 ms. From here the intensity is quickly ramped with the AOM to let no light through in 0.1 ms. To prepare the atoms in the $|F = 1\rangle$ ground state, the repump frequency is turned off, by switching the RF driving the EOM, 4 ms before the cooling intensity is reduced to zero. The timing of the state selection is characterised in figure 5.7. Many other sequences rely on frequency jumps and power ramps to achieve the coldest temperatures [9, 11, 67, 105]; however, this scheme proved to be optimal for this setup. It was found that the bulk shielding is sufficient to remove external fields and compensation coils were not required to achieve the coldest temperatures, requiring only the optimisation of the intensity and frequency of the cooling laser.

Initial state preparation

To prepare the atoms in the hyperfine ground state $|F = 1\rangle$, the repump frequency is turned off before the end of the molasses sequence. Without the repump light present, atoms are quickly lost from the cooling process as they excite the $|F = 2\rangle \leftrightarrow |F' = 2\rangle$ transition, and may fall into the $|F = 1\rangle$ state. Once in the $|F = 1\rangle$ state, no further cooling may take place, requiring this step to be as short as possible. To characterise the time required to prepare the atoms in the $|F = 1\rangle$ state, the time before the end of the molasses sequence is scanned and shown in figure 5.7. Here the population ratio of atoms in the $|F = 2\rangle$ state is compared to the total population, and

*Without this step, the cloud temperature was limited to 20 µK.

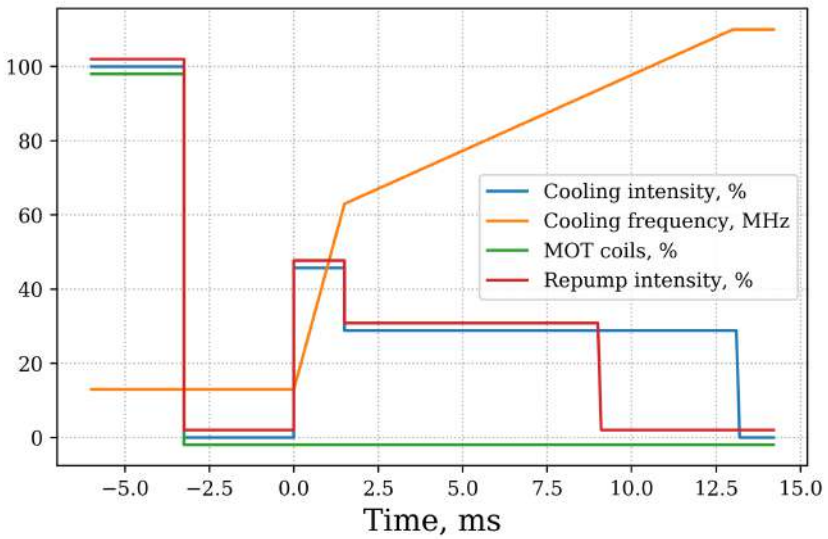

Figure 5.6 Optical molasses sequence used in the gradiometer. Cooling light is turned off for 3.25 ms to allow the eddy currents from the switching of the MOT coils to decay. The frequency is ramped from the MOT frequency by 50 MHz further red detuned in 1.5 ms, then by an additional 47 MHz in 11.5 ms. The repump frequency is turned off 4 ms before the end of the molasses sequence to prepare the atoms in the hyperfine ground state $|F = 1\rangle$. Magnetic fields and repump intensity have been offset by 2 % for clarity.

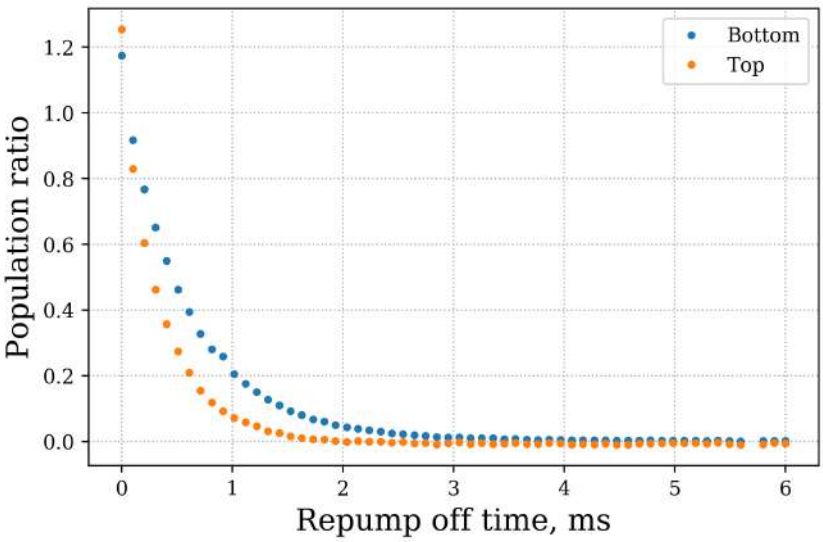

Figure 5.7 Changing the length of time the repump light is turned off for before the end of the optical molasses changes the population ratio of atoms in each state, using the detection sequence described in section 5.2. To ensure a clean state, the repump is turned off for the final 4 ms of the molasses sequence.

measured using the detection scheme described later in section 5.2. Despite some biasing seen in the detection here, the atoms are prepared into the $|F = 1\rangle$ ground state after the repump frequency is turned off 4 ms before the end of the molasses sequence. In the complete molasses sequence, a pulse of resonant cooling light is used to blow-away any atoms remaining in the $|F = 2\rangle$ ground state.

Temperature determination by time of flight measurement

To optimise the molasses parameters, the temperature and number of atoms in the ensemble is used as the figure of merit. This can be achieved with a time of flight measurement, performed using a light sheet [106, 107]. In this scheme, a thin sheet of resonant light is used to measure the fluorescence versus time from the atom cloud as it passes through the sheet. For sufficiently long fall times, the temporal width from the output of the light sheet

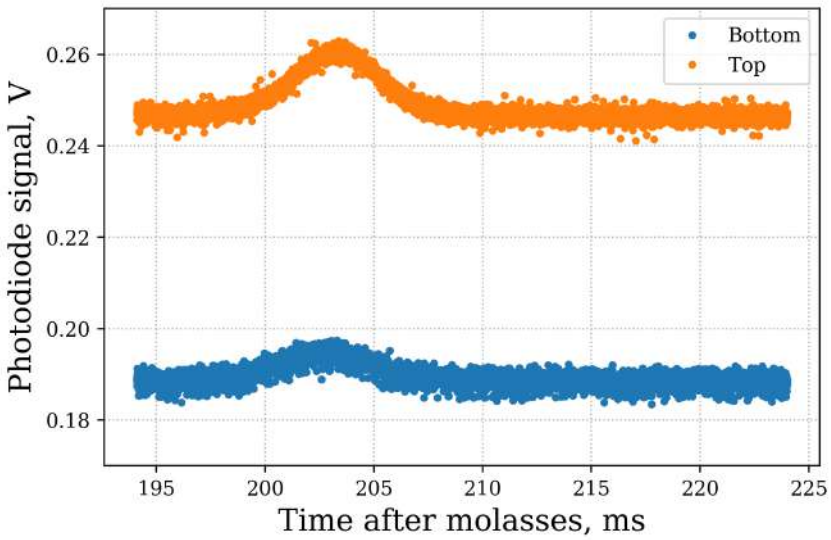

Figure 5.8 A typical light sheet trace for the two atom clouds. The atomic signal is small compared with the background light level.

follows a Gaussian profile and is related to the spatial width of the cloud at that time. Since the atomic ensemble expands ballistically, the spatial width of the cloud may be used to determine the ensemble temperature when an initial cloud size, σ_0, is given. From the spatial width, a velocity width, σ_v, may be calculated when the evolution time, t, is known. For evolution times of a few hundred milliseconds, and a finite width of the light sheet, σ_a, the distribution is better approximated by including terms dependent upon the light sheet thickness [107]. In the light sheet measurement, the integrated signal is proportional to the atom number, and when the Gaussian signal becomes taller and thinner, the atomic ensemble is becoming colder.

Following a molasses sequence, atoms prepared in $|F = 1\rangle$ are dropped and measured in the detection region, 203 ms after the end of the molasses sequence. A typical trace is seen in figure 5.8, where the atomic cloud signal for each chamber is shown. The size of the atomic signal compared to the background signal is small in this detection configuration, due to scattering

of the light sheet as it passes through the vacuum chamber. Since this measurement is only used to determine the temporal width of the ensemble, the poor signal to noise does not present an issue. The signal in figure 5.8 can be fit with

$$y(t) = \frac{A}{\sqrt{\sigma_a^2 + \sigma_v^2 t^2}} \exp\left[-\left(\frac{g(t_0^2 - t^2)}{2\sqrt{2}\sqrt{\sigma_0^2 + \sigma_v^2 t^2}} \right)^2 \right], \tag{5.6}$$

to extract the velocity width of the atom cloud, by fixing the initial cloud size and light sheet thickness. The ensemble temperature can then be calculated from the velocity width, using

$$T = \frac{m}{k_B} \sigma_v^2. \tag{5.7}$$

Using CCD cameras in the MOT region, the initial cloud radius is estimated as 2 mm. Based on the slit size used in the optic delivery, the light sheet should have a thickness of 300 µm; changing the value of the light sheet thickness in equation 5.6 an order of magnitude either way has less than 0.1 % effect on the calculated temperature. Using these parameters, the system performs repeated light sheet measurements, to establish the atom cloud temperature stability, with the results of a three hour measurement seen in figure 5.9. The temperature and arrival time show good stability over the measurement time, despite a 10 % cooling laser intensity drift (see appendix D.1 for further analysis of the light sheet data), demonstrating the robustness of the prism MOT geometry. The ensemble temperatures are averaged over the measurement time and calculated to be (2.86 ± 0.09) µK for the top chamber and (3.7 ± 0.2) µK for the bottom chamber, with the error calculated as the standard deviation of the temperatures. These ensemble temperatures are limited by the relative alignment of the MOT beams, but

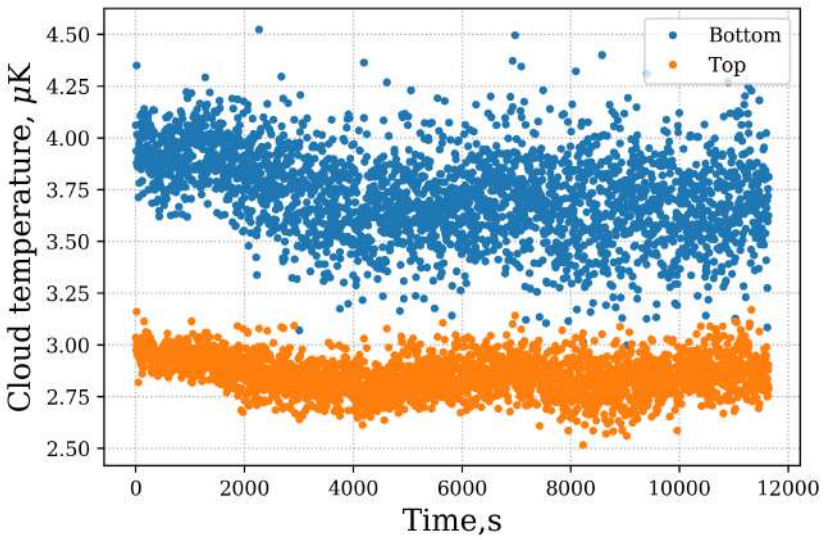

Figure 5.9 Extracted temperature for time of flight measurements, with figure 5.8 showing a representative measurement. Ensemble temperature is calculated as $(2.86 \pm 0.09)\,\mu\text{K}$ for the top chamber and $(3.7 \pm 0.2)\,\mu\text{K}$ for the bottom chamber with the error calculated as the standard deviation of the temperatures. The ensemble temperature shows resilience against 10 % cooling power fluctuations.

meet the requirements outlined in section 2.9.4.

5.2 Detection scheme

For the primary detection output, light delivered vertically from both cooling laser telescope ports is used to induce fluorescence in the atoms. Atoms in the $|F = 2\rangle$ hyperfine ground state may be measured by exciting the transition $|F = 2\rangle \leftrightarrow |F' = 3\rangle$. Because of the laser scheme implementation, light resonant with the transition $|F = 1\rangle \leftrightarrow |F' = 2\rangle$ is not available from the cooling laser, so to detect atoms in the $|F = 1\rangle$ ground state, the cooling light with the repump sideband is used to excite both states. Using these modes to produce fluorescence from the two ground states, the signal from atoms in each states can be represented as a relative population ratio through use of several light pulses.

The pulse sequence for the vertical detection scheme is shown in figure 5.10, where seven peaks, numbered from left to right, are presented. The first two peaks are used to measure the fluorescence from the atoms in the dropped clouds, where peak 1 contains light exciting atoms in state $|F = 2\rangle$, and peak 2 contains light exciting atoms in states $|F = 1\rangle$ and $|F = 2\rangle$. In order to separate the atomic signal from the background light level, and fluorescence from background atoms, peaks without the atomic cloud present are subtracted from the signal peaks. To prepare these background peaks, peak 3 excites both states, and is used as a blow-away to remove the atoms left from the atomic cloud, leaving only scattered light and background fluorescence. The pulses used in peaks 1 and 2 are repeated in peaks 4 and 5, and are used for background subtraction. As well as removing an offset, the subtraction also removes intensity fluctuations on that timescale, since the level of the background peak is removed from the atom peaks.

To calculate the relative population of $|F = 2\rangle$, as a ratio of the population in states $|F = 1\rangle$ and $|F = 2\rangle$, the following is done

$$\frac{P_{|F=2\rangle}}{P_{|F=1\rangle} + P_{|F=2\rangle}} = \frac{V_1 - V_4}{V_2 - V_5}, \tag{5.8}$$

where V_i is the integrated voltage of peak i from the detection photodiode signal. This calculation of population ratio requires that atoms detected in peak 1 are also present in peak 2, which can be achieved by using pulses that are short and closely spaced in time. In this implementation, each detection pulse is 100 µs, which has been found to satisfy the normalisation of equation (5.8). Longer detection pulses would need a scale factor to account for atoms lost in peak 1. For detection, the intensity is reduced from that used in the MOT, with an intensity of 2.4 mW cm^{-2} used.

As well as being used to blow-away the atomic cloud, peak 3 may be used

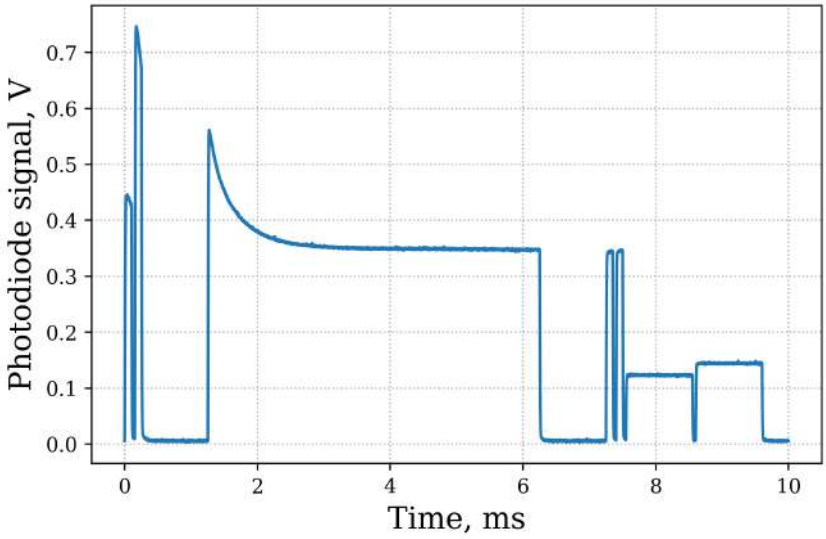

Figure 5.10 Dection scheme used to characterise population ratio. Background subtracted ratio of peak 1 over peak 2 gives the relative population of atoms in $|F = 2\rangle$. Peak 1 has light resonant with $|F = 2\rangle$, peak 2 has light resonant with $|F = 1\rangle$ and $|F = 2\rangle$. Peak 3 blows away any remaining atoms, and peaks 4 and 5 are background peaks for 1 and 2. Peaks 6 and 7 are diagnostic peaks for each Raman port.

as a diagnostic peak, and is especially useful to determine if the detection frequency is resonant, by looking at the shape of the decay. Peak 6 and 7 are pulses from Raman laser 1 and 2, and are used as diagnostic peaks, to determine the stability of the Raman laser with time.

During the detection pulse sequence, the frequency of the cooling laser is tuned to maximise the amount of fluorescence in peak 2, for an atom cloud prepared in the $|F = 1\rangle$ ground state. A scan of the detection frequency is shown in figure 5.11, shows a sharp asymmetric peak at a frequency of 3 MHz detuned from the resonance of a stationary atom. For this drop time, two broad peaks are expected from the Doppler shift at ± 5 MHz, corresponding to the atoms being resonant with light from cooling beams delivered from the top and bottom telescopes. This suggests additional mechanics are at play, with a possible explanation being that the vertical detection is forming a 1D-molasses with the falling atom cloud.

Henceforth, this detection scheme will be used for detection and characterisation of the atoms. The suitability of this detection scheme is evaluated in section 6.2.1. The number of atoms can be determined using this detection scheme, shown in appendix D.3.

5.3 Spectrum scan

As discussed in section 2.2.1, the two-photon Raman transition is a velocity selective process, when delivered in a counter-propagating configuration, with the Fourier limited linewidth inversely proportional to the pulse length. For Raman pulses with a length of tens of µs, the linewidth is of order 50 kHz. This means that the frequency of the Raman transition must be set to within a few kHz in order to be resonant with the transition. In addition, Doppler sensitivity requires the Raman frequency to be linearly chirped to maintain

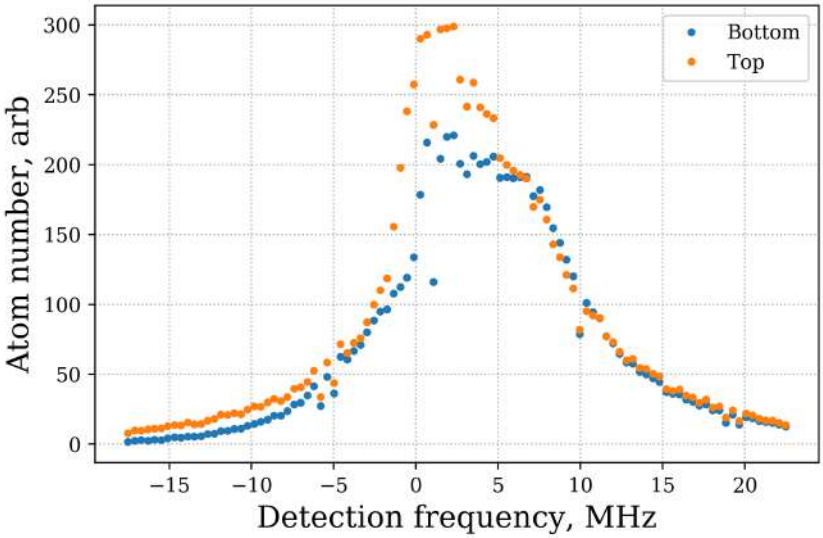

Figure 5.11 Scan of the detection frequency from the resonance of a stationary atom. The optimum detection frequency is set to 3 MHz. The asymmetric shape suggests additional mechanisms to simple Doppler shift.

the resonance condition with the falling cloud.

Through scanning the two-photon detuning δ, the resonant frequencies of the available Raman transitions can be observed. The atoms have been prepared in the hyperfine ground state $|F = 1\rangle$, distributed amongst the three degenerate m_F sub-levels. An applied bias field along the Raman beam direction sets the quantisation axis and lifts the energy degeneracy of the m_F magnetic sub-levels, with a frequency change of $0.7\,\mathrm{MHz\,G^{-1}}$ as determined in section 2.3 [47]. With the polarisation of the variable retarder plates set to produce perpendicular linear light, the only allowed transitions will be those driving $\sigma_+\sigma_+$ and $\sigma_-\sigma_-$ transitions. Through inspection of the Clebsh-Gordon coefficients [47], the only allowed transitions between the two hyperfine ground states, via the intermediate state, are $\Delta F = 1$ and $\Delta m_f = 0$. These correspond to the transitions: $|F = 1, m_F = -1\rangle \rightarrow |F = 2, m_F = -1\rangle$, $|F = 1, m_F = 0\rangle \rightarrow |F = 2, m_F = 0\rangle$, $|F = 1, m_F = 1\rangle \rightarrow$

$|F = 2, m_F = 1\rangle$. When scanning the two-photon detuning, three peaks are therefore expected in the spectrum, separated by a frequency difference of $2 \times 0.7\,\mathrm{MHz\,G^{-1}}$, proportional to the applied bias field.

To enable narrow peaks, and thus accurate determination of the central peak position, the Raman power is attenuated to reduce the Rabi frequency and enable the use of longer pulses. In the following, a Raman pulse length of 10 µs was used, with a frequency resolution of 15 kHz. A sample spectrum scan is shown in figure 5.12, where measurements are taken simultaneously in both chambers. In each case the three m_F states are clearly resolved. The populations are distributed evenly between the m_F peaks in the top chamber, but a gradient is seen across the m_F peaks for the bottom chamber. This is likely due to optical state pumping, which may be caused by beam alignment in the molasses. The Lorentzian line shapes are well approximated by Gaussian functions, therefore multiple Gaussian peaks are fit to the data. From this, the central resonant frequency for each state, as well as the amplitude and width, are extracted.

5.3.1 Chirp rate calculation

By measuring the resonance of the transition $|F = 1, m_F = 0\rangle \leftrightarrow |F = 2, m_F = 0\rangle$ at each drop time, the chirp rate applied to the Raman sideband can be measured for both interferometers.* Spectrum scans scans are taken for different atom fall times, from 7 ms to 196 ms, relative to the end of the molasses sequence. This is the maximum usable drop region of the 203 ms atom fall time, since time is required for state preparation and ramping to

*The chirp rate may also be determined by a more traditional gravimeter chirp rate scan, where the rate is changed for different T time fringes in order to find the rate at which cancels the phase shift due to gravity. To determine the chirp rate to more decimal places than extracted from the spectrum scan, good vibration isolation and low phase noise is required.

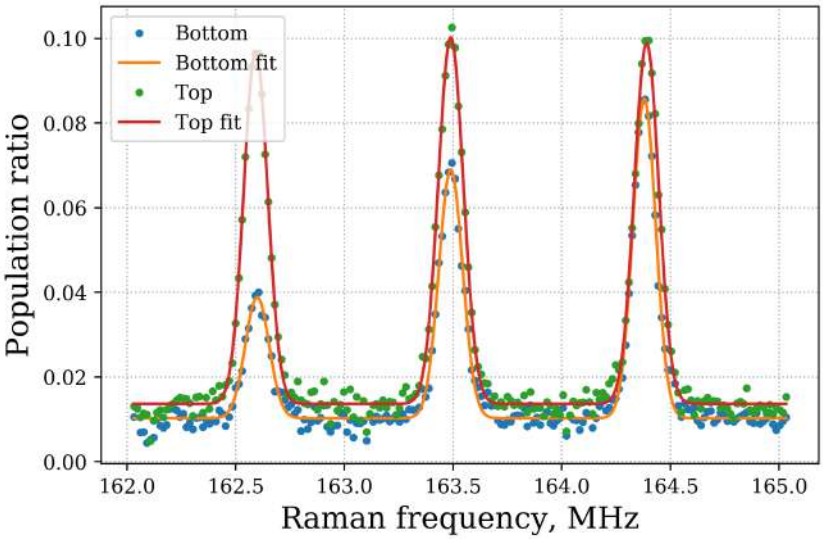

Figure 5.12 Single spectrum scan taken at 70 ms after the end of molasses, with a 10 µs Raman pulse and a bias coil current of 100 mA. Three peaks can be seen for each chamber, corresponding to the resonances of the three transitions between the m_F states, separated by the first order Zeeman shift of $2 \times 0.7\,\mathrm{MHz\,G^{-1}}$. The top chamber shows an equal population between each of the sub-levels; however, the bottom chamber shows uneven population distribution, suggesting optical state pumping.

the detection frequency.

The peak centres are obtained for all three m_F resonance peaks, and are shown in figure 5.13 for the top chamber. A linear fit to the central $m_F = 0$ peak centre, gives a value of the chirp rate to first order. The same analysis is performed for both chambers, yielding a value of the chirp rate as $(-25.164 \pm 0.003)\,\mathrm{MHz\,s^{-1}}$ and $(-25.152 \pm 0.003)\,\mathrm{MHz\,s^{-1}}$, for the top and bottom chamber respectively. There is slight difference in the chirp rate in each chamber which is larger than would be expected due gravity gradient alone. This suggests that the two clouds are being accelerated differently, which is also evidenced in the arrival time data (appendix D.1). This may be caused by a cloud having unbalanced cooling force in the molasses sequence, creating an additional acceleration to gravity, or by shifts in the resonant frequency due to light shifts. For the Raman pulse lengths of interest, the chirp rate is not needed to a higher precision, since short pulses are used to address a large fraction of the atoms; this also mitigates the issue with the chirp rates being slightly different for each interferometer. Over the length of the drop, the difference in chirp rate only corresponds to a shift of 2 kHz; negligible fidelity will be lost in the Raman pulses due to this difference.

5.3.2 Magnetic field profile

The difference between the transition frequencies of the $\Delta F = 1, \Delta m_F = 0$ peaks is used to calculate the magnetic field strength at the atom cloud position for that measurement time. The measurement at that point is of the total field, not just the z magnetic field, with the separation given by the first order Zeeman shift, $2 \times 0.7\,\mathrm{MHz\,G^{-1}}$ in the low field regime [47]. The magnetic field profile measured by the atoms is shown in figure 5.14, for a bias coil current of 100 mA. The field profile shows a variation of 3 % across

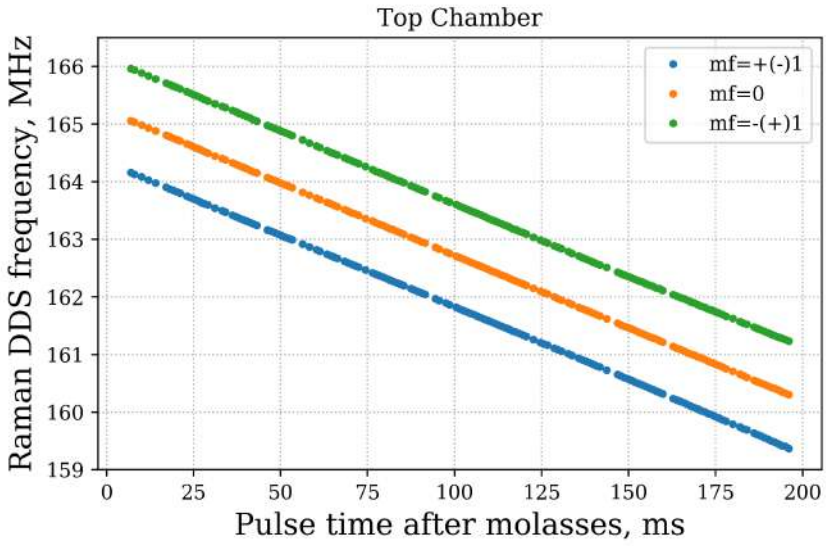

Figure 5.13 Peak positions from the spectrum scans at each height, shown for the top chamber. The peak corresponding to a transition between $m_F = 0$ is the resonance of interest. Fitting a line to the peak centre values gives the chirp rate to first order. The chirp rate is measured as $(-25.164 \pm 0.003)\,\mathrm{MHz\,s^{-1}}$ for the top chamber, and $(-25.152 \pm 0.003)\,\mathrm{MHz\,s^{-1}}$ for the bottom chamber.

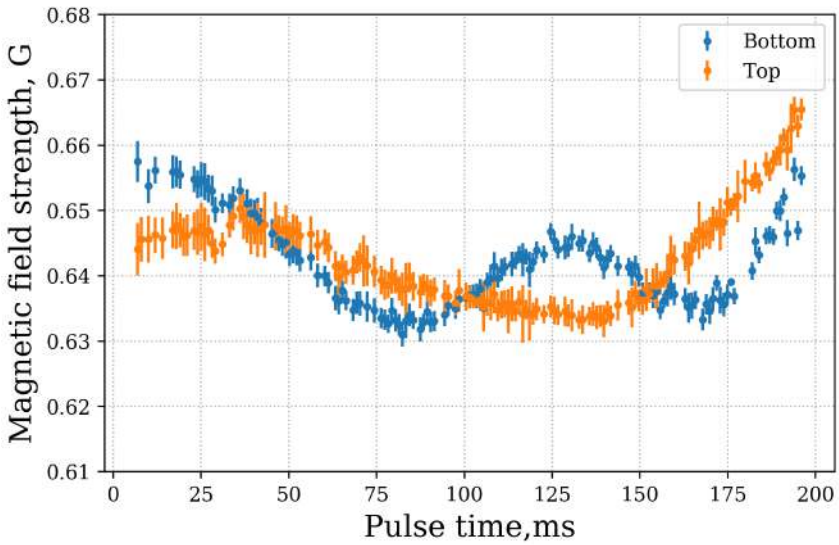

Figure 5.14 Total magnetic field measured by the atoms using the difference in resonant frequency between the peaks corresponding to transitions between $|F = 1, m_F = 2\rangle \rightarrow |F = 0, m_F = 0\rangle$ and $|F = 1, m_F = 2\rangle \rightarrow |F = \pm 1, m_F = \pm 1\rangle$. The total field contains the bias coil field and any background magnetic field

the total measurement region. To investigate the variation, the spectrum scan across the drop region is repeated at half the bias coil current (50 mA). Since the measurement of magnetic field, B_{tot}, contains the magnetic field from the bias coil, B_{bias}, and the background field, B_{back}, measurements at half the current allow for the extraction of the background field and bias field using the equation

$$B_{back} = 2B_{tot}^{half} - B_{tot}^{full} = 2\left(\frac{B_{bias}}{2} + B_{back}\right) - \left(B_{bias} + B_{back}\right), \quad (5.9)$$

and

$$B_{bias} = B_{tot}^{full} - B_{back}. \quad (5.10)$$

Applying equation (5.9) to the data gives the background field profile for the atomic trajectory. This is seen in figure 5.15, where the background field is plotted with respect to the distance from each MOT centre. The bottom chamber shows a slightly stronger variation in the background field than the top chamber, which could suggest that the shielding is less effective in that region, or that there is some magnetisation of the shields.

Using the background field profile, the bias field profile for each chamber is calculated using equation (5.10). This is shown in figure 5.16, where the distance is given with respect to each MOT region. The profiles show a variation from uniformity with the total variation across the length of interferometry region equal to $(3.1 \pm 0.4)\,\%$ and $(4.7 \pm 0.4)\,\%$ for the top and bottom chamber respectively. This variation is 6 to 9 times larger than that from the modelled field profile of the bias coil, and is seen more clearly when overlayed with the modelled field profile (figure 5.17). It should be

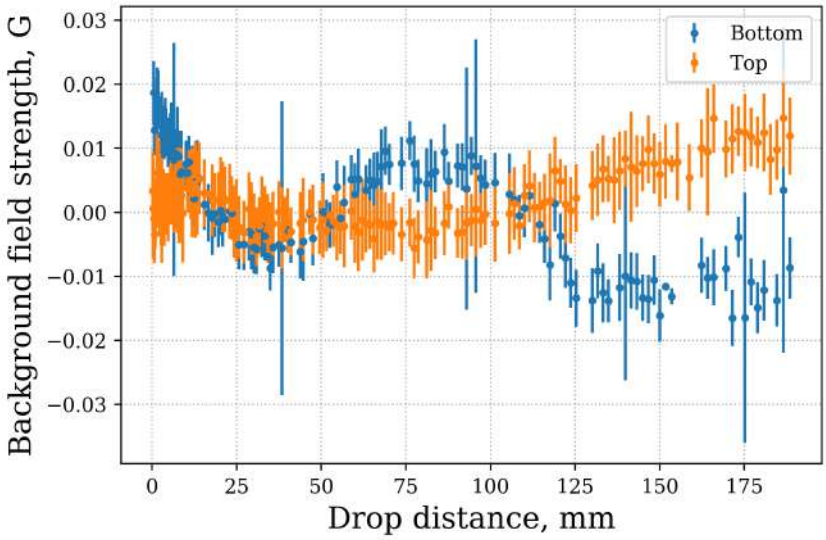

Figure 5.15 Background field calculated using equation (5.9) by taking spectrum scan measurements at two different bias coil currents.

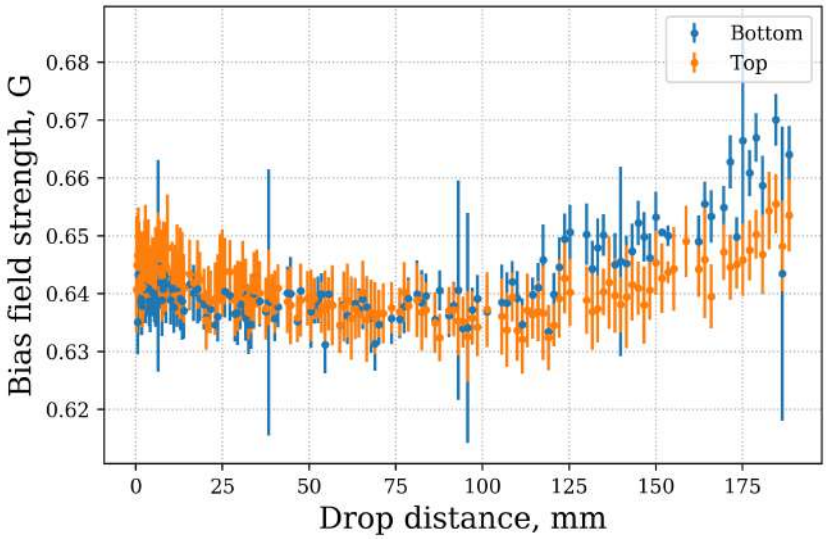

Figure 5.16 Bias field calculated using equation (5.10) by subtracting the background values shown in figure 5.15 from the total magnetic field measured in figure 5.14. Top chamber shows a peak to peak variation of $(3.1 \pm 0.4)\,\%$ across its interferometry region. Bottom chamber shows a larger variation of $(4.7 \pm 0.4)\,\%$ over its interferometry region.

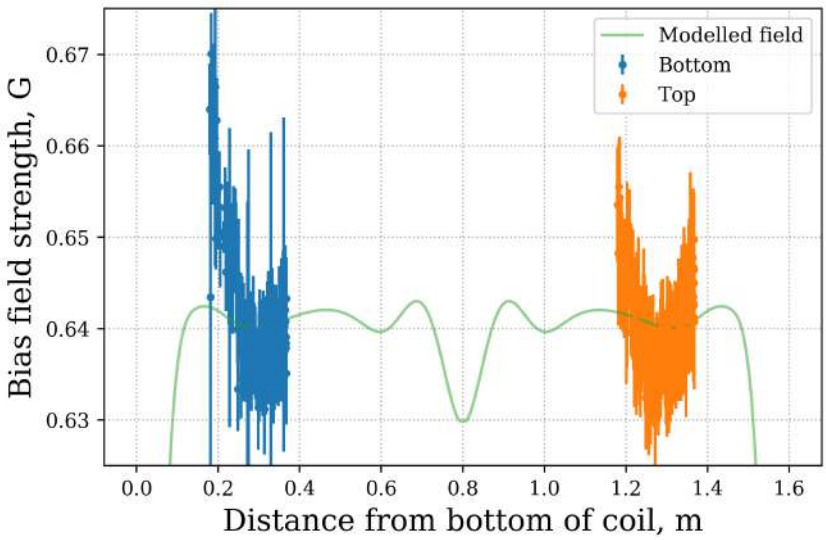

Figure 5.17 The measured bias field profile overlayed with the modelled field profile calculated in section 3.5. A significant deviation from the modelled field profile is seen, suggesting that inside the system, the field profile is distorted when integrated into the experiment and the mumetal shielding.

noted that the method used to calculated the field at each point can only be used to distinguish the magnitude of the field; the variation may be therefore be induced by the coil along the radial axes, instead of the z direction.

Regardless of the calculated bias field profile, the magnetic field profile seen by the atoms is that of figure 5.14, where the background field is present. Since the bias field is generated by a single coil of fixed dimensions, the field profile cannot be tuned by changing the current density. The gradients in the magnetic field will result in phase shifts between the two interferometers, as discussed in section 2.3, and so must be kept constant to avoid phase noise. A stable, low noise power supply is therefore used to drive the bias coil in order to maintain the stability of the field profile.

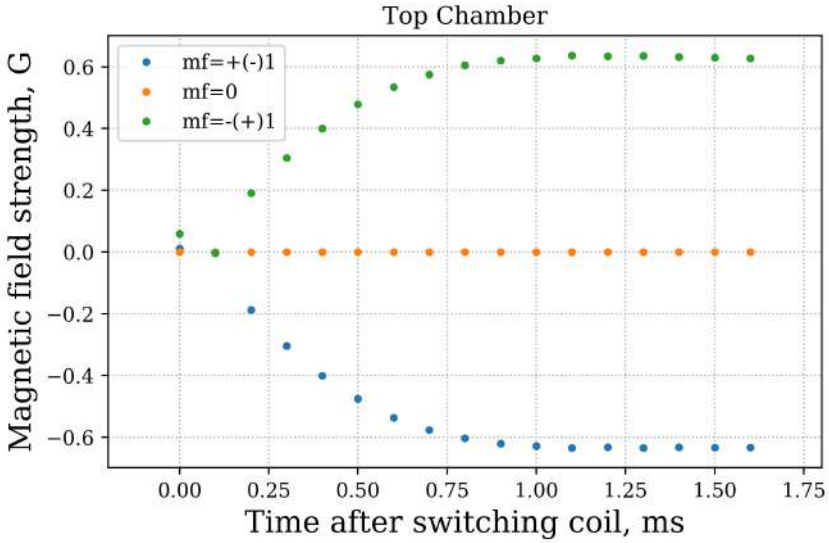

Figure 5.18 Bias coil switching on, as measured by the separation of the m_F peak centres. Current is switched on, via the IGBT, into the bias coil at 0.1 ms, stabilising after roughly 1 ms. Fitting the rate equation (5.11) to the data gives a time constant of (0.2680 ± 0.0004) ms.

5.3.3 Bias field switching speed

Due to the high level of integration, a representative switching speed of the bias coil is best determined by using the state separation of the atom's m_F states. The switching speed of the bias coil, for both on and off, is measured using this method. m_F peak centres are shown versus time for the coil switching on, in figure 5.18. The data is fit to the rate equation

$$B_z(t) = A(1 - \exp(-t/\tau_{coil}), \tag{5.11}$$

to determine the switching time constant τ_{coil}. For the coil switching on, the time constant is measured as (0.2680 ± 0.0004) ms. In the measurements of the coil switching off, the time constant is measured as (0.211 ± 0.005) ms. To ensure a stable magnetic field strength, times of a few millisecond should be allocated, to allow for many time constants after the current is switched.

5.4 Rabi oscillations

To determine the length of the $\pi/2$ and π pulses used in the interferometer, Rabi oscillations are measured at each of the pulse heights in the interferometry region. Each oscillation is taken by scanning the Raman pulse length, and measuring the ratio of population transferred between $|F = 1, m_F = 0\rangle \leftrightarrow |F = 2, m_F = 0\rangle$. As, in these measurements, the system is operating without sub-level state selection, discussed later in section 5.5, the maximum expected ratio is when all of the atoms from $m_F = 0$ are transferred. The population ratio is therefore limited to approximately $1/3$, based upon the the state occupation levels shown in figure 5.12. To measure the π time and the fidelity of the Rabi oscillation, the data at each pulse height is fit using the damped oscillation

$$\frac{P_{F=2}}{P_{F=1} + P_{F=2}}(t) = A\left\{\left[\sin^2\left(\frac{\Omega t}{2}\right) - \frac{1}{2}\right]\exp(-\Gamma t) + \frac{1}{2}\right\}. \tag{5.12}$$

The damping term is required as the atomic cloud diameter, approximately 4 mm, is a comparable size to the Raman beam width, 15 mm, which as discussed in section 2.9.4, creates multiple Rabi frequencies across the atom cloud. In lieu of a multi-frequency model, a damping term reproduces the behaviour between Raman pulse times $0 - \pi$. In the damped oscillation case, the π time is taken to be the Raman pulse length that provides the maximum population transfer, with the $\pi/2$ time defined similarly at half the population transferred. The damping means that the π time is no longer twice the length of the $\pi/2$ time and is observed to be nearly four times longer. The maximum population transfer is also measured at the π time.

The Raman pulse length scans are recorded at every 1 ms of the atom fall

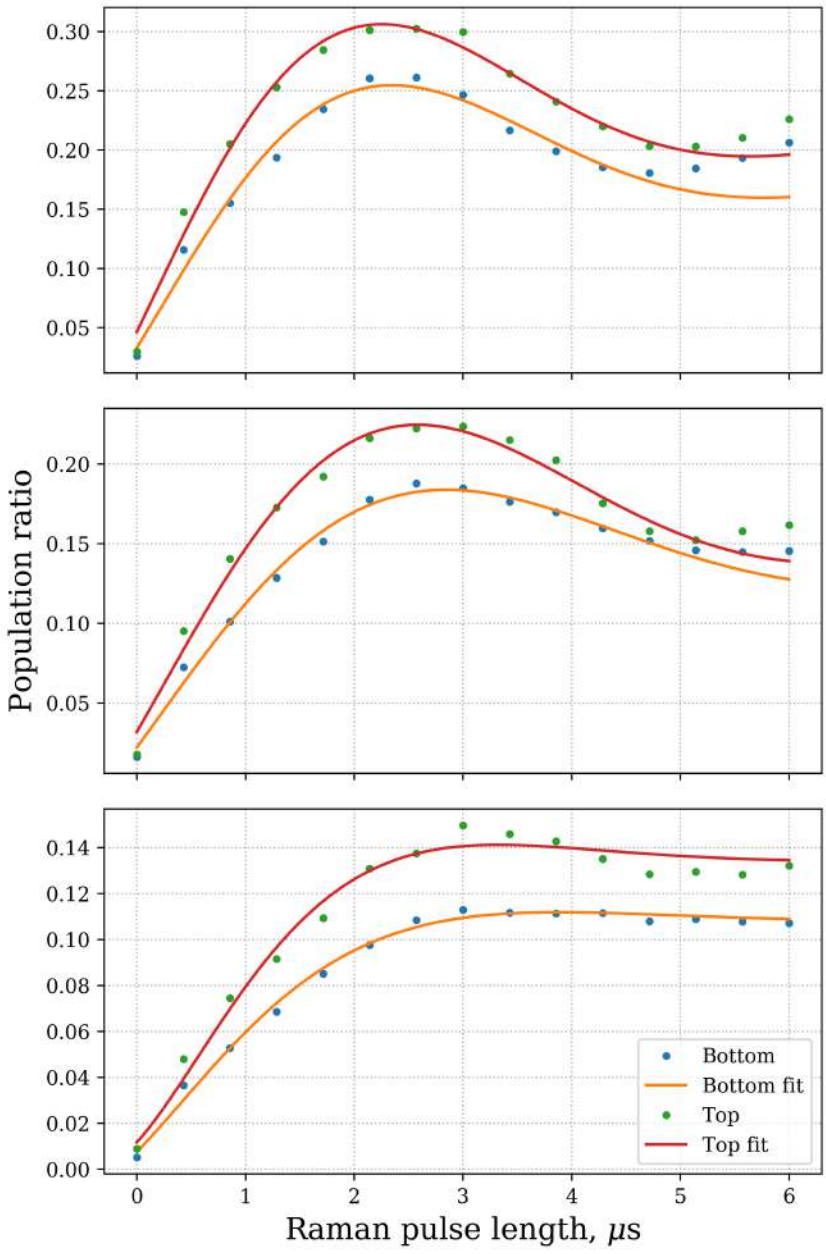

Figure 5.19 Typical Rabi oscillations at the three interferometer pulse times of 17 ms, 102 ms, and 187 ms, without sub-level state preparation. π time of the Rabi oscillation becomes longer as the Raman pulse scan is taken later in the interferometry region. Additionally, a drop in transfer of population is observed due to dephasing.

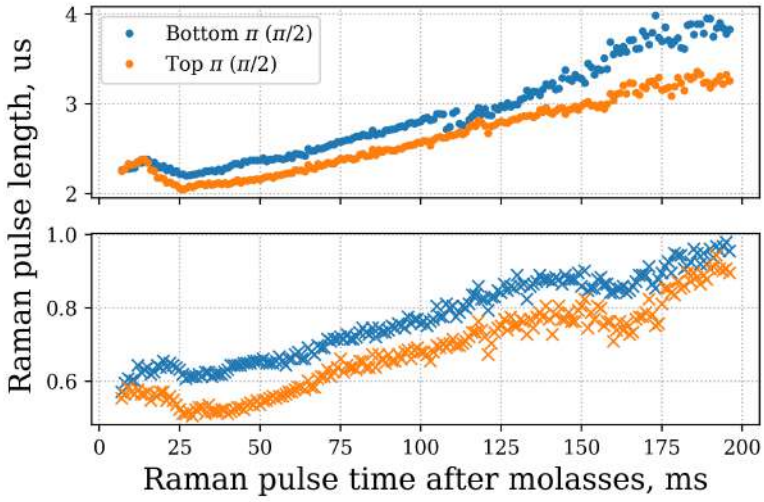

Figure 5.20 Measured π (top panel), and $\pi/2$ (bottom panel), times from the Raman pulse scans at 1 ms intervals of the interferometry region.

time, with the resulting oscillations fit using equation (5.12). Only the first 5 µs are used in the fit, since the damped oscillation fails to reproduce the Rabi dynamics at later times. Three example Rabi oscillations are shown in figure 5.19, at the pulse heights used later in the interferometry sequence. The Raman pulse scans show that as the atoms fall through the chamber, the Rabi π time becomes longer, and the population excited by the Raman transition, at the π time, decreases. The measured π and $\pi/2$ times are shown for the full interferometry regions in figure 5.20, which shows a predominantly linear increase of the π, and $\pi/2$, time with fall time, for both the top and bottom interferometer, consistent with a linear change in the atom cloud size.

The maximum population transfer is shown in figure 5.21 and shows a mostly linear drop off at times past 25 ms. In the time between 7 ms and 25 ms, an increase in transfer is seen. This is due to overlap with the co-propagating transitions, which can be removed by further optimisation of

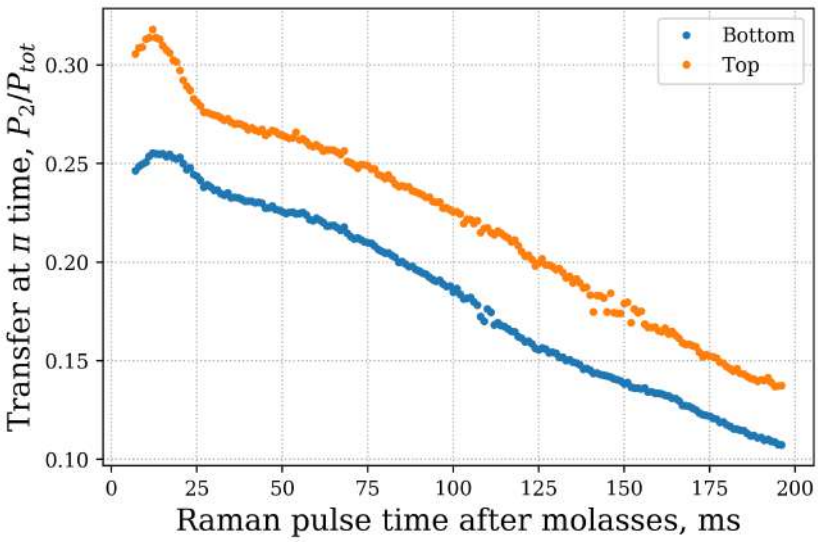

Figure 5.21 Population ratio of atoms in $|F = 2\rangle$ state to the total population at the measured π time. The transfer decreases with drop length due dephasing.

the output of the Raman telescopes, to produce perpendicular linear polarisations. Coupling of the co-propagating transitions will result in a reduction of contrast in the interferometer and should generally be avoided. This may be achieved by changing the bias field strength to move the resonance position, or by using a pulse time outside of this range. Alternatively, additional sub-level state preparation may be used to remove those atoms in the first instance.

The population transferred by a Raman pulse at each interferometer pulse correlates to the final fidelity of the interferometer fringe, or, in the gradiometer case, the ellipse. Generally speaking, the final ellipse fidelity is a product of the population transferred by the three pulses. It has been observed that low transfer in the Rabi oscillations, particularly in the last pulse, limit the final fidelity of the ellipse, although a deterministic scaling has not been found. In order to maximise the transfer in both ellipse axes, the Rabi frequency has been optimised to be as similar as possible for each

interferometer. This ensures that the single π time set in the experiment sequence addresses both interferometers equally.

5.5 Sub-level state selection

As seen in the spectrum scan data, shown in section 5.3, the population is divided amongst the three m_F sub-levels. Whilst roughly 1/3 of the atoms are resonant with the Raman transition in the interferometer, the remaining two thirds of the atoms are not, and still appear in the detection signal. It was noted that some of these atoms in the co-propagating peaks were responsible for additional transfer in the Rabi oscillations when the states overlap (figure 5.21). As discussed later in section 6.2.1, photons in detection not from the atoms partaking in the interferometer contribute to the overall detection noise. To combat this, the atoms in the $m_F = \pm 1$ states may be removed via sub-level state selection. In addition to removing background signal, the sub-level state selection removes any atoms that have a magnetic sensitivity. The consequence of this is the bias field is not required to separate the m_F states, only providing a bias field strong enough to define a quantisation axis. This reduction in field strength will reduce additional phase shifts in the interferometer, as discussed in section 2.3.

Here, the state selection will be performed optically, with the steps required outlined below. The sub-level state selection can also be performed using the RF antenna installed into the vacuum system, requiring a modification to the RF chain. At the end of the optical molasses sequence, the atoms are prepared into the $|F = 1\rangle$ ground state. As before, a resonant blow-away is used to remove any atoms left in $|F = 2\rangle$, ensuring a clean hyperfine ground state. To prepare the magnetic sub-level state, the bias field is switched on, and a Raman transition transfers atoms from

$|F = 1, m_F = 0\rangle \leftrightarrow |F = 2, m_F = 0\rangle$ with a π pulse. Light resonant with atoms in the state $|F = 1\rangle$ is required to blow-away the remaining atoms left in the sub-level states $|F = 1, m_F = \pm 1\rangle$. A Raman transition can then be used to transfer atoms back to the state $|F = 1, m_F = 0\rangle$, and an $|F = 2\rangle$ blow-away can be used to remove any atoms not transferred back. This sequence allows for the magnetic sub-level to be selected through removing the other states.

It is possible to further improve the number of atoms in an $m_F = 0$ state by instead using optical pumping to prepare the state $|F = 2, m_F = 0\rangle$. This technique requires an additional π polarised light, resonant with the transition $|F = 2\rangle \leftrightarrow |F' = 2\rangle$, on the ^{87}Rb D_1 line [47, 108].

5.5.1 Creating an F=1 blowaway

In order to efficiently perform the state selection, atoms need to be removed when in the state $|F = 1\rangle$. This is achieved with the sideband added to the Raman laser, as described in section 4.1.2, where a frequency resonant with atoms in the $|F = 1\rangle$ ground state can be independently generated without the cooling frequency.

Careful tuning of the RF frequency found that the most efficient blow-away achieved at an applied RF frequency of 8383 MHz. This corresponds to exciting the transition $|F = 1\rangle \leftrightarrow |F' = 1\rangle$, although the closely grouped hyperfine excited states cause scattering into the $|F = 2\rangle$ ground state. In order to reduce the number of atoms transferred into the $|F = 2\rangle$ state, the optical power must be tuned in addition to the frequency. To characterise the $|F = 1\rangle$ blow-away, atoms are prepared into $|F = 1\rangle$ and the length of the resonant pulse is scanned. The results can be seen in figure 5.22, where efficient blow-away is achieved for the bottom atom cloud. The top atom

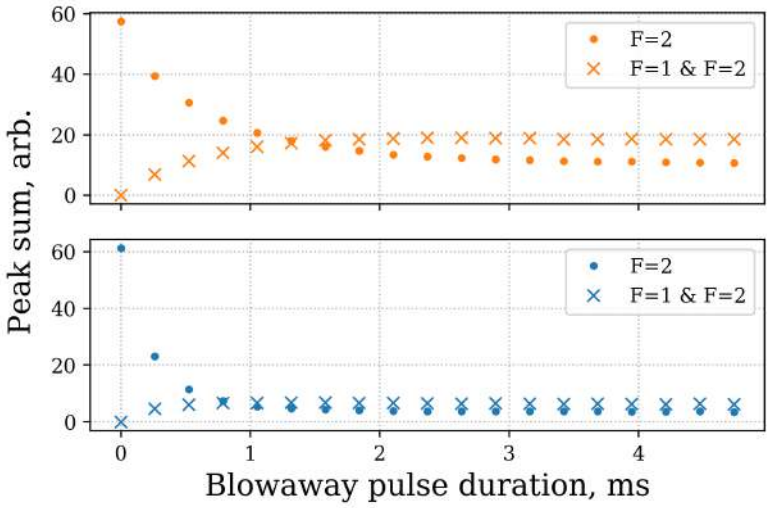

Figure 5.22 Characterisation of the $|F = 1\rangle$ blow-away. Despite being on resonance, and with low power, there are still atoms transferred from $|F = 1\rangle$ to $|F = 2\rangle$, rather than being blown away. An additional blow-away resonant with atoms in $|F = 2\rangle$ will be required to achieve a clean state.

cloud shows pumping of the atoms into the $|F = 2\rangle$ state, which reduces the efficiency of the blow-away pulse. An additional blow-away resonant with $|F = 2\rangle$ will remove the atoms transferred by $|F = 1\rangle$ blow-away.

5.5.2 State characterisation and velocity selectivity

To characterise the state, the atoms are prepared, using the sequence described in the previous section, into the $|F = 1, m_F = 0\rangle$ ground state. Performing a spectrum scan, both with and without the state preparation, is seen in figure 5.23, where, it shows that the population of atoms in the $m_F = \pm 1$ states have been removed. Note that since the detection method takes the ratio of population in the state $|F = 2\rangle$ and the total population, that the ratio now may go to a maximum possible 1. The state selected data is fit with the sinc^2 function of equation (2.29), as discussed in section 2.2.6.

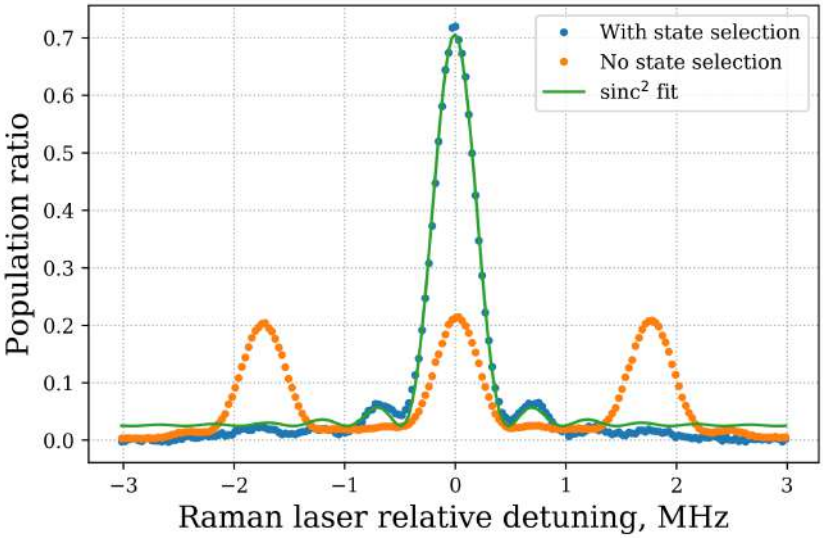

Figure 5.23 Atomic spectrum, with and without the optical m_f state preparation, shown for the top chamber only. The state selected data is fit with a sinc^2 function due to the velocity selective pulse having a square intensity profile.

The sinc^2 function does not fully describe the data, as seen in the higher order lobes. To do so, a more thorough treatment of the velocity selection is required, where the output is the product of the atomic cloud's momentum distribution, and the sinc^2 function from the square Raman pulse [63, 109]. As well as selecting the atomic state, the Raman pulses used have also selected a particular velocity class determined by the longest velocity selective Raman pulse.

With the state preparation verified, Rabi oscillations can be taken using the clean state. Figure 5.24 compares two Rabi oscillations, with and without sub-level state preparation. As well as increased transfer and coherence, the effect of spontaneous emission, from scattering with the other m_F states, is reduced, as evidenced by the upward slope of the data without state selection.

With increased fidelity of each Rabi oscillation, the amplitude of the el-

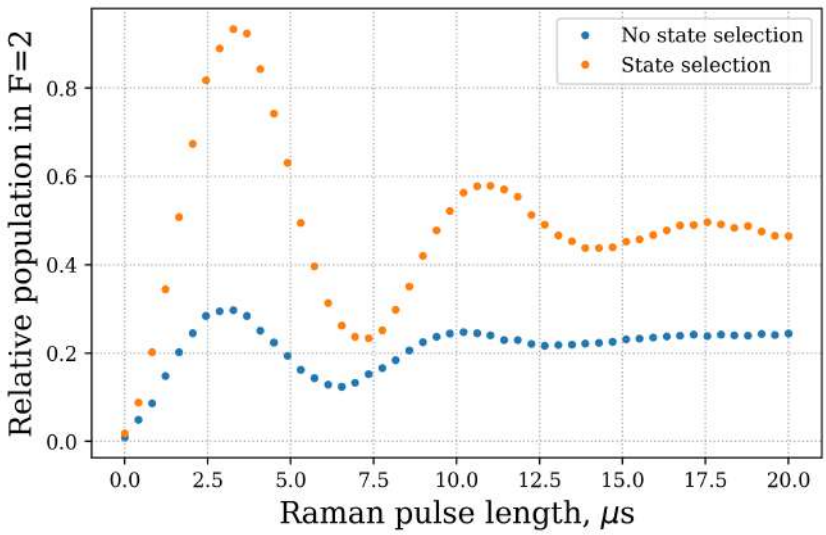

Figure 5.24 Rabi oscillation at the height of the first Raman pulse of the interferometer. State selection removes the atoms previously present, effectively re-scaling the oscillation from 1/3 to 1. The effect of spontaneous emission is reduced as evidenced by the upward slope of the un-selected data.

lipse is expected to increase. Figure 5.25 shows two comparable ellipses, with and without sub-level state selection. The ellipse contrast is increased as expected for the state selected atoms, but analysis of the ellipse phase noise showed little improvement to the signal to noise. This is because although the pedestal created by the atoms not participating in the interferometer is gone, the state selection process loses atoms in its inefficient blow-away pulses and Raman transitions. Despite no improvement to the ellipse at this time, the state selection allows for the removal of the other states, which previously provided mechanisms for the coupling of magnetically sensitive states into the interferometer. Using sub-level state selection is therefore required for use in a field environment.

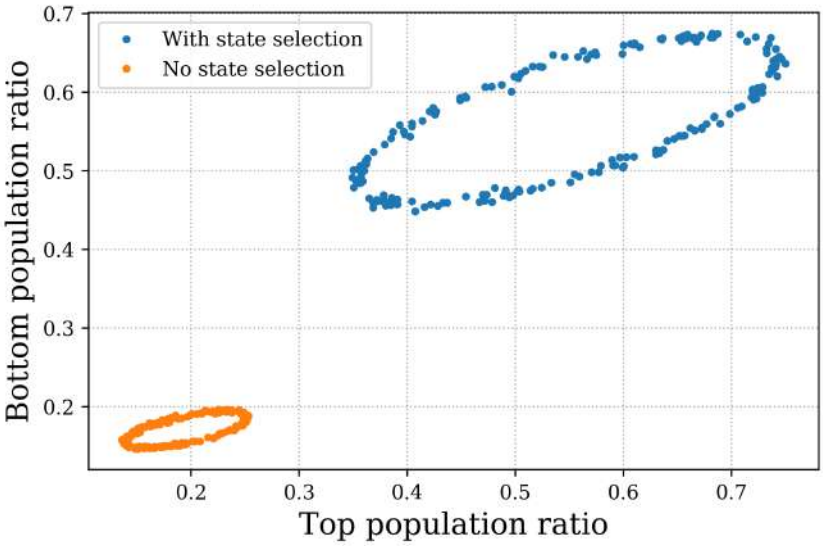

Figure 5.25 Comparison of ellipses with and without state selection. The ellipse contrast increases when sub-level state selection is used.

5.6 Ellipse phase shift

In an ideal cold atom gravity gradiometer, the phase shift between the two interferometers is accumulated due to the local gravity gradient. In practice, systematic effects such as those due to light shifts, magnetic field gradients, and instrument orientation, introduce additional phase shifts between the two interferometers. For a gravity gradiometer targeting civil engineering applications, systematic phase shifts which have a time dependence or change with external factors should be removed to ensure repeatable measurements. Static systematic phase shifts do not present an issue for a relative gravity gradient measurement, where the absolute phase shift is not required.

5.6.1 Applied relative phase shift

Analysis of the ellipse fitting, discussed in section 4.3, shows that a bias to the ellipse phase is introduced to the measured phase value when the ellipse phase is anything other than $\pi/2$. It is particularly strongest at phases close to multiples of π, where the ellipse forms a line. At these phase values, the points overlap and the fit to the data becomes unreliable. When the ellipse measurement is run at the maximum interferometer T time, it was found that the ellipse phase was a multiple of 2π, creating a "closed" ellipse. To "open" the ellipse, and facilitate fitting, an applied phase shift is used to introduce a relative phase shift between the two interferometers, recovering an open ellipse.

Introducing an applied phase shift to the gradiometer must be performed in a way that does not compromise the overall stability of the phase measurement. It should not therefore introduce additional phase noise, or drift, into the measurement, lest the overall sensitivity be affected. The implementation of an applied phase shift should have an element of tunability to allow for any changes in the phase, such as the optimisation of the light shift [11].

The methods to introduce an applied phase shift can be split into passive and active approaches. A passive solution requires a constant change to be applied to the system and could be achieved by the following: a magnetic field gradient added to the bias field profile, particularly in one interferometry region; alteration of the alignment to the Raman beams; changing the chirp rate applied to the Raman sideband, in order to change each interferometer's phase; or any further method which break the symmetry of the two interferometers.

A more deterministic phase shift may be added to the system using an

adapted method to one found in the literature [110]. The technique was developed to enable the extraction of the absolute gravity gradient between the two interferometers, by introducing a phase shift proportional to an applied frequency shift, $\Delta\nu$. The phase shift occurs through changing the value of $|\mathbf{k}_{\text{eff}}|$ by an amount $|\mathbf{k}_{\text{eff}}| = 4\pi\Delta\nu/c$, during the central π pulse only. Here, there is no requirement to know the absolute gradient value, but the technique may be applied to introduce a deterministic phase shift to the ellipse.

To implement this shift in the current laser architecture, a shift to the cooling laser frequency causes the offset locked Raman laser frequency to move by the applied frequency shift. Since the RF system can already ramp the cooling frequency without issue, the frequency of the central π pulse is easily changed in the time between the first two pulses, and ramped back in time for the final $\pi/2$ pulse. An example of the applied phase shift is seen in figure 5.26, where a frequency shift of 10 MHz is applied to the Raman laser frequency for the central π pulse. The applied frequency shift changes the ellipse phase by approximately $\pi/4$.

To characterise the resultant phase shift from the applied frequency shift, several ellipses were measured with different frequency shifts. The ellipses are fit to every 25 shots, and the mean phase of each 450 shot ellipse is taken. Each phase value is repeated twice, separated by approximately 2 hours. Figure 5.27 shows the expected linear dependence of ellipse phase to the applied frequency shift, measured to be $(0.076 \pm 0.001)\,\text{rad}\,\text{MHz}^{-1}$.

5.6.2 Ellipse parameters vs. interferometer T time

The sensitivity of the gravity gradient measurement is determined by both the signal to noise ratio from the atoms and additional noise sources. In an

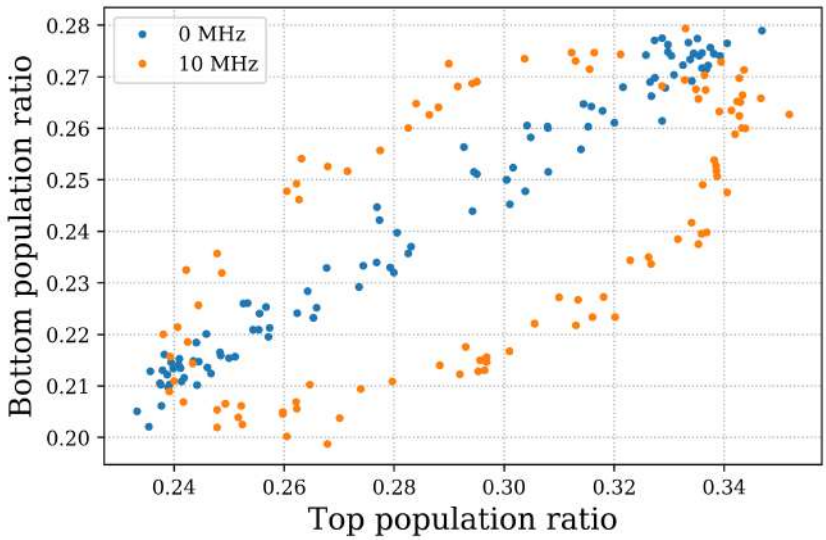

Figure 5.26 First 100 points of ellipses at T =85 ms, with and without an applied phase shift. A frequency shift of 10 MHz is sufficient to open the ellipse for fitting.

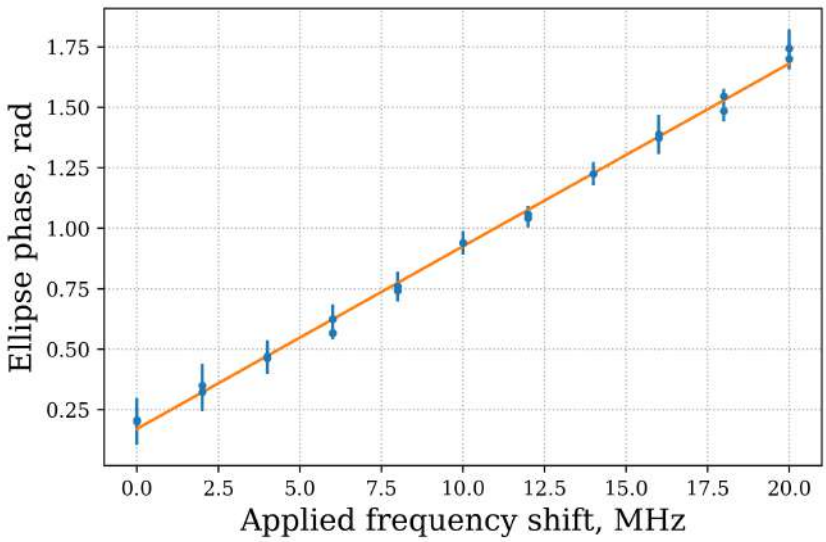

Figure 5.27 A frequency shift is applied to the single photon detuning, Δ, during the central interferometer π pulse. The resulting phase shift is characterised as $(0.076 \pm 0.001)\,\mathrm{rad\,MHz^{-1}}$ for a $T = 85$ ms ellipse. The error of the phase from fitting is lowest when the ellipse phase is close to 1 rad.

ideal gravity gradient measurement, a larger interferometer T time always results in an increase in instrument sensitivity; however, due to dephasing this may not be the case in the experiment. The ellipse parameters are characterised for different interferometer T times, to determine the phase evolution, contrast, and phase noise. Each ellipse measurement consists of 2350 shots, and 8 different T times, with a fixed first pulse time of 17 ms after the molasses sequence. Each T time is repeated between 8 and 9 times over a continuous three day measurement period. Again, each ellipse is split into 25 shot sub-ellipses and fit to, discussed in section 4.3. The run was performed with a 10 MHz applied frequency shift to ensure an open ellipse at 85 ms, as per the results of section 5.6.1, and operate without sub-level state preparation. The consequence of this is that ellipses at T times of around 40 ms are omitted, since their phase is close to a multiple of π and were not reliably fit to.

Ellipse contrast

The ellipse contrast can be measured by taking the maximum and minimum population ratio for each fringe axis, and is plotted against each interferometer T time in figure 5.28. The contrast is seen to decay, in both axes, by a rate of $(-0.0013 \pm 0.0003)\,\mathrm{ms}^{-1}$, due to the dephasing seen in the Rabi oscillations in figure 5.21. In order to reduce this drop in contrast, the fidelity of the Rabi oscillations should be improved.

Measurements of the amplitude noise can be estimated from the standard deviation of the ellipse contrast data, shown in figure 5.29, where the noise is shown to be approximately constant for the interferometer T times measured. The amplitude noise does not increase with T; however, the ellipse amplitude decreases, contributing to an overall lower signal to noise

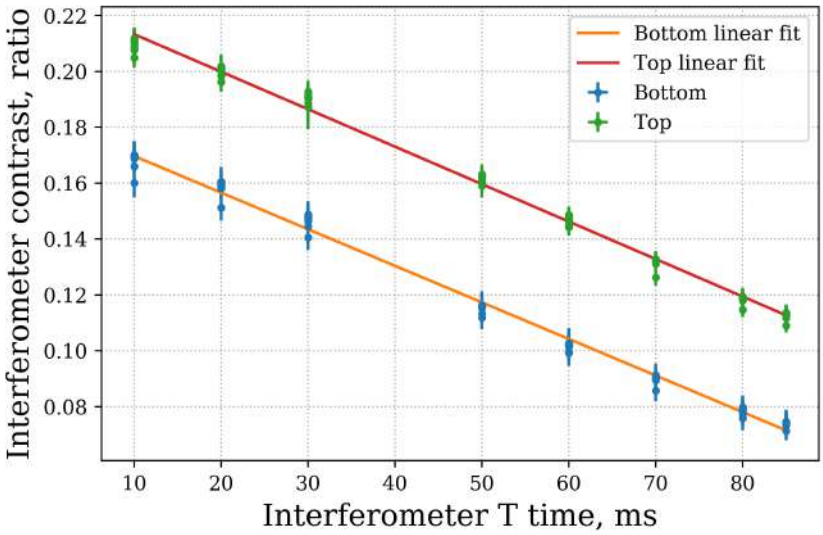

Figure 5.28 Ellipse contrast vs interferometer T time. Both show a linear trend with a gradient of $(-0.0013 \pm 0.0003)\,\mathrm{ms}^{-1}$.

ratio.

Ellipse phase

The phase evolution against T time is shown in figure 5.30. Each of the eight or nine repeats are shown on each point with their error, demonstrating the repeatability of each phase measurement. The scaling is approximately linear with a gradient of $(41.4 \pm 0.6)\,\mathrm{mrad\,ms}^{-1}$, where deviations from the linear suggest additional systematic phase contributions. The phase change expected from gravity gradient alone may be calculated using equation (2.44), where, for the gravity gradient of Earth, the phase shift is quadratic in T, with a change of $49\,\mu\mathrm{rad\,ms}^{-2}$. At 85 ms, the change in phase expected from gravity gradient is 0.359 rad, ten times smaller than the change in phase seen. The instrument is therefore dominated by systematic shifts; further testing will determine if they are independent of external factors.

To determine the evolution of the ellipse phase noise with interferometer

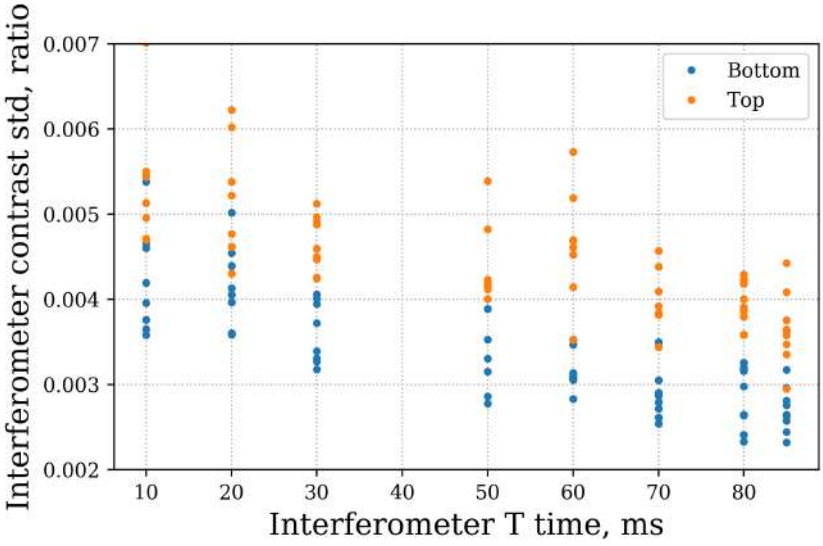

Figure 5.29 Standard deviation of the interferometer contrast vs T time. There is a negligible change in the spread over the different T times, suggesting that amplitude noise in the interferometer does not scale with the T time.

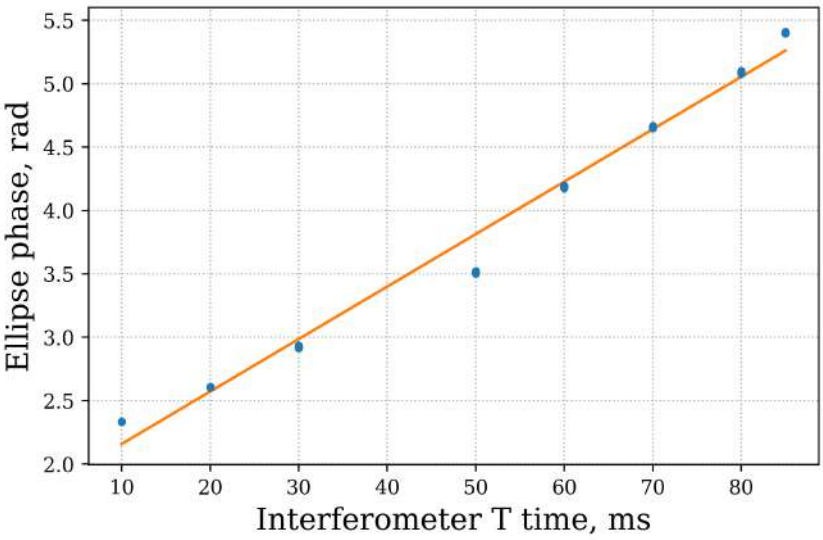

Figure 5.30 Ellipse phase evolution vs interferometer T time for 8-9 repeated measurements, shown with error bars. The approximate relationship between the two is obtained by a linear fit and is calculated as $(0.0414 \pm 0.0006)\,\mathrm{rad\,ms^{-1}}$. A 10 MHz frequency shift has been applied to all ellipses.

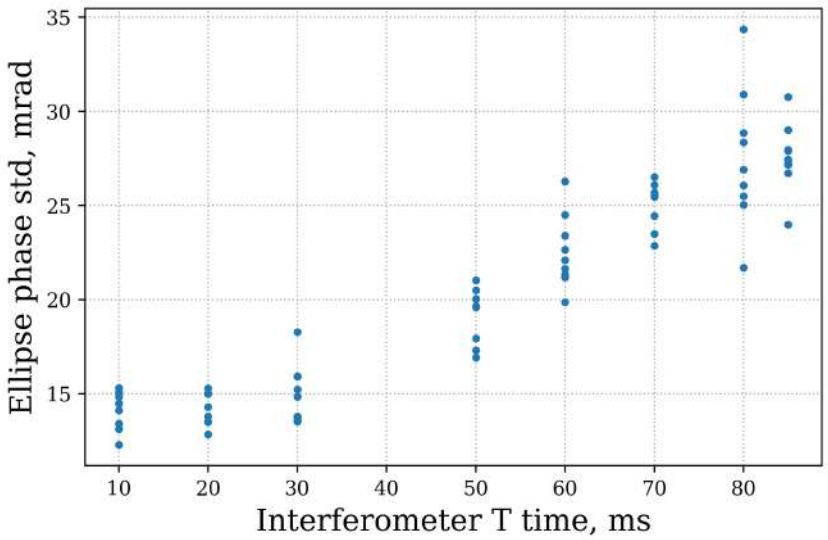

Figure 5.31 Standard deviation of the phase vs T time. Increase in uncertainty of the phase by a factor of 2 over measured T times with approximate scaling $0.2\,\mathrm{mrad\,ms^{-1}}$.

T time, the standard deviation of the ellipse phase is plotted in figure 5.31. The certainty in the phase does degrade with T time, increasing by a factor of 2 over the T times tested, which may be caused by the decreased ellipse contrast. Since the gradient measurement is contained in the ellipse phase, a larger T time produces a phase change proportional to T^2. This means that operation at an increased T time is more favourable at this level of phase uncertainty. To aid the understanding of this, for a fixed signal size of 500 E, the change in ellipse phase can be calculated from equation (2.44) for different T times. In the case of $T = 10\,\mathrm{ms}$, the ellipse phase shift would be less than 1 mrad, however for $T = 85\,\mathrm{ms}$, the ellipse phase shift would be approximately 58 mrad. A factor of two increase in the ellipse phase uncertainty is small compared with the factor of 70 increase in the signal size in this case.

5.7 Characterisation summary

The cold atom interferometry sequence has been broken down into its constituent parts and characterised. The cold atom system is shown to produce ensemble temperatures of $(2.86 \pm 0.09)\,\mu\text{K}$ and $(3.7 \pm 0.2)\,\mu\text{K}$ for the top and bottom interferometers respectively. The atomic state has been cleanly prepared in the $|F = 1\rangle$ hyperfine ground state, with the option to prepare further into the $|F = 1, m_F = 0\rangle$ state. Rabi π times of as low as $2\,\mu\text{s}$ have been demonstrated, enabling large velocity classes to be addressed. Two state detection allows for the output of the interferometer to be analysed, using a vertical detection scheme. Interferometer T times of up to $85\,\text{ms}$ have been demonstrated and will be used to measure the performance of the gravity gradiometer.

Chapter 6

System performance and measurement campaigns

This chapter examines the performance of the instrument, in and out of the laboratory, and uses the modulation of test masses to benchmark the instrument and test functionality. When discussing the ellipse phase, the values will be converted into units of gravity gradient using equation (2.44), to facilitate comparison between instrument performance and gravitational signal sizes more readily. The terms ellipse phase and gravity gradient are used interchangeably since differences in signal are being investigated.

6.1 Instrument sensitivity

The sensitivity may be determined through measurements of the ellipse phase, with the instrument operating in static conditions. From this, the short term noise can be determined as well as any drift in the ellipse phase. In order to assess the performance of the instrument, an Allan deviation is used to determine the average measured phase value after a given averaging

time. Typically, equally spaced, time series data is required. To measure the instrument noise and stability, the phase measurements should be taken with a static instrument, with static environmental conditions. In the treatment of this sensor, overlapping estimates are used for higher confidence in the Allan deviation, as recommended by IEEE standards, and follows the descriptions within them [111–113]. The formula used to estimate the Allan deviation and the error is derived in appendix B.2.

6.1.1 Short term vs long term sensitivity

Plotting the Allan deviation for all available values of the averaging time, reveals features of the data and can help to identify the types of noise present in the system, provided the physics is well understood. Provided the instrument drift occurs over long timescales, the sensitivity of the instrument is measured by simply taking the first point from the Allan deviation and dividing by the square root of the integration time.* This is the average short term sensitivity of the instrument and describes the minimum measurement uncertainty after a given integration time. For Poisson distributed noise, the measurement uncertainty will average with $1/\sqrt{\tau}$ behaviour, where τ represents the integration time.

As noise sources and drift accumulate in the measurement set, the deviation from Poisson distributed noise reveals instability in the system. The uncertainty in the phase after a certain integration time may be used to estimate a long term sensitivity of the instrument. After a certain integration time, accumulated noise and drift mean that averaging the phase values no longer decreases the uncertainty in the measurement. Where possible, the resolution of the sensor will be estimated from the point of lowest phase

*Fitting to the data and calculating the intercept at τ =1 s may provide a more accurate determination of the sensitivity.

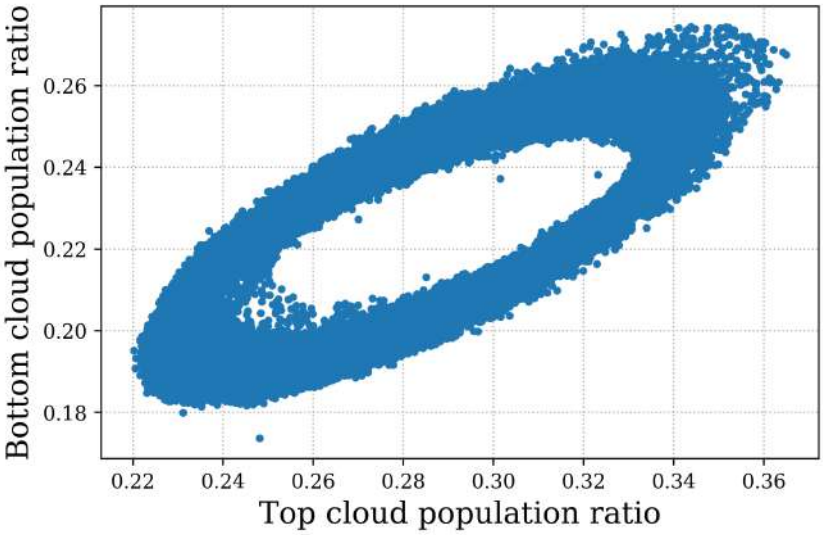

Figure 6.1 66 hour gravity gradient measurement with all points shown. Every 25 shots would be fit with an ellipse and the phase value extracted. Smearing of the ellipse is seen due to drift in the ellipse contrast over timescales longer than a 25 shot measurement. Ellipse is opened using an applied phase shift described in section 5.6.1

uncertainty.

6.1.2 Sensor evaluation

The system was ran continuously for several days in a static condition and temperature controlled lab, in order to gauge the instrument performance. The resulting ellipse is shown in figure 6.1 for all 155300 shots. The ellipse measurement is split into sequential 25 shots, and a phase value for each is extracted, using the fitting method described in section 4.3. The phase versus time, shown in units of gravity gradient (equation (2.44)), is displayed in figure 6.2.

An Allan deviation of the data, shown in figure 6.3, measures the short term sensitivity of the instrument at $(861 \pm 7)\,\mathrm{E}/\sqrt{\mathrm{Hz}}$, averaging to a minimum value of 16 E in 23 757 s. It is clear from the Allan deviation that the

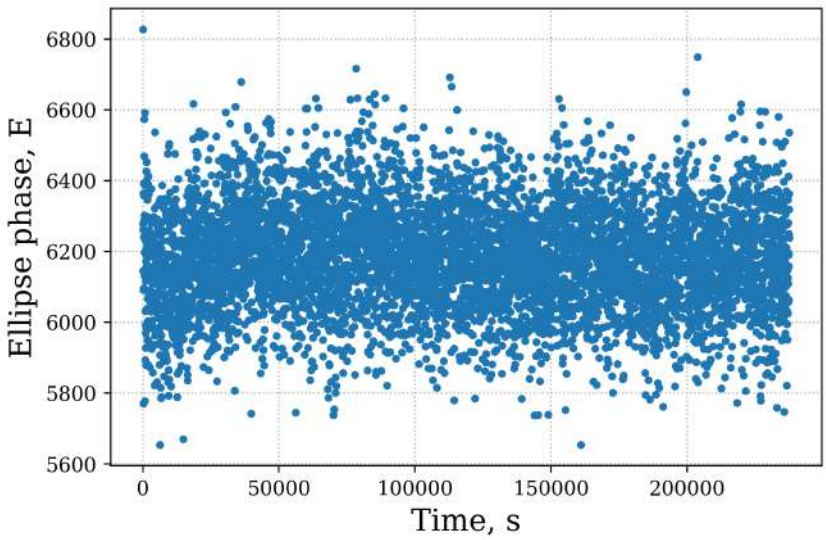

Figure 6.2 Ellipse phase in units of gravity gradient vs time for a $T = 85\,\text{ms}$ ellipse run of 66 hours in a temperature stabilised lab.

averaging quickly diverges from $1/\sqrt{\tau}$. This is indicative that there is drift in the instrument, which prevents white noise averaging at longer integration times.

The main noise sources will be investigated to determine the current sensitivity limits of the instrument.

6.2 Noise sources

In this section, the main contributions to the noise budget will be investigated. As discussed in section 2.4.5, the quantum projection noise is the fundamental sensitivity limit of this atom interferometer, where the gravity gradient uncertainty is limited by the statistical uncertainty from the number of atoms. It is estimated in appendix D.3 that 10^5 atoms partake in each interferometer, at a time between measurements of 1.5 s; the repetition rate of the interferometer is 0.67 Hz. Equation (2.47) is then used to

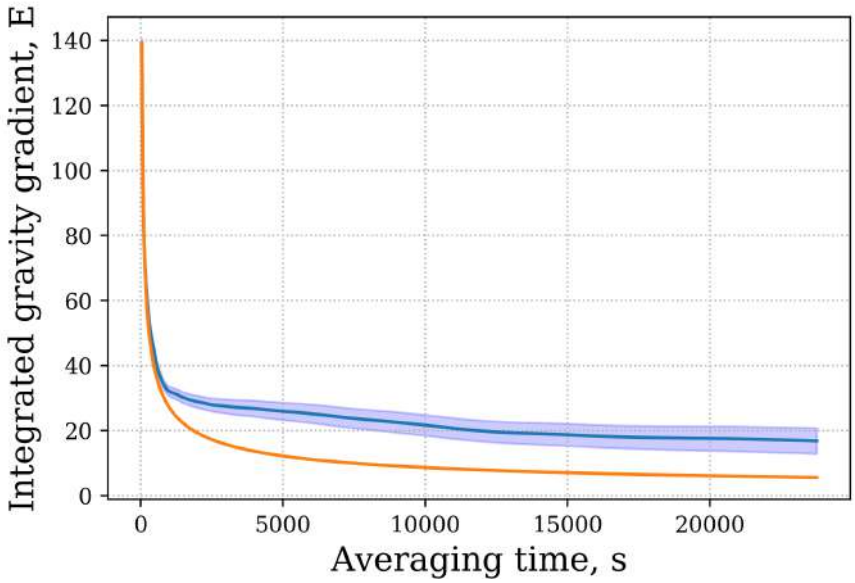

(a)

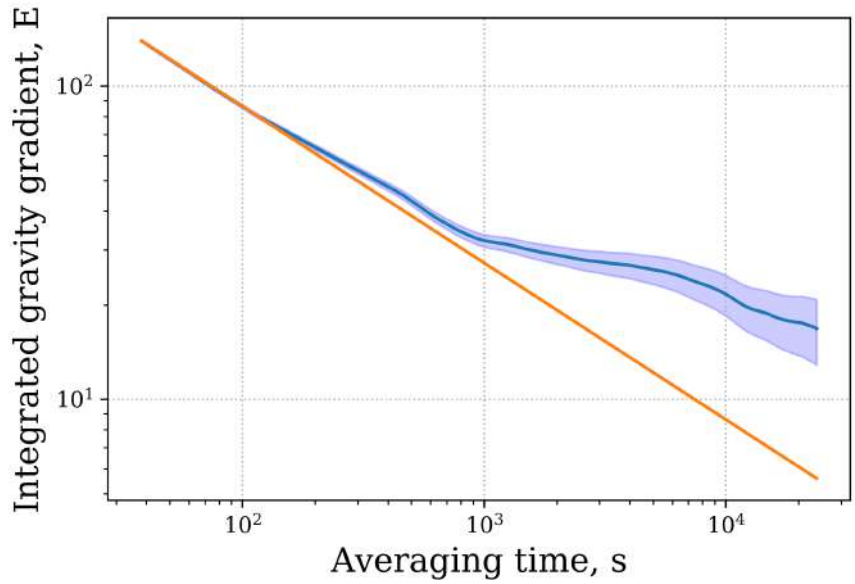

(b)

Figure 6.3 Allan deviation for 66 hour data set shown in a) linear and b) log-log plot. The shaded section shows the percentage error calculated with equation B.13. The orange line denotes the $1/\sqrt{\tau}$ average from the first point.

calculate the atom shot noise limited, short term sensitivity as $47\,\mathrm{E}/\sqrt{\mathrm{Hz}}$, for an interferometer T time of 85 ms. Comparing this to the short term sensitivity measured in figure 6.3, the instrument is operating at 18 times the atom shot noise limit.

6.2.1 Detection scheme evaluation

The detection scheme is that described in section 5.2, where fluorescence from resonant pulses is detected by an amplified photodiode. The ratio of the state populations is calculated using equation (5.8), from integrating the voltage for each detection peak, as measured by the detection photodiode assembly.

The measured voltage from each detection peak is a result of the amplification of several current sources, which can be broken down into their contributions. The main current sources are: the photoelectrons due to fluorescence from the atoms I_a, the photoelectrons due to background light I_b, and current noise in the photodiode I_{ph}. The total measured current may therefore be written as a sum of these sources, such that

$$I = I_a + I_b + I_{ph}. \tag{6.1}$$

Each current source I_x, can be written as the sum of an averaged value $\overline{I_x}$, and a current noise δ_x, allowing the separation of offsets and noise. The ratio taken with the detection scheme can remove offsets if they are applied to the detection and background peaks equally, which, for example, is used to remove intensity fluctuations, to first order. Ideally, the only current contribution should come from the current associated with the atom fluorescence, I_a, so to achieve lower noise overall, contributions from other sources should be reduced.

Signal current sources

The photocurrent from the atoms is given by the responsivity of the photodetector, R, multiplied by the optical power, P. The optical power is given by the photon flux, Φ, multiplied by the energy of the photon, $\hbar\omega$. For an atom emitting radiation, the photon flux is given by the detection solid angle, Ω, the scattering rate from the atoms, Γ_{scatt}, and the number of scattering atoms N. Thus the photocurrent can be written as

$$I_a = RP = R\Omega\Gamma_{scatt}\hbar\omega N. \tag{6.2}$$

The maximum photocurrent occurs when the scatting rate is maximum, equal to half the transition linewidth, $\Gamma_{\mathrm{scatt}} = \Gamma/2 = 2\pi{\cdot}6.06\,\mathrm{MHz}/2$. Silicon detectors for 780 nm have a typical responsivity of $0.45\,\mathrm{A\,W^{-1}}$. Thus, for 10^5 atoms in a detection solid angle of 0.21 sr, the value of photocurrent expected from the atoms alone is on the order of 10 nA, meaning that a large amplification is needed to generate voltages with a significant signal to noise. With the photodiode gain 10^7, a voltage of 100 mV is generated from the atomic signal.

Detector noise

Fundamentally, the noise of the detector is limited by Johnson noise, where fluctuations due to thermal motion of charge carriers cause an uncertainty in the measured voltage or current. The average value of the current caused by Johnson noise is zero, so the current contribution becomes $I_{PH} = \delta_{PH} = \overline{i_N^2}(\nu)$. Johnson current noise depends on the temperature of the circuit $\mathcal{T}$, the resistance R, and the bandwidth of the sensor $\Delta\nu$ [114], and is given by

$$\overline{i_N^2}(\nu) = \frac{4k_B\mathcal{T}\Delta\nu}{R}. \tag{6.3}$$

Since the detection sequence used requires short pulses, the bandwidth becomes an important contribution to the Johnson noise. Using the short pulse detection scheme therefore puts increased technical requirements on the detection amplification circuit, to avoid adding additional noise. For the photodiode used in detection (section 3.7.1), the Johnson noise is quoted as 50 pA, which when the photodiode amplifier gain is applied, without filtering, gives a noise of 0.5 mV.

To ascertain the contribution of detector noise to the overall noise, the voltage noise when the light is off can be examined. A small sample of this is available between detection pulses 2 and 3, and can be seen in figure 5.10. Taking the standard deviation of the voltage between these peaks gives 0.74 mV and 0.85 mV, for the top and bottom signals respectively. This may be compared to the Johnson noise by taking the difference in quadrature between the two. This estimates that 0.54 mV and 0.7 mV of each signals detection noise comes from the photodiode amplification and electrical pickup noise. The data acquisition device is specified to 61 µVrms noise at range ±1 V, contributing only a small fraction of the noise [115].

Background light

The majority of the detection signal is from background light levels, as seen in figure 5.10. Optical access for ambient light has been removed, and variations from ambient light levels are not observed. Contributions from background light are therefore coming from within the chamber, due to scattering from the vacuum chamber walls and prism MOT assembly, as well as fluorescence from background rubidium vapour. Scattered light contri-

butions may be reduced through the use of an absorptive coating applied to the science chambers. Testing of a blackened coating at ultra-high vacuum was still in the early stages during design and construction of the vacuum system, so the coating was forgone. The results of the blackened coating have since been published, and may be found in reference [116], where it was demonstrated that a coating applied to the vacuum chamber does not impact the vacuum quality, or react with the rubidium vapour, whilst providing a reduction in scatter. Background light levels may also be reduced through use of a mask over the photodiode, or imaging detection to digitally mask the signal. Delivery of the detection light through the side port may also reduce the contributions from background light.

An additional source of photocurrent comes from atoms not partaking in the interferometer sequence. This can arise due to choice of state preparation and dephasing, where up to 2/3 of the detected atoms do not participate in the interferometer. Each additional atom has an associated atom and photon shot noise which contributes to the overall detection noise.

The background light level noise can be estimated when the light is on, but no atoms are present, from the flat tail of peak 3. Here the standard deviation of the signal is 0.90 mV and 0.88 mV for the top and bottom signals, respectively. From these two measurements, the background light level noise is estimated as 0.5 mV when taking the difference in quadrature between the measured noise and the detector noise.

Sensitivity limit with identified noise sources

Using the measured noise amplitudes, Gaussian noise is applied to synthetic peak data in order to calculate the noise on the state population ratio, using equation (6.4) shown in the next section. The resulting population ratio

noise is applied as uncorrelated amplitude noise to two synthetic interferometer fringes, whose contrast is representative of the ellipse in figure 6.1. Fitting to the synthetic ellipse data, allows for the ellipse phase noise to be estimated; no phase drift is applied to the phase data. Taking an Allan deviation of the synthetic data, reveals that the detection noise limits the average short term sensitivity of the instrument to $(258 \pm 5)\,\mathrm{E}/\sqrt{\mathrm{Hz}}$. When the noise from the background light is included, using the noise profile from the end of peak 3, the sensitivity shows an increase to $(284 \pm 6)\,\mathrm{E}/\sqrt{\mathrm{Hz}}$.

Clearly this does not account for all the noise in the instrument; however, the detection noise is a large component of the total noise, limiting the performance to six times the current atom shot noise limit.

6.2.2 Data processing optimisation

With the detection noise characterised, the data processing can be optimised in order to further lower the noise. The noise from the detection is normally distributed, so the error can be reduced by increasing the number of measurements. A general treatment of errors from the detection sequence can be derived, and here is applied to the detection noise.

The state population ratio used in detection is calculated through the use of background subtracted peaks (equation (5.8)), with a noise associated with each peak. The voltage error from each detection peak may be written as σ_i, where i is the index assigned to each peak. The error on the ratio R, may be calculated from

$$\sigma_R = R\sqrt{\frac{\sigma_1^2 + \sigma_4^2}{(V_1 - V_4)^2} + \frac{\sigma_2^2 + \sigma_5^2}{(V_2 - V_5)^2}}, \tag{6.4}$$

where V_i is the mean peak value. Taking the limiting case where the detection noise is the dominant noise source, the noise on each peak is the

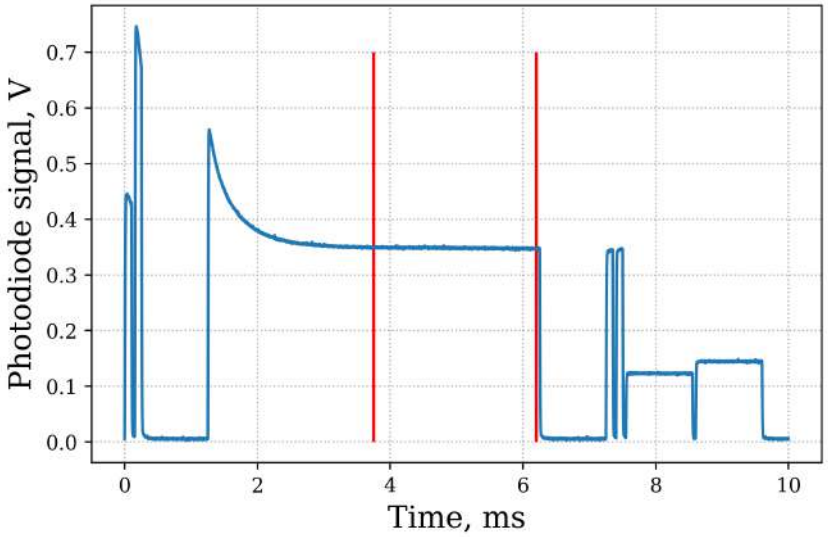

Figure 6.4 Using the tail end of the blow-away pulse P3, lower noise background subtraction is realised by integrating between the red markers. This provides a section of data 24.5 times longer than the background peak 4 or 5, and integrates the noise of the background peaks.

same, such that $\sigma \approx \sigma_i$. If the noise associated with any peak is reduced, then the error in the ratio is reduced, leading to lower ellipse phase noise. The noise in the atom peaks cannot be easily lowered, but the background subtraction may instead be performed with a longer integration time. To test this on existing data, the tail section of peak 3, 24.5 times longer than original peak length, is used as the new background peak. It is highlighted in figure 6.4. The increased length reduces the error of the background peak by a factor of five. In order to maintain the ratio, the new background peak's mean value is normalised to the mean values of peaks 4 and 5.

Using this new background subtraction on the data set presented in section 6.1.2, a reduction in short term sensitivity from $(973 \pm 8)\,\mathrm{E}/\sqrt{\mathrm{Hz}}$, to the presented value of $(861 \pm 7)\,\mathrm{E}/\sqrt{\mathrm{Hz}}$, is achieved. From this, each background peak length should be increased from 100 µs to 2.5 ms to reduce the instrument noise.

6.3 Systematics due to external effects

The instrument is assessed for systematic shifts and noise introduced by external factors.

6.3.1 Tilt

As discussed in section 1.2.2, a measurement of the gravity gradient has a suppression of the sensitivity to tilt compared to gravimetry. This suppression may be reduced from the ideal case due to implementation in the gradiometer, e.g. atom clouds following different trajectories.

In this section, the sensitivity to tilt is characterised by measurements of the change in ellipse phase, for a range of instrument tilts. The instrument tilt can be maintained to within 2.5 mdeg in the laboratory using the adjustable sensor head feet, as measured by an inclinometer.* The ellipse phase, in units of gravity gradient, is shown for a range of instrument tilts in figure 6.5. The mean gradient value of each measurement is shown, with an error bar of the standard error of the mean gradient. Fitting to the small angle measurements, the sensitivity to tilt is calculated as (-1.2 ± 0.8) E/mdeg from the gradient of the fit. For the example anomaly modelled in figure 2.12, the signal would be obscured by a tilt of 31 mdeg. Comparing this to the ideal gradiometer case shown in section 1.2.2, the tilt suppression is not as effective; however, compared to the ideal gravimeter, the same anomaly would be obscured by a tilt of 13 mdeg. The gradiometer has demonstrated a reduction in tilt sensitivity of three times when compared to an ideal gravimeter. An additional consequence of this is that drift in the instrument tilt provides less of a signal change, relaxing the constraints on stable

* *Jewell Instruments* DMP-2-10-232-AMP

flooring.

To explain the systematic shifts removing some of the tilt suppression, the atom interferometry scheme is examined further. The atom cloud's trajectory during the fall is relative to the Raman beam alignment. Due to the finite Raman beam size, and Gaussian intensity profile, systematic shifts are introduced to the interferometer when the the atom cloud samples different parts of the Raman beam. In the extreme case, an atom beginning in the middle of the Raman beam will travel to the 7.5 mm radius Raman beam for an angle of 2.9 degrees over the interferometer drop length, meaning that the interferometer sequence is not completed for that atom. In order for the tilt to be effectively suppressed at smaller tilts, the two atom clouds must follow the same trajectory and sample the same intensity profile of the Raman beam. This can be achieved with better cloud positioning and beam alignment, which all help to reducing dephasing in the interferometers.

6.3.2 Ambient temperature

Operation of the instrument in ambient temperatures in excess of 30 °C results in overheating of the control package, generally causing loss of laser lock and, in some cases, laser components exceeding their operating temperatures. Within operating temperatures, the instrument is tested for ellipse phase drift with ambient temperatures, by measuring in a temperature controlled laboratory for several days. The temperature is changed between temperatures of 20 and 34 °C, through use of the air conditioner, whilst the instrument continuously runs ellipse measurements.

The output phase, shown in units of gravity gradient, is plotted with the temperature measured at the base of the sensor package by the inclinometer, and can be seen in figure 6.6. There seems to be some dependence on the

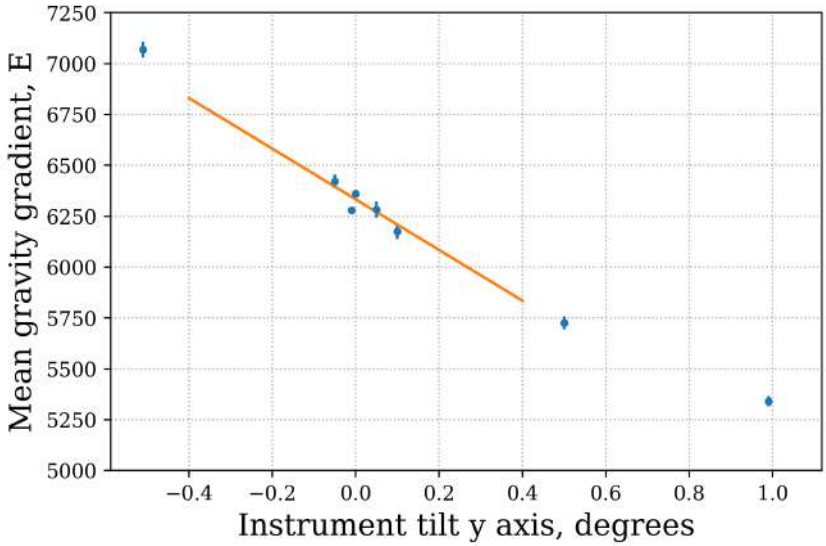

Figure 6.5 Instrument tilt parallel to the inclinometer y axis. Multiple phase measurements are averaged and the standard error shown for different instrument tilts. To allow for the averaging of the phase, the tilt is separated into box sizes of 10 mdeg and averaged. The resulting sensitivity for small angle changes is obtained from the fit over the central angles, with a calculated gradient of (-1.2 ± 0.8) E/mdeg. Expected cosine behaviour is not seen since the absolute gradient is not measured.

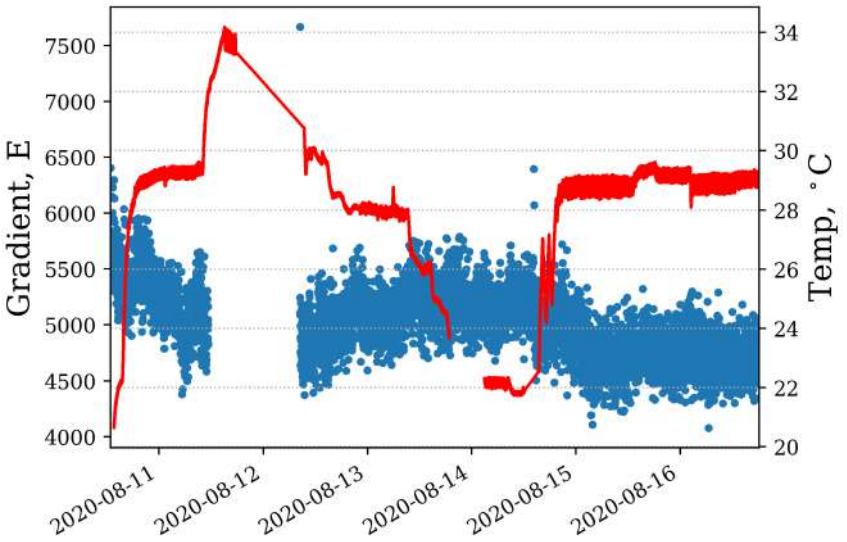

Figure 6.6 Ellipse phase (blue, left axis) vs ambient laboratory temperature (red, right axis), as measured by the inclinometer. There is certainly a change in the phase when temperature changes; however, no deterministic coupling is obvserved.

temperature; however, figure 6.7 shows that there is not a simple correlation. This is perhaps not surprising - changes in temperature have long timescales and cause spatial and temporal gradients. The effect of temperature on the entire instrument is complex, and dependent on several drift channels. The instrument should aim to remain at a lower operating temperature in order to have less measurement noise.

6.3.3 Coriolis measurement

As discussed in section 2.7, the Coriolis effect may contributes significant shifts to the instrument when rotated. When operating in spatial mapping, the instrument will need to be moved between measurement points, relying on the ability to repeatably align the rotation of the sensor each time.

To measure the effect of Coriolis on the instrument, the sensor package was placed on a large rotation stage, graduated at 1° increments. Ellipse

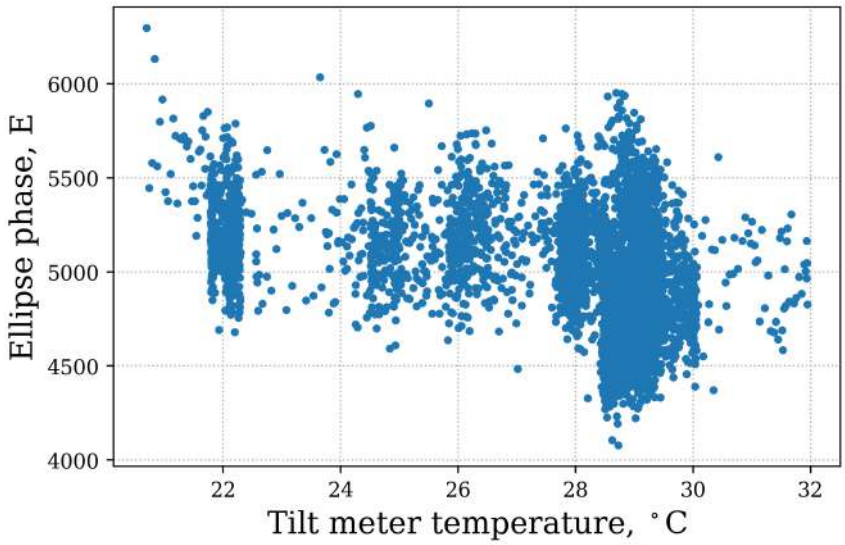

Figure 6.7 No correlation with ambient temperature and ellipse phase up to 28.3 °C, where the phase uncertainty increases. Above ambient temperatures of 30 °C, instrument operation is intermittent, due to overheating in the control package.

phase measurements were taken at 30° increments and the standard error of the mean is taken for each point. The data is fit using equation 2.52 and the value for the relative cloud velocity is measured from the fit parameters. This is shown in figure 6.8, where the measured value for v_{hor} is $(7.6 \pm 0.1)\,\mathrm{mm\,s^{-1}}$.

Whilst with careful optimisation the differential horizontal velocity can be reduced, the effect of orientation can be controlled with good survey practice. The laying of a survey line and ensuring good alignment to the line can allow for repeatability of a few degrees when done by eye. Further measurements with common civil engineering tools, such as laser scanners, can ensure repeatability to within fractions of a degree. At the maximum slope, the Coriolis effect introduces a change in gravity gradient of roughly 11 E/ deg to the instrument, when aligned to the East-West axis. This shift is therefore negligible when well controlled for, making it small compared to

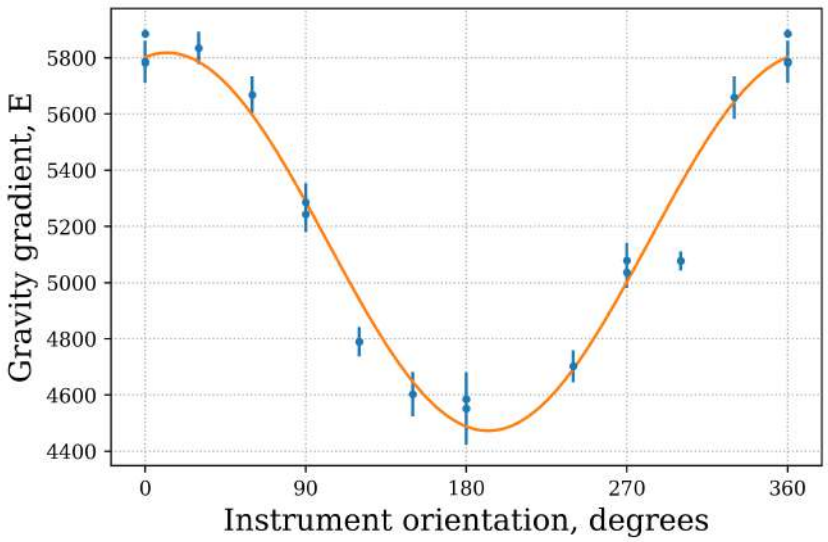

Figure 6.8 Measurements of the ellipse phase under rotation of the instrument with fit. Using equation 2.52, the relative cloud velocity is measured as $(7.6 \pm 0.1)\,\mathrm{mm\,s^{-1}}$. Points at 0 and 360 degrees are the same measurements, repeated for completeness.

prospective civil engineering targets.

6.4 Measurements of test masses

Benchmarking the instrument performance in the laboratory can be done using the modulation of test masses, whilst keeping the instrument stationary. Two trolleys are loaded with lead bricks and positioned around the centre of the sensor to generate the largest gravity gradient. Modulation is performed manually, with the trolleys pulled radially away from the sensor, to generate the change in gravity gradient.

The setup of this experiment is shown in figure 6.9, where the two trolleys are positioned on a pair of optical tables, with a gap for the sensor package. Aluminium bars are used to guide and position the trolleys, as well as prevent them falling from the tables. Each trolley is stacked with lead bricks

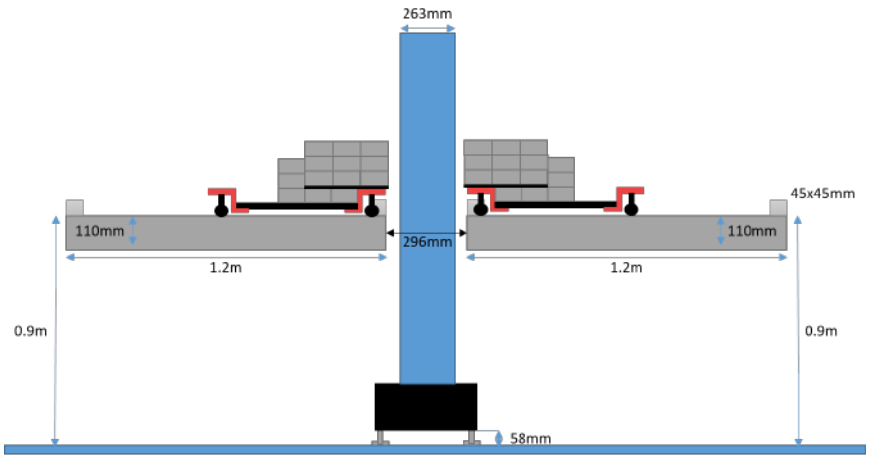

Figure 6.9 Benchmarking of the instrument using modulation of test masses. Two trolleys stacked with lead bricks, shown positioned at their closest to the centre of the instrument. Each trolley holds (487 ± 1) kg of lead and can be manually positioned in one of seven distances from the sensor package.

to a total mass of (487 ± 1) kg, in a wide, shallow, rectangular pattern to maximise the signal. The distance between the two extreme trolley positions is 550 mm, limited by the length of the optical tables. Modelling of the test masses in these two positions gives an expected change in gravity gradient of 198 E.*

Testing was performed over 5 days, with the mass modulated every 25 minutes by manually moving the trolleys into position each time. Modulation was performed between the two extremes on the first two days, an additional centre point on the third day, and two new points added on the final days, forming a total of seven distances from the sensor head. The trolleys being positioned close to the sensor are defined as the base station, with gravity gradient values defined relative to the base station. When these measurements took place, the performance of the instrument was approximately $1000\,\mathrm{E}/\sqrt{\mathrm{Hz}}$, degrading to $1300\,\mathrm{E}/\sqrt{\mathrm{Hz}}$ over the measurement days.

*With thanks to A. Rodgers and D. Boddice.

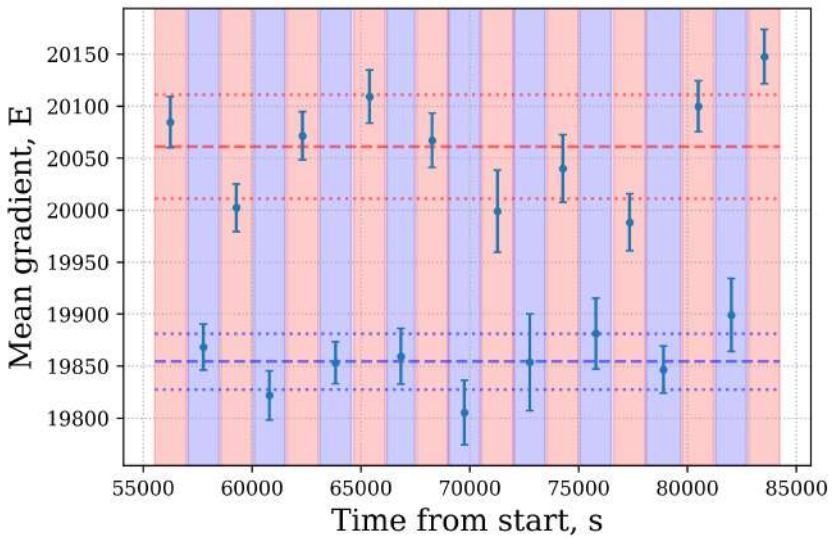

Figure 6.10 Test masses modulated between the two extreme points. Data is shown for a single day's modulation which excludes the overnight measurements for clarity in presentation. Red and blue bands show the time over which the measurement point was averaged, with the colour representing the mass being closest and furthest from the system, respectively. Dashed and dotted lines show the mean values of both positions with the standard deviation. There is some evidence of the phase drifting, some of which may be removed using a base station method [3]. Ellipse phase is put into gravity gradient using equation (2.44).

The results from the first day of modulation is shown in figure 6.10, where the test masses are moved from the base station to the furthest away point. Repeats are required since the phase appears to be drifting, likely due to changing lab temperature. Including the overnight measurements in the two positions, the difference between the two mass positions is calculated to be (202 ± 39) E.

Data from each subsequent day is analysed in the same way, calculating a mean value for the gravity gradient difference between each measurement point and the base station. The error ascribed to each value is calculated from the standard errors of each measurement point and the base station added in quadrature, meaning that the error from each base station measurements is transferred onto the other points. The results from all mea-

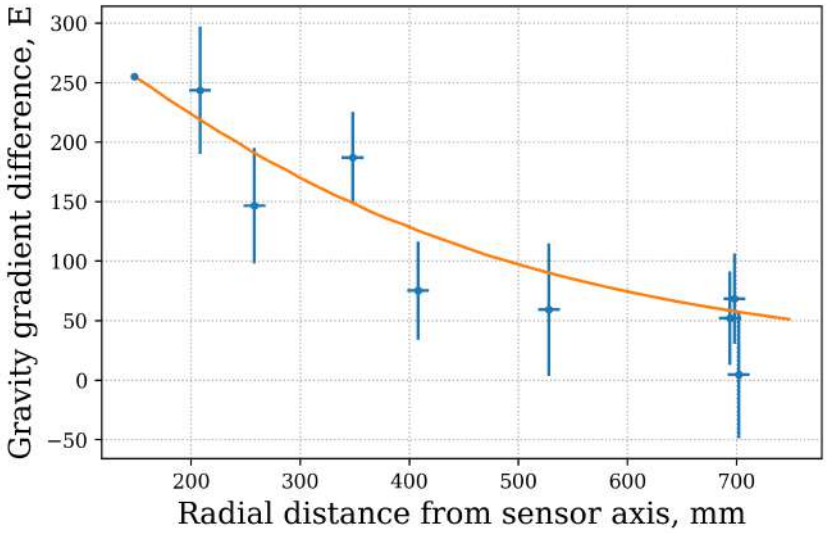

Figure 6.11 Average difference between the base station, defined as the masses closest to the system, and the measurement points for all data sets. Error on each measurement point is calculated from the standard error of the difference measurement, compared to the base station point. Overlaid is the model of the expected gravity gradient signal, aligned to the base station.

surements are shown in figure 6.11, where the results are shown relative to the defined base station point. The model of expected gradient from the test masses overlays the data and is aligned relative to the base station.

The noise and positions of the points can be further reduced by splitting the measurements into sets. Treating the data this way assumes a linear phase drift between each base station measurement, which can be subtracted from each of the measurement in the set [3]. This method of analysis is a little more involved and the phase drift was not observed to be linear between measurement sets; therefore, only the average difference of each days points will be presented here.

To ensure the effect is caused by the presence of the test mass and not the movement of the trolleys, a two point null test is performed with empty trolleys. Following the same measurement procedure as before, the

trolleys are moved every 25 minutes between the two extreme points. The difference between the closest and furthest measurement points gives an offset of (32 ± 41) E, showing that there is no adverse effect from the trolleys themselves.

The data shows a trend close to that expected from the model, with all data points within two sigma. The model also has its limitations, only considering an approximation of the lead shape along a single horizontal measurement line. From this benchmarking test, with the instrument operating at around $1000\,\mathrm{E}/\sqrt{\mathrm{Hz}}$, gravity gradient signals of 200 E were measured at a signal to noise ratio of around three, when all data sets were considered. Resolution of approximately 50 E was demonstrated, enabling the shape of the gravity gradient to be measured, in addition to the total signal size.

6.5 Outdoor measurements

6.5.1 Comparative performance

To compare the performance of the system outdoors and indoors, the system was left to run statically outside of the building on brick flooring. Seven hours of measurements were obtained outside and can be compared with the performance in a temperature controlled laboratory from figure 6.3.

The average short term performance of the instrument outside was measured as $(838 \pm 23)\,\mathrm{E}/\sqrt{\mathrm{Hz}}$, averaging to the lowest value of 34 E in 1224 s. This means that the short term performance is the same, to within errors, as the temperature stabilised laboratory results, up to averaging times of 2000 s. For averaging times over 2000 s, the outdoor performance begins to diverge from the laboratory data. Without a longer data set, it is not possible to determine the truly long term performance; however, the com-

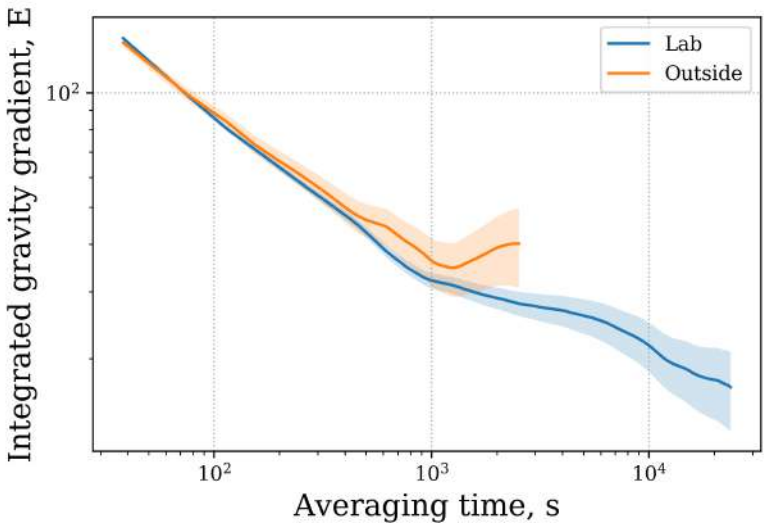

Figure 6.12 Comparison of the laboratory performance in a temperature controlled environment with the performance when measuring outdoors. The sensitivity is the same; however, some long term drift is seen in the outdoor data likely due to thermal effects. The data set is also shorter, consisting of 7 hours of measurement.

parison shows the beginnings of some additional long term drift, likely due to thermal effects. Operation outside of the laboratory sees little effect on the static performance, demonstrating robustness of the instrument.

6.5.2 Out of lab performance and repeatability

To assess the performance in a civil engineering environment, the instrument is transported in a van to a prospective survey site. Overnight and during weekends, the instrument is stored in the workshop's generator house, where further optimisation of the sensor was performed.

Best performance out of the laboratory

Further optimisation of the experiment sequence was performed while in the field to improve the instrument performance. An improvement in the short term sensitivity was achieved when: the Raman beam intensity was

lowered, the bias coil current was reduced, and sub-level state selection was used. Lowering the Raman intensity creates longer Rabi π times, with an increase from 2.4 µs to 4 µs providing the best result. Through using the z compensation coils to split the magnetic sub-levels during sub-level state selection, the bias coil current was lowered from 100 mA to 25 mA. The reduction in current reduces the phase shifts and phase noise associated with magnetic field gradients, as discussed in section 2.3. Characterisation of the coupling between these parameters versus the instrument performance are still ongoing. The resulting Allan deviation, taken inside the generator house, is shown in figure 6.13, where the average short term sensitivity is measured as $(466 \pm 8)\,\mathrm{E}/\sqrt{\mathrm{Hz}}$, again at a 0.67 Hz measurement bandwidth, reaching the lowest value of 17 E in 5126 s. This is nearly a factor of two improvement to the best average short term sensitivity measured in the laboratory (figure 6.3), and within a factor of two of the detection limited sensitivity, calculated in section 6.2.1. All points of the ellipse, and the extracted phase from the fit of each 25 shots, is shown in figures 6.14 and 6.15.

Repeatability of measurements

An important consideration for a portable instrument is the repeatability of measurements taken at a certain point. Early attempts to measure the tare of the instrument showed jumps in the ellipse phase of 1000 E when picked up, moved, and placed back on the same point. Issues with rigidity in the structure were identified and improved; however, tare in the ellipse phase value was found to be lowest when moved from point to point using the trolley pictured in figure 3.14b.

Repeatability was determined through measurements on a base station

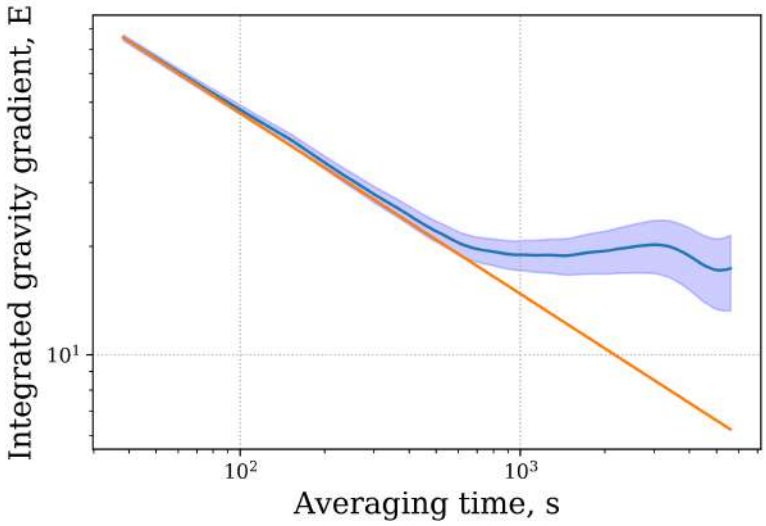

Figure 6.13 Performance of instrument out of the lab following optimisation. Short term sensitivity is calculated as $(466 \pm 8)\,\mathrm{E}/\sqrt{\mathrm{Hz}}$, at a 0.67 Hz measurement bandwidth, reaching the lowest value of 17 E in 5126 s for which the long term sensitivity may be estimated as $1200\,\mathrm{E}/\sqrt{\mathrm{Hz}}$.

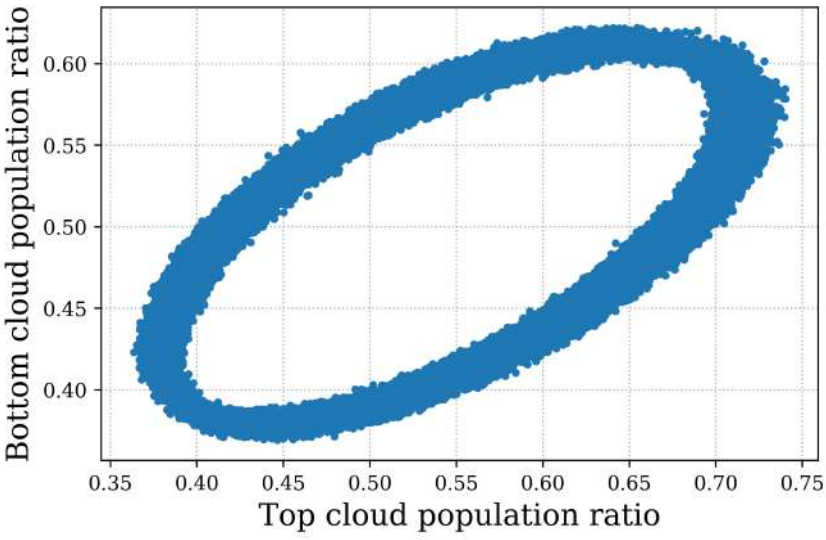

Figure 6.14 All points for 16 hour ellipse measurement performed in the generator house. Each 25 shots is fit with an ellipse to extract the phase value. The thickness of the total points suggests some drift in the ellipse centre or amplitudes.

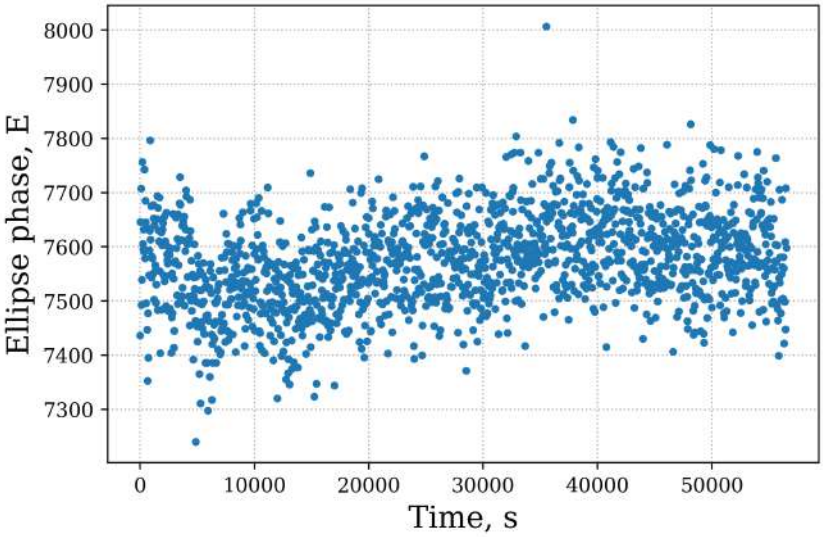

Figure 6.15 Extracted phase values of ellipse shown in units of gravity gradient. Some drift in the measured phase can be seen, which explains the thickness of the ellipse data in figure 6.14

point, moving from the point by several meters, left for tens of minutes, then returning to the base station. The measurements were performed outdoors, in order to assess the repeatability in a civil engineering setting. When moved, the instrument must be lifted from its tripod feet onto the trolley, moved by pushing the trolley, then the feet are re-engaged and levelled in the new measurement position. When returning to the base station, the process is repeated.

The resulting assessment of measurement repeatability will be a combination of instrument tare and phase drift, which cannot be easily separated. The results of the repeatability test are shown in figure 6.16, where the standard deviation of all measurements of phase is 58 E.

To examine the point to point repeatability, the phase difference between successive measurements is compared. The largest phase difference between successive points is measured as (132 ± 28) E. The mean difference between

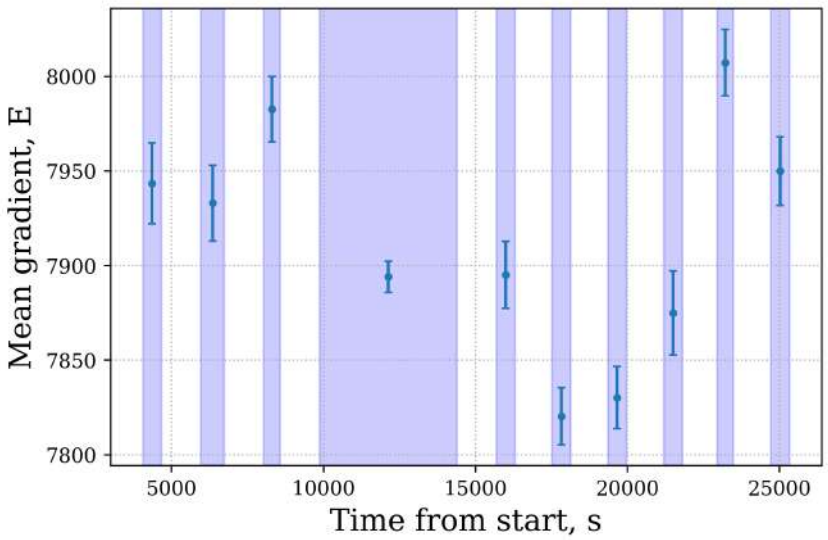

Figure 6.16 Instrument is moved from the base station and returned after a time, to evaluate the repeatability of the gradient measurement, whilst in a field environment. Average difference between successive points is (52 ± 40) E, showing excellent repeatability.

successive points however, is equal to (52 ± 40) E, showing a small shift. To isolate phase drift and instrument tare, the measurement should be repeated with long measurements on the base station to assess phase drift, followed by fast moves away and back to look for discontinuities in the ellipse phase caused by tare. In an instrument dominated by phase drift, a base station measurement technique may be used to remove linear drift when surveying across multiple spatial points.

6.6 Road to a civil engineering detection

The developed gravity gradiometer has demonstrated the performance required to measure anomalies, with gravity gradient signals of order 100 E expected to be resolved. This covers most of the civil engineering features outlined in the introduction (figure 1.2), including large pipes, tunnels and

sinkholes. In order to take the next step, one of these features must be detected in a relevant setting, requiring a few further system characterisations.

Many environmental challenges presented by operation in the field have been addressed in the design of this instrument, and where possible tested. A remaining challenge is the assessment of external magnetic fields on the performance of the instrument. Measurement sites typically contain lots of metal in the buried infrastructure, or have metal components forming the structures of the feature of interest. These tend to produce spatially varying magnetic fields whose shapes follow similar trends to the those from the gravitational field. In order to measure only the gravity gradient, effects from external magnetic fields must be well characterised and nulled. The inclusion of bulk shielding removes the majority of the field, but early testing suggests some sensitivity to external magnetic fields. To mitigate this, the implementation of area reversal, discussed in section 4.1.2, in the experiment sequence should allow for any phase shifts due to residual field effects to be removed. This requires first that the interferometer sequence be optimised for the reversed $\mathbf{k}_{\text{eff}}$, in order to have comparable performance in each direction, and to couple the effects equally into each sequence.

In order to increase the sensitivity of the instrument further, upgrading the laser and RF system will allow for the implementation of the atom launch, enabling an increase in the interferometer T time. The vacuum chamber has been designed to allow for a free-fall limited T time of 230 ms, which, accounting for similar reductions in the T time due to state preparation, will allow for ellipse measurements of up to T =215 ms. If the instrument operates at the same noise level as presented in section 6.5.2, the short term sensitivity could be improved to $73\,\mathrm{E}/\sqrt{\mathrm{Hz}}$, enabling averaging to within a few Eötvös in seconds. The atom launch can be achieved through

use of a moving molasses launch, demonstrated in the system for the throwing of one cloud; or a Bloch elevator [117], which uses a moving optical lattice to accelerate both atom clouds [118].

The ellipse contrast will decrease as the atom evolution time increases, as seen in section 5.6.2, resulting in a reduction of signal to noise, which will in turn impair the sensitivity. Improvements to the Rabi transfer at each height will reduce this effect, where techniques such as composite pulses may aid in recovering the dephasing due to thermal expansion.

Chapter 7

Conclusion

The instrument presented in this thesis, has demonstrated a short term sensitivity of $(466 \pm 8)\,\mathrm{E}/\sqrt{\mathrm{Hz}}$ in the field, at a measurement bandwidth of 0.67 Hz. A minimum resolution of 17 E was achieved, after an integration time of 5126 s. This represents the best sensitivity reported by any portable, atom interferometry based, gravity gradiometer. Measurements of test masses were used to demonstrate the ability to detect changes in gravity gradient, with good repeatability between measurement days observed, and general agreement to the modelled signal size. Evaluation of the instrument's performance, in and out of the laboratory, was consistent to within error, demonstrating no adverse effects from transportation, or operation outside of a temperature controlled environment. Point to point repeatability was assessed in the field, demonstrating an average shift between measurement points of $(52 \pm 40)\,\mathrm{E}$, when the instrument is moved on a trolley. A three times reduction in sensitivity to instrument tilt was measured, when compared to an ideal gravity measurement. Together, these tests show that the common mode noise suppression provided by atom interferometry is a viable technology for civil engineering applications, when

an instrument of this type is robustly packaged. The resulting sensitivity may be used to target most near surface civil engineering applications, with signal sizes of order 100 E.

Further work is required to characterise the instrument's sensitivity to external effects, such as background magnetic fields, before detection of a civil engineering feature, using measurements of the gravity gradient, can be convincingly made. The current short term sensitivity of the instrument is limited by technical noise to ten times the atom shot noise limit, evaluated at an interferometer T time of 85 ms. Roughly half of the instrument noise may be described by the atomic state detection, with the remaining noise yet to be characterised. The high level of system integration required for portable operation, saw traditional optimisation of the interferometer sequence forgone, making identification of noise sources difficult, and requiring some parameters to be measured indirectly.

At the instrument's current performance, many of the remaining challenges target the robustness and stability of the instrument. Following a detection of a civil engineering feature, transfer of knowledge into industry, may see commercial instruments of this type operational in the field within five years.

As well as the reduction of the remaining technical noise, the instrument provides a test bed for further physics challenges, in order to reach improved sensitivities. Implementation of an atom launch, either through a moving molasses or a Bloch elevator, may enable the use of longer interferometer T times, providing a T^2 improvement to the sensitivity. The sensor package has been designed to accommodate an atom free-fall limited interferometer T time of 230 ms, providing a seven times reduction in the atom shot noise limited sensitivity. Additionally, techniques such as large momentum beam

splitting will increase the interferometer area further within the same free-fall time. If the measurement bandwidth is prioritised, detection methods, such as phase shear, can be used to readout the interferometer within a single shot. All these methods to improve the sensitivity, demonstrated in the laboratory, have yet to translate to portable sensors and will face challenges in robust implementation. The sensor package, designed here, may facilitate the development of such advances, enabling instrument sensitivities never seen in any portable or transportable gravity sensors.

Appendix A

Rubidium-87 data

Table A.1 Properties of ^{87}Rb isotope and its D_2-line transition, values from [47], table reproduced with permission from [9]

Property	Symbol	Value
Mass	m	$1.443\,16 \cdot 10^{-25}\,\mathrm{kg}$
Wavelength (in vacuum)	λ	$780.24\,\mathrm{nm}$
$\lvert F=2\rangle \leftrightarrow \lvert F'=3\rangle$ transition frequency	$\nu_{2\leftrightarrow 3}$	$384.228\,115\,20\,\mathrm{THz}$
$\lvert F=1\rangle \leftrightarrow \lvert F'=2\rangle$ transition frequency	$\nu_{1\leftrightarrow 2}$	$384.234\,683\,23\,\mathrm{THz}$
$5^2S_{1/2}$ hyperfine splitting frequency	f_{hfs}	$6.834\,682\,610\,9\,\mathrm{GHz}$
Excited state lifetime	τ_0	$26.235\,\mathrm{ns}$
Natural linewidth	Γ	$2\pi{\cdot}6.067\,\mathrm{MHz}$
Recoil velocity	v_{recoil}	$5.8845\,\mathrm{mm\,s^{-1}}$
Doppler temperature	$\mathcal{T}_{\mathrm{D}}$	$145.6\,\mu\mathrm{K}$
Saturation intensity $\lvert F=2\rangle \leftrightarrow \lvert F'=3\rangle$ transition (isotropic light polarisation)	I_{sat}	$3.5771\,\mathrm{mW\,cm^{-2}}$

Appendix B

Additional theory

B.1 Propagation phase

As discussed in section 2.4.1, the phase of the interferometer is derived from three components: the propagation phase, the internal evolution phase, and the atom light interaction laser phase.

To demonstrate that the phase contribution due to propagation, even with a gravitational field, is zero, a path integral approach will be outlined. The paths the atoms take are close to those expected from a classical mechanics viewpoint, and it can be shown that the average path taken by the atoms does indeed follow the classically expected path [69].

The Lagrangian for a particle evolving in a gravitational field is given by

$$L\Big(z, v\Big) = \frac{1}{2}mv(t)^2 - mgz(t), \tag{B.1}$$

and the classical paths derived from the Lagrange equations are simply

the equations of motion with initial velocity v_0, position z_0, and time t_0,

$$v(t) = v_0 - g(t - t_0), \tag{B.2a}$$

$$z(t) = z_0 + v_1(t - t_0) - \frac{1}{2}g(t - t_0)^2. \tag{B.2b}$$

The path taken is determined by the principle of least action, where the action is defined as the integral of the Lagrangian for a given path. The classical path is a function of only the start and end point

$$S_{cl}(z_1, t_1, z_0, t_0) = \int_{t_0}^{t_1} L(z, v)dt = \int_{t_0}^{t_1} \left[\frac{1}{2}mv(t)^2 - mgz(t)\right]dt. \tag{B.3}$$

By defining the final velocity and position, as well as the initial velocity as

$$v_1 = v_0 - g(t_1 - t_0), \tag{B.4a}$$

$$z_1 = z_0 + v_1(t_1 - t_0) - \frac{1}{2}g(t_1 - t_0)^2, \tag{B.4b}$$

$$v_0 = \frac{z_1 - z_0}{t_1 - t_0} + \frac{1}{2}g(t_1 - t_0), \tag{B.4c}$$

it can be shown that the action is equal to [69]

$$S_{cl}(z_1, t_1, z_0, t_0) = \frac{m}{2}\frac{(z_1 - z_0)^2}{t_1 - t_0} - \frac{mg}{2}(z_1 + z_0)(t_1 - t_0) - \frac{mg^2}{24}(t_1 - t_0)^3. \tag{B.5}$$

Using the points defined in figure B.1, the trajectories are changed to be quadratic in time due to the influence of the gravitational acceleration. The

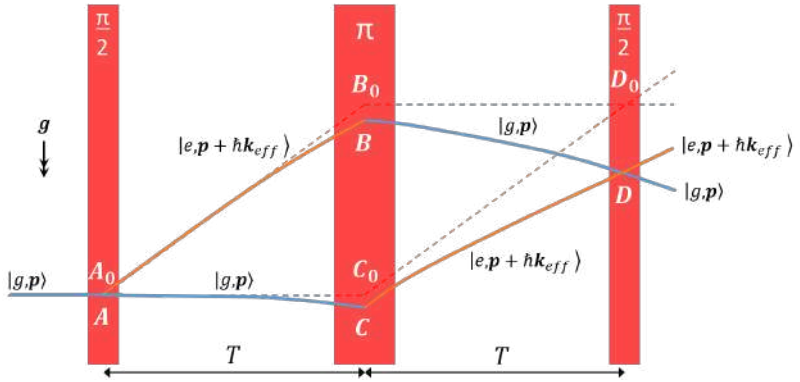

Figure B.1 With the presence of gravity, the space-time diagram of figure 2.7 has the classical paths shifted due to the atoms falling.

points are given by

$$A = A_0 \tag{B.6a}$$

$$B = B_0 - \frac{gT^2}{2} \tag{B.6b}$$

$$C = C_0 - \frac{gT^2}{2} \tag{B.6c}$$

$$D = D_0 - \frac{g(2T)^2}{2} = D_0 - 2gT^2. \tag{B.6d}$$

The phase shift between the arms of the interferometer, due to propagation, is given by the difference in the actions along the paths

$$\delta\phi^{prop} = S_{cl}(AB) + S_{cl}(BD) - \big(S_{cl}(AC) + S_{cl}(CD)\big). \tag{B.7}$$

Using equation (B.5) and the point definitions of equations (B.6), equation

(B.7) becomes

$$\delta\phi^{prop} = \frac{M}{T}(z_B - z_C)(z_B - z_A - z_D + z_c - gT^2) \tag{B.8a}$$

$$= \frac{M}{T}(z_B - z_C)(z_{B_0} - z_{A_0} - z_{D_0} + z_{C_0}). \tag{B.8b}$$

Since $A_0B_0C_0D_0$ is a parallelogram, the side lengths follow the relationship $z_{B_0} - z_{A_0} = z_{D_0} - z_{C_0}$, and thus

$$\delta\phi^{prop} = 0. \tag{B.9}$$

Therefore, it is has been shown that there is no phase contribution from the propagation of the atoms travelling different paths during the interferometer.

B.2 Calculating an Allan deviation

Let a time series measurement of the ellipse phase at time t be denoted $\phi(t)$, where each measurement is separated uniformly in time by a value τ_0 and a total number of samples N. The data can be grouped into clusters of length $\tau = k\tau_0$, for $k = 1 \longrightarrow k < N$. Let the average phase of each cluster containing m data points be defined as

$$\overline{\phi}_k(\tau) = \frac{1}{T}\int_{t_k}^{t_k+\tau} \phi(t)dt, \tag{B.10}$$

between times t_k and $t_{k+\tau}$.

The Allan variance is then defined as the infinite time average of the difference between neighbouring clusters over the limits $k = 1$ to $k = \infty$.

$$\sigma^2(\tau) = \frac{1}{2}\left\langle (\overline{\phi}_{k+1} - \overline{\phi}_k)^2 \right\rangle . \tag{B.11}$$

Since the data set is not an infinite time length and thus N is restricted, the Allan variance must be estimated for a sum over the available data sets. Thus equation B.11 becomes a discrete sum over the cluster averages.

For the overlapping average, which requires the sampling period is fixed, the estimate of the Allan variance of length τ is

$$\sigma^2(\tau) = \frac{1}{2(N-2m)} \sum_{k=1}^{N-2m} \left[\overline{\phi}_{k+m}(\tau) - \overline{\phi}_k(\tau)\right]^2 . \tag{B.12}$$

The Allan deviation $\sigma(t)$, is simply the square root of the Allan variance.

For the finite data set, the error of estimation of the Allan deviation increases as τ is increased, since there is not sufficient data to compare with. It can be shown that the percentage error of the Allan deviation is

$$\sigma(\delta) = \frac{1}{\sqrt{2\left(\frac{N}{m} - 1\right)}}, \tag{B.13}$$

and so as the number of data points in each cluster m, approaches the total number of samples N, the percentage error tends to infinity [113]. In practice, it is sensible to restrict the averaging time τ to less than 10 % of the total measurement set length, when calculating an Allan deviation using this functional form [111].

Appendix C

Vacuum assembly

C.1 Cleaning for UHV

Construction begins with cleaning parts from the workshop, and any contaminated pieces, using an adapted CERN cleaning procedure to prepare for UHV applications [119].

The steps performed were:

1. Hand scrub in detergent (*Alconox* Liquinox).
2. Rinse with tap water to remove detergent.
3. Ultrasonic bath in acetone to soften grease.
4. Ultrasonic bath in detergent to remove grease (*Ultrawave* Ultraclean M2).
5. Ultrasonic bath in distilled water, followed by a rinse in distilled water, to remove any remaining detergent.
6. Ultrasonic bath in acetone to de-grease further.

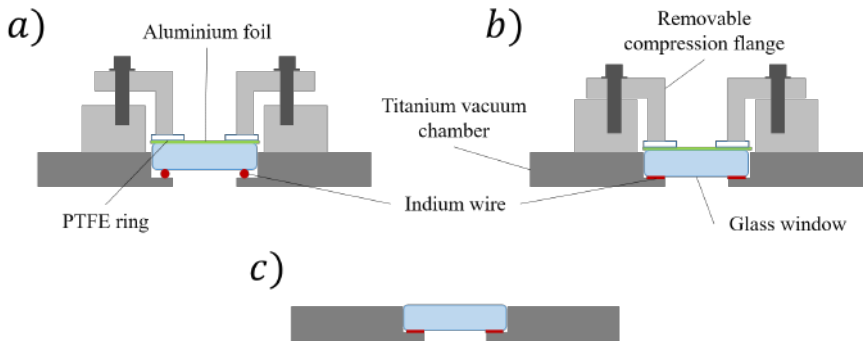

Figure C.1 a) Window and indium are placed onto chamber with the compression flange holding them loosly in place. b) Even torque is gradually applied to the screws in a star pattern with a torque wrench. c) After baking the vacuum chamber, the compression flange may be removed, leaving a sealed window with no additional bulky screws.

7. Ultrasonic bath in iso-propanol as a final clean.

8. Bake component in vacuum oven to remove moisture from within the bulk material, until ready for assembly.

Commercial parts from vacuum companies were not cleaned due to being shipped ready for UHV.

C.2 Indium sealing

Windows are attached to the science chambers using indium sealing. Clean, 1mm indium wire was compressed between a glass window and the chamber sealing surface using an evenly applied mechanical force. This is applied using a compression flange with a layer of clean foil, and PTFE ring between the flange and the window. The process is shown in figure C.1.

Each compression flange for the 1" windows consist of a mount bolted directly to the chamber, and the compression disc which has 6 M3 bolts, with spring washers, distributed evenly around a circle. The disc only touches the window around the edge, directly above the sealing surface.

Once the tension has been set correctly, the entire compression flange assembly can be removed, leaving a sealed window. The larger 60 mm top windows were compressed using a larger aluminium ring with 8 M3 bolts, with spring washers, fixing directly into the chamber. Each screw is incrementally torqued in a star pattern to 1.5 N m, with leaky windows having each screw tightened to no more than 1.9 N m to avoid smashing the glass.

By placing a layer of foil directly on the window, temperature gradients across the window generated during baking should be reduced. The PTFE ring serves to spread the pressure of the compression flange across the entire circumference of the window. Both are used to reduce the risk of the glass breaking due to stress.

C.3 UHV epoxy testing

To determine the suitability of EpoTek's UHV epoxy H21D, a test prism mount was made. The criteria for the glue to be suitable is ease of application, strength, and relative prism alignment.

A pair of mirror coated glass prisms* were bonded to the mount using the epoxy and cured at 150 °C for 1 hour in a vacuum oven, following the recommended procedure [120]. The horizontal alignment of the prisms is set by small alignment posts on the mount and should be on the order of manufacturing tolerances. Since the epoxy is bonding the prisms to a flat surface, the angular displacement of the prisms can be measured using the measurement scheme in figure C.2. The two prisms form a parallel reflector so changes in the distance between the input and output beam over a length will give an estimate of the angular prism alignment. Over a length of 1.77 m, a change in beam centres of 2.5±0.7 mm was measured. This gives

*Thorlabs, MRA15-E03

an estimate of the angular alignment of the prisms to be (80 ± 20) mdeg, which can be ascribed to the angular tolerance of the prisms [121]. Consequently, bonding the prisms to a machined flat surface with this UHV epoxy is suitable.

Repeat bakeouts of the vacuum chamber may be necessary, so the epoxy bond is tested for thermal stress failure. To do this, the room temperature mount is placed in the oven pre-heated to 150 °C. After a period of thermalisation, the mount is removed, allowed to cool, then inspected. Changes in the appearance of the epoxy were monitored, then each prism tapped and twisted to check the bond strength. After three repetitions, there was no noticeable change in the bond strength or appearance of the epoxy.

A final destructive test was performed to test for robustness, to mechanical shock. To simulate the instrument being dropped, the mount was dropped from increasing heights against an optical table while lightly being held from the sides. This stopped the prisms hitting the table first which would not be a feasible failure method in the instrument. In steps of roughly 10 cm, the bond on one prism failed at a drop height of 30 cm. Of the four bonding points, the epoxy failed on two points and the glass broke on the other two. The second prism which used slightly more epoxy was unaffected.

From these tests, the UHV epoxy, EPO-TEK®H21D, is deemed to be a suitable choice to bond the prisms to the mount, since it has demonstrated that it is at least a durable as the glass prisms, under thermal and mechanical shocks far outside of usual operating conditions.

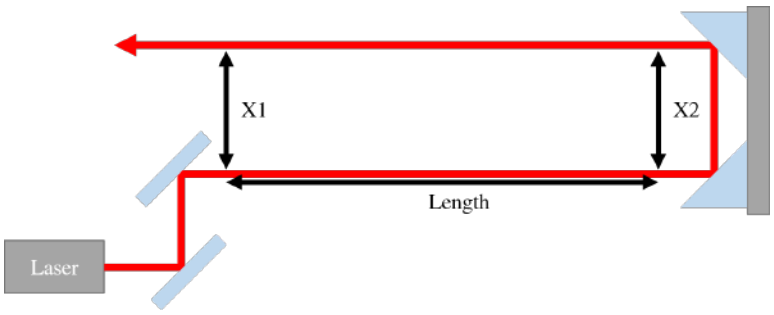

Figure C.2 Angular deviation of (80 ± 20) mdeg is measured for a test prism mount. Calculated using a change in X_n over a distance of 1.77 m.

Appendix D

Additional characterisation data

D.1 Time of flight

As well as the temperature of the atomic ensembles, characterised in section 5.1.2, the light sheet can also give an indication of the cooling power and atom number fluctuations. In the fit parameters, any change in the offset would correspond to a change in cooling laser power, and any change in the amplitude will be related to the atom number. Figure D.1 shows how the offset changes with time, showing a peak to peak change of 11 % for both chambers. This change in offset also appears to have similar features to the atomic signal amplitude. This may be explained by trapping fewer atoms due to lower power and seeing less fluorescence due to lower optical power.

To determine the effect of the initial cloud radius on the final ensemble temperature, different initial cloud radii were tested when fitting using formula 5.6. The objective was to determine the maximum cloud temperature

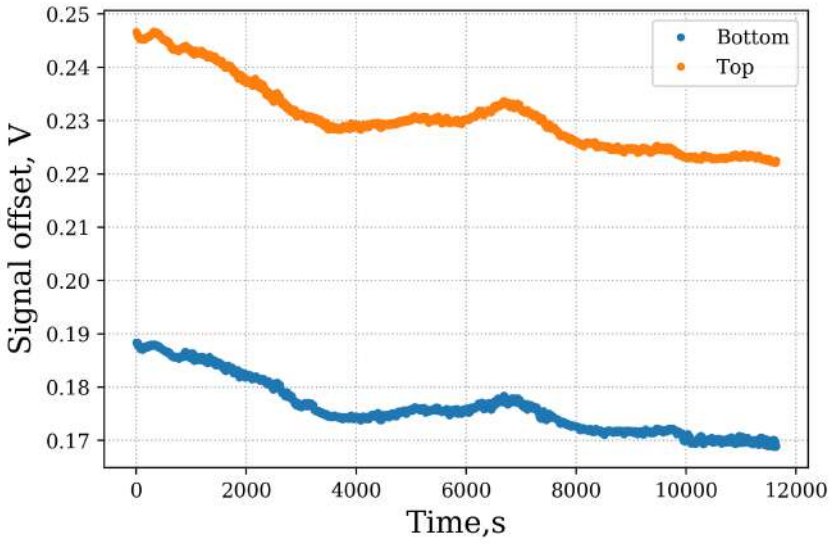

Figure D.1 Change in the offset seen with time for the time of flight light sheet measurements. Cooling intensity drift is the cause, showing a peak to peak change of 11 % in both chambers.

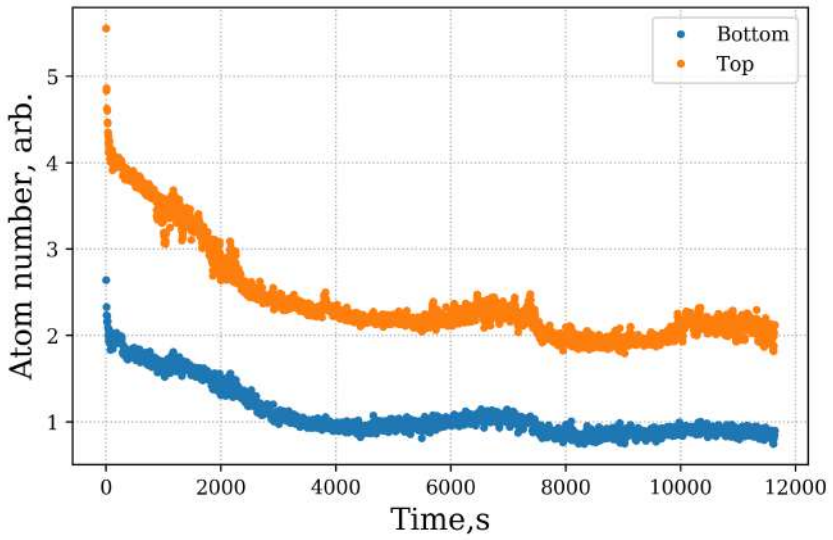

Figure D.2 Change in the fitting parameter a, proportional to the atom number. A decay is seen from the beginning as the instrument warms up into the duty cycle. Similar features are seen in the atom number data to the offset (figure D.1), showing that the changes in the cooling intensity affects the detected number of atoms.

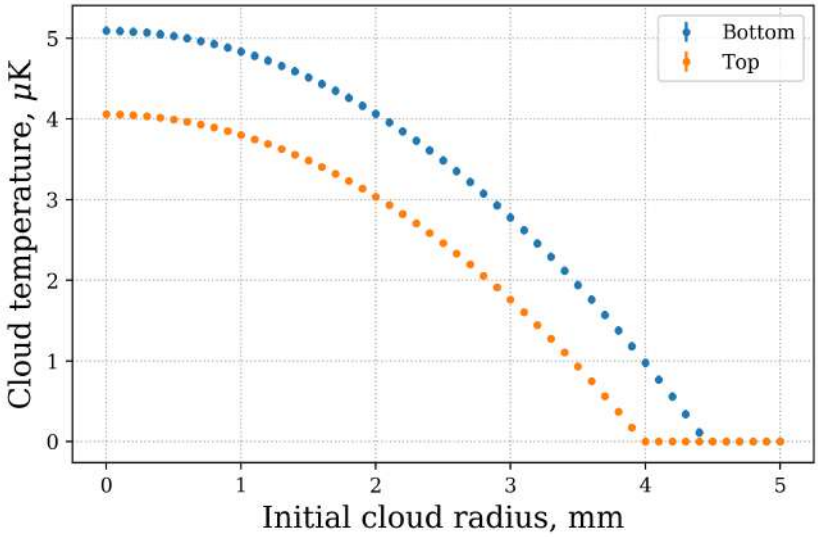

Figure D.3 Calculated cloud temperature for different initial cloud radii shows a biasing of the temperature measurement. Setting the initial radius to zero provides an upper bound on the ensemble temperature. As the initial radius approaches the final cloud radius, the temperature goes to zero. The error in temperature for each cloud radius is 0.017 µK and 0.05 µK for the top and bottom, and is the same to within 1 % for non-zero cloud temperatures.

as well as determine any increase in error in fitting. The result of changing the cloud radii is shown in figure D.3 and shows that the maximum cloud temperature when the initial cloud radius is zero is 4.1 µK and 5.1 µK for the top and bottom cloud respectively. The error in temperature is unchanged for non-zero values of temperature, and is evaluated at 0.017 µK and 0.05 µK for the top and bottom atom clouds, respectively. Therefore, the initial cloud radius adds only a bias to the calculated cloud temperature.

D.2 VRP settings

The variable retarder plates (VRPs), used in the telescopes set the polarisation of the light for the cooling and Raman beams. With the voltage value determined off the system, optimisation on system can be done using the

atomic signal. To optimise the linear polarisation required for the Raman beams, co-propagating transitions were minimised in the spectrum scans. This means that one of the beams is polarised to be as linear as possible. Varying the applied voltage to the other beam allows for the maximising of the counter-propagating transitions. This method allows for the setting of two perpendicular linear beams.

To determine the switching speed of the variable retarder plates, Rabi oscillations were taken and the π times evaluated. When circularly polarised, the π time is very slow, and when linear polarised, it is fast. The switching time is determined by the length of time required for the π time to stabilise. Measurements show that the polarisation can be switched from circular to linear polarisation in less than 1 ms; however, the switch from linear to circular polarisation takes around 20 ms.

D.3 Atom number in detection

The detection scheme used in the experiment makes an accurate determination of an atom number tricky. Scanning the detection frequency (presented in figure 5.11), revealed that the frequency which gave the most signal was not as one would expect from simple calculations of resonance frequency, with a possible explanation of a 1D-molasses being formed. Consequently, the determination of the detuning at detection may not accurately be calculated, but will be modelled as resonant light to give an estimate of the atom number.

To calculate the atom number from a photodiode signal, the key parameters are: measured voltage, V, photodiode resistance or amplifier circuit gain, G, the responsivity of the photodetector, R, the light intensity, I, and the detection solid angle, Ω. In order to reach the calculation of atom

number, some values must be derived from these parameters.

An atom will scatter a number of photons based on the intensity, I, and detuning, δ, of the light. To calculate the scattering rate, Γ_{scatt}, take the excited state population fraction, ρ_0, and times by the transition linewidth, Γ. For ^{87}Rb, the transition linewidth for the D_2-line is $2\pi \cdot 6.067\,\mathrm{MHz}$, and the saturation intensity I_0, for the transition $|F = 1\rangle \leftrightarrow |F' = 2\rangle$ is $1.67\,\mathrm{mW\,cm^{-2}}$ [47]. The scattering rate is therefore given by

$$\Gamma_{scatt} = \Gamma\rho_0 = \frac{\Gamma}{2} \cdot \frac{I/I_0}{1 + I/I_0 + (2\delta/\Gamma)^2}. \tag{D.1}$$

The maximum fraction of the atoms in the excited state is 0.5, and occurs when the intensity is much greater than the saturation intensity and the detuning is zero. For the detection intensity of $2.4\,\mathrm{mW\,cm^{-2}}$ from each vertical detection beam, the excited state fraction is estimated as 0.41. The number of photon scatters per atom, per second is therefore $1.57 \cdot 10^7$.

Assuming the atoms radiate evenly over a sphere, the photodiode can capture photons through a solid angle of 0.159 sr, dictated by the chamber diameter and viewport. The quantum efficiency of the photodiode may be calculated from the responsivity and the wavelength. For a responsivity of $0.43\,\mathrm{A\,W^{-1}}$, the quantum efficiency is 68 % at 780 nm. The gain of the amplifier circuit is $1 \cdot 10^7$. For N atoms, the voltage measured from the photodiode circuit is therefore

$$V = e\Gamma_{scatt}\Omega NQG, \tag{D.2}$$

where e is the electron charge. Rearranging for the number of atoms gives

$$N = \frac{V}{e\Gamma_{scatt}\Omega QG}. \tag{D.3}$$

An estimation of the number of atoms transferred in the interferometer comes from the number of atoms transferred at the Rabi π time. For the top and bottom chambers, the photodiode voltage changes by 52 mV and 33 mV at this time. From this, an estimated $2.7 \cdot 10^5$ and $1.7 \cdot 10^5$ atoms take part in the top and bottom interferometers, respectively.

List of Figures

1.1 Instrument tilt on the projection of gravity and gravity gradient 10

1.2 Civil engineering parameter space 13

2.1 MOT mechanism 18

2.2 ^{87}Rb energy level diagram 20

2.3 Raman transition 23

2.4 Rabi oscillation theory 31

2.5 Velocity selective Raman pulses 36

2.6 Sub-level splitting in a magnetic field 37

2.7 Mach-Zehnder sequence 39

2.8 Gravity gradiometer using two Mach-Zehnder interferometers. 45

2.9 Fringe to ellipse 50

2.10 Shift due to Coriolis force 53

2.11 Allan deviation of phase data for initial demonstrator 55

2.12 Modelled gravity anomaly 57

2.13 Baseline and T time optimisation for the new sensor head . . 59

2.14 Prism MOT operation 63

2.15 Sensor head design 65

3.1 Launch and drop configurations 69

3.2 Vacuum chamber overview 70

3.3 Prism mount 73
3.4 Interfacing module 75
3.5 Bakeout graph 77
3.6 Mounting structure 80
3.7 MOT coil spacing 81
3.8 Measured MOT coil field 82
3.9 Compensation coil measurements 83
3.10 Bias coil optimised field geometry 86
3.11 Bias coil winding and fitted to system. 87
3.12 Measurements of bias coil field profile using Gauss meter . . . 88
3.13 Exploded view of telescopes 91
3.14 Assembled sensor head with and without bias coils and shields 98

4.1 Laser and RF diagram 101
4.2 Ellipse fitting on synthetic data 107
4.3 Ellipse phase with time 108
4.4 Fitting bias of ellipse phase for number of shots 109
4.5 Short term sensitivity vs number of shots per ellipse 109
4.6 Fitting bias of ellipse phase for applied phase shift 110

5.1 MOT detuning 114
5.2 Comparison of MOT loading curves 116
5.3 MOT fluorescence vs dispenser current 116
5.4 MOT loading time vs dispenser current 118
5.5 Equilibrium atom number vs loading time 119
5.6 Molasses sequence 121
5.7 F=1 state preparation 122
5.8 Light sheet trace 123

5.9 Temperature vs. time 125
5.10 Detection scheme 127
5.11 Detection frequency scan 129
5.12 Spectrum scan for both chambers 131
5.13 Peaks from all spectrum scans 133
5.14 Total magnetic field from spectrum 133
5.15 Background magnetic field measurement using the atoms . . 135
5.16 Bias field measurement using the atoms 135
5.17 Bias field overlayed with the modelled field profile. 136
5.18 Bias coil switching 137
5.19 Rabi oscillation at each interferometer pulse time 139
5.20 Rabi oscillations at all heights through the chamber 140
5.21 Rabi transfer at all heights 141
5.22 F=1 blow-away 144
5.23 State selection spectrum. 145
5.24 Rabi oscillation with state selection 146
5.25 State selection ellipse comparison 147
5.26 Ellipses with and without applied phase shift 150
5.27 Ellipse phase with applied frequency shift 150
5.28 Ellipse contrast vs T time 152
5.29 Interferometer contrast standard deviation vs T time 153
5.30 Phase shift vs interferometer T time 153
5.31 Phase standard deviation vs T time 154

6.1 Long-term ellipse measurement. 158
6.2 Ellipse phase vs time in temperature stabilised lab 159
6.3 Long term Allan deviation 160
6.4 New background subtraction 166

6.5 Measured sensitivity to instrument tilt 169
6.6 Ellipse phase and the ambient temperature vs time 170
6.7 Ellipse phase versus ambient temperature 171
6.8 Coriolis effect as measured by instrument. 172
6.9 Test mass setup 173
6.10 Two point test mass results 174
6.11 Test mass data with model 175
6.12 Comparison of system performance in and out of the lab . . . 177
6.13 Performance of instrument out of lab 179
6.14 Generator house ellipse measurement 179
6.15 Generator house phase values 180
6.16 Repeatability of the gradient measurement 181

B.1 Mach-Zehnder with gravity 191

C.1 Indium sealing 195
C.2 Prism mount alignment 198

D.1 Time of flight offset 200
D.2 Time of flight atom number 200
D.3 Bias in cloud temperature for time of flight measurement . . 201

List of Tables

2.1 Light shift contributions in interferometer phase 40

A.1 Properties of rubidium-87 and D_2-line transition 188

List of References

[1] H. Kater and T. Young, *IV. An Account of experiments for determining the length of the pendulum vibrating seconds in the latitude of London*, Philosophical Transactions of the Royal Society of London **108** 33–102 (1818).

[2] L. Eötvös, *Über die Anziehung der Erde auf verschiedene Substanzen*, Mathematische und naturwissenschaftliche Berichte aus Ungarn **8** 65–68 (1890).

[3] W. Torge and J. Müller, *Geodesy*. Walter de Gruyter (2012).

[4] J. A. Hammond and J. E. Faller, *A laser-interferometer system for the absolute determination of the acceleration due to gravity*, Precision Measurement and Fundamental Constants; Proceedings **343** 457 (1971).

[5] J. M. Goodkind, *The superconducting gravimeter*, Review of Scientific Instruments **70** (11), 4131–4152 (1999).

[6] L. Timmen and O. Gitlein, *The capacity of the Scintrex Autograv CG-3M no. 4492 gravimeter for "absolute-scale" surveys*, Revista Brasileira de Cartografia **56** (2004).

[7] M. Kasevich and S. Chu, *Atomic interferometry using stimulated Raman transitions*, Phys. Rev. Lett. **67** 181–184 (1991).

[8] M. Kasevich and S. Chu, *Measurement of the gravitational acceleration of an atom with a light-pulse atom interferometer*, Applied Physics B **54** (5), 321–332 (1992).

[9] C. V. Rammeloo, *Optimisation of a compact cold-atoms interferometer for gravimetry.* Ph.D. thesis, University of Birmingham (2018).

[10] A. Lamb, *Cold atom gravity gradiometer for field applications.* Ph.D. thesis, University of Birmingham (2019).

[11] M. Schmidt, *A mobile high-precision gravimeter based on atom interferometry.* Ph.D. thesis, Humboldt-Universität zu Berlin, Mathematisch-Naturwissenschaftliche Fakultät I (2011).

[12] T. Lévèque, L. Antoni-Micollier, B. Faure and J. Berthon, *A laser setup for rubidium cooling dedicated to space applications*, Applied Physics B **116** (4), 997–1004 (2014).

[13] Q. Wang, Z. Wang, Z. Fu, W. Liu and Q. Lin, *A compact laser system for the cold atom gravimeter*, Optics Communications **358** 82–87 (2016).

[14] V. Schkolnik et al., *A compact and robust diode laser system for atom interferometry on a sounding rocket*, Applied Physics B **122** (8), 217 (2016).

[15] V. Ménoret, P. Vermeulen, N. Le Moigne, S. Bonvalot, P. Bouyer, A. Landragin and B. Desruelle, *Gravity measurements below* 10^{-9} g

with a transportable absolute quantum gravimeter, Scientific reports **8** (1), 1–11 (2018).

[16] X. Wu, Z. Pagel, B. S. Malek, T. H. Nguyen, F. Zi, D. S. Scheirer and H. Müller, *Gravity surveys using a mobile atom interferometer*, Science Advances **5** (9), (2019).

[17] D. Boddice, N. Metje and G. Tuckwell, *Capability assessment and challenges for quantum technology gravity sensors for near surface terrestrial geophysical surveying*, Journal of Applied Geophysics **146** 149 – 159 (2017).

[18] K. Bongs, M. Holynski, J. Vovrosh, P. Bouyer, G. Condon, E. Rasel, C. Schubert, W. P. Schleich and A. Roura, *Taking atom interferometric quantum sensors from the laboratory to real-world applications*, Nature Reviews Physics 1–9 (2019).

[19] H. D. Young, R. A. Freedman, T. Sandin and A. L. Ford, *University physics*, vol. 9. Addison-Wesley New York (1996).

[20] M. Caputo and L. Pieri, *The normal gravity formula and the polar flattening according to geodetic reference system 1967*, Annals of Geophysics **21** (1), 123–149 (1968).

[21] W. Torge, *Gravimetry*. W. de Gruyter (1989).

[22] A. Hinton et al., *A portable magneto-optical trap with prospects for atom interferometry in civil engineering*, Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences **375** (2099), 20160238 (2017).

[23] G. de Villiers, K. Ridley, A. Rodgers and D. Boddice, *On the use of*

the profiled singular-function expansion in gravity gradiometry, Journal of Applied Geophysics **170** 103830 (2019).

[24] A. D. Rodgers, *Bayesian inference from terrestrial gravimetry measurements of near-surface anomalies using a bespoke reversible-jump Markov chain Monte Carlo algorithm.* Ph.D. thesis, University of Birmingham (2017).

[25] R. Warburton, H. Pillai and R. Reineman, *Initial Results with the New GWR iGrav™ Superconducting Gravity Meter*, (2010).

[26] T. Niebauer, R. Billson, B. Ellis, B. Mason, D. van Westrum and F. Klopping, *Simultaneous gravity and gradient measurements from a recoil-compensated absolute gravimeter*, Metrologia **48** (3), 154 (2011).

[27] J. E. Faller, *Thirty years of progress in absolute gravimetry: a scientific capability implemented by technological advances*, Metrologia **39** (5), 425–428 (2002).

[28] I. A. Robinson and S. Schlamminger, *The watt or Kibble balance: a technique for implementing the new SI definition of the unit of mass*, Metrologia **53** (5), A46 (2016).

[29] R. Reudink, R. Klees, O. Francis, J. Kusche, R. Schlesinger, A. Shabanloui, N. Sneeuw and L. Timmen, *High tilt susceptibility of the Scintrex CG-5 relative gravimeters*, Journal of Geodesy **88** (6), 617–622 (2014).

[30] *CG-6 AUTOGRAV TM SURVEY GRAVITY METER.* Rev. C, accessed 10/01/2021.

[31] A. Peters, K. Y. Chung and S. Chu, *High-precision gravity measurements using atom interferometry*, Metrologia **38** (1), 25–61 (2001).

[32] Z.-K. Hu, B.-L. Sun, X.-C. Duan, M.-K. Zhou, L.-L. Chen, S. Zhan, Q.-Z. Zhang and J. Luo, *Demonstration of an ultrahigh-sensitivity atom-interferometry absolute gravimeter*, Phys. Rev. A **88** 043610 (2013).

[33] P. Gillot, O. Francis, A. Landragin, F. P. Dos Santos and S. Merlet, *Stability comparison of two absolute gravimeters: optical versus atomic interferometers*, Metrologia **51** (5), L15 (2014).

[34] G. D'Amico, F. Borselli, L. Cacciapuoti, M. Prevedelli, G. Rosi, F. Sorrentino and G. Tino, *Bragg interferometer for gravity gradient measurements*, Physical Review A **93** (6), 063628 (2016).

[35] P. Altin et al., *Precision atomic gravimeter based on Bragg diffraction*, New Journal of Physics **15** (2), 023009 (2013).

[36] H. Müller, S.-w. Chiow, Q. Long, S. Herrmann and S. Chu, *Atom Interferometry with up to 24-Photon-Momentum-Transfer Beam Splitters*, Phys. Rev. Lett. **100** 180405 (2008).

[37] T. Mazzoni, X. Zhang, R. Del Aguila, L. Salvi, N. Poli and G. M. Tino, *Large-momentum-transfer Bragg interferometer with strontium atoms*, Phys. Rev. A **92** 053619 (2015).

[38] J. M. McGuirk, G. T. Foster, J. B. Fixler, M. J. Snadden and M. A. Kasevich, *Sensitive absolute-gravity gradiometry using atom interferometry*, Phys. Rev. A **65** 033608 (2002).

[39] R. H. Parker, C. Yu, W. Zhong, B. Estey and H. Müller, *Measurement of the fine-structure constant as a test of the Standard Model*, Science **360** (6385), 191–195 (2018).

[40] G. Rosi, F. Sorrentino, L. Cacciapuoti, M. Prevedelli and G. M. Tino, *Precision measurement of the Newtonian gravitational constant using cold atoms*, Nature **510** 518–521 (2014).

[41] L. Zhou et al., *Test of Equivalence Principle at* 10^{-8} *Level by a Dual-Species Double-Diffraction Raman Atom Interferometer*, Phys. Rev. Lett. **115** 013004 (2015).

[42] S. Dimopoulos, P. W. Graham, J. M. Hogan and M. A. Kasevich, *Testing General Relativity with Atom Interferometry*, Phys. Rev. Lett. **98** 111102 (2007).

[43] P. W. Graham, J. M. Hogan, M. A. Kasevich, S. Rajendran and R. W. Romani, *Mid-band gravitational wave detection with precision atomic sensors*, arXiv preprint arXiv:1711.02225 (2017).

[44] L. Badurina et al., *AION: an atom interferometer observatory and network*, Journal of Cosmology and Astroparticle Physics **2020** (05), 011–011 (2020).

[45] P. W. Graham, J. M. Hogan, M. A. Kasevich and S. Rajendran, *New Method for Gravitational Wave Detection with Atomic Sensors*, Phys. Rev. Lett. **110** 171102 (2013).

[46] Health and Safety Executive, *Pipelines and gas supply industry: FAQs*. `https://www.hse.gov.uk/pipelines/faqs.htm`. Accessed: 2020-05-07.

[47] D. A. Steck, *Rubidium 87 D Line Data*, Technical Report, University of Oregon (2019). `http://steck.us/alkalidata`. (revision 2.2.1, 21 November 2019).

[48] D. Becker et al., *Space-borne Bose–Einstein condensation for precision interferometry*, Nature **562** (7727), 391–395 (2018).

[49] T. van Zoest et al., *Bose-Einstein Condensation in Microgravity*, Science **328** (5985), 1540–1543 (2010).

[50] J. Rudolph et al., *A high-flux BEC source for mobile atom interferometers*, New Journal of Physics **17** (6), 065001 (2015).

[51] G. Gattobigio, T. Pohl, G. Labeyrie and R. Kaiser, *Scaling laws for large magneto-optical traps*, Physica Scripta **81** (2), 025301 (2010).

[52] H. J. Metcalf and P. van der Straten, *Laser Cooling and Trapping*. Springer (1999).

[53] C. J. Foot et al., *Atomic physics*, vol. 7. Oxford University Press (2005).

[54] P. D. Lett, R. N. Watts, C. I. Westbrook, W. D. Phillips, P. L. Gould and H. J. Metcalf, *Observation of Atoms Laser Cooled below the Doppler Limit*, Phys. Rev. Lett. **61** 169–172 (1988).

[55] A. Steane, M. Chowdhury and C. Foot, *Radiation force in the magneto-optical trap*, JOSA B **9** (12), 2142–2158 (1992).

[56] J. Schoser, A. Batär, R. Löw, V. Schweikhard, A. Grabowski, Y. B. Ovchinnikov and T. Pfau, *Intense source of cold Rb atoms from a pure two-dimensional magneto-optical trap*, Phys. Rev. A **66** 023410 (2002).

[57] D. S. Goodman, I. Sivarajah, J. E. Wells, F. A. Narducci and W. W. Smith, *Ion–neutral-atom sympathetic cooling in a hybrid linear rf Paul and magneto-optical trap*, Phys. Rev. A **86** 033408 (2012).

[58] J. Barry, D. McCarron, E. Norrgard, M. Steinecker and D. DeMille, *Magneto-optical trapping of a diatomic molecule*, Nature **512** (7514), 286–289 (2014).

[59] M. H. Anderson, J. R. Ensher, M. R. Matthews, C. E. Wieman and E. A. Cornell, *Observation of Bose-Einstein condensation in a dilute atomic vapor*, science **269** (5221), 198–201 (1995).

[60] T. Lévèque, A. Gauguet, F. Michaud, F. Pereira Dos Santos and A. Landragin, *Enhancing the Area of a Raman Atom Interferometer Using a Versatile Double-Diffraction Technique*, Phys. Rev. Lett. **103** 080405 (2009).

[61] B. Young, M. Kasevich and S. Chu, *Precision Atom Interferometry with Light Pulses*, in P. R. Berman *Atom Interferometry*, 363 – 406. Academic Press, San Diego (1997).

[62] A. Niggebaum, *Towards mobile quantum sensors for gravity surveys.* Ph.D. thesis, University of Birmingham (2015).

[63] K. Moler, D. S. Weiss, M. Kasevich and S. Chu, *Theoretical analysis of velocity-selective Raman transitions*, Physical Review A **45** (1), 342 (1992).

[64] J. Stammers, *An atom interferometer for measuring horizontal accelerations.* Ph.D. thesis, Imperial College London (2018).

[65] M. Kasevich and S. Chu, *Laser cooling below a photon recoil with three-level atoms*, Phys. Rev. Lett. **69** 1741–1744 (1992).

[66] A. J. Kerman, V. Vuletić, C. Chin and S. Chu, *Beyond Optical Molasses: 3D Raman Sideband Cooling of Atomic Cesium to High Phase-Space Density*, Phys. Rev. Lett. **84** 439–442 (2000).

[67] L. Zhu, *A cold atoms gravimeter for use in absolute gravity comparisons*. Ph.D. thesis, University of Birmingham (2018).

[68] K. P. Zetie, S. F. Adams and R. M. Tocknell, *How does a Mach-Zehnder interferometer work?*, Physics Education **35** (1), 46–48 (2000).

[69] P. Storey and C. Cohen-Tannoudji, *The Feynman path integral approach to atomic interferometry. A tutorial*, Journal de Physique II **4** (11), 1999–2027 (1994).

[70] B. Barrett, P. A. Gominet, E. Cantin, L. Antoni-Micollier, A. Bertoldi, B. Battelier, P. Bouyer, J. Lautier and A. Landragin, *Mobile and remote inertial sensing with atom interferometers*, in M. Kasevich and G. Tino *Atom Interferometry*, vol. 188 of *Proceedings of the International School of Physics "Enrico Fermi"*, 493 – 555. IOS Press (2014).

[71] D. S. Weiss, B. C. Young and S. Chu, *Precision measurement of $\hbar/m$ Cs based on photon recoil using laser-cooled atoms and atomic interferometry*, Applied physics B **59** (3), 217–256 (1994).

[72] A. Roura, *Circumventing Heisenberg's Uncertainty Principle in Atom Interferometry Tests of the Equivalence Principle*, Phys. Rev. Lett. **118** 160401 (2017).

[73] G. T. Foster, J. B. Fixler, J. M. McGuirk and M. A. Kasevich,

Method of phase extraction between coupled atom interferometers using ellipse-specific fitting, Opt. Lett. **27** (11), 951–953 (2002).

[74] G. Rosi, L. Cacciapuoti, F. Sorrentino, M. Menchetti, M. Prevedelli and G. Tino, *Measurement of the gravity-field curvature by atom interferometry*, Physical review letters **114** (1), 013001 (2015).

[75] S.-w. Chiow, J. Williams and N. Yu, *Noise reduction in differential phase extraction of dual atom interferometers using an active servo loop*, Physical Review A **93** (1), 013602 (2016).

[76] F. P. Dos Santos, *Differential phase extraction in an atom gradiometer*, Phys. Rev. A **91** 063615 (2015).

[77] A. Sugarbaker, S. M. Dickerson, J. M. Hogan, D. M. S. Johnson and M. A. Kasevich, *Enhanced Atom Interferometer Readout through the Application of Phase Shear*, Phys. Rev. Lett. **111** 113002 (2013).

[78] B. Wu et al., *Dependence of the sensitivity on the orientation for a free-fall atom gravimeter*, Opt. Express **27** (8), 11252–11263 (2019).

[79] K. Ridley, *Private communication.*

[80] P. Kearey, M. Brooks and I. Hill, *An introduction to geophysical exploration.* John Wiley & Sons (2013).

[81] G. M. Tino and M. A. Kasevich, *Atom interferometry*, vol. 188. IOS Press (2014).

[82] W.-J. Xu et al., *Effects of wave-front tilt and air density fluctuations in a sensitive atom interferometry gyroscope*, Opt. Express **28** (8), 12189–12200 (2020).

[83] A. Trimeche, M. Langlois, S. Merlet and F. P. Dos Santos, *Active control of laser wavefronts in atom interferometers*, Physical Review Applied **7** (3), 034016 (2017).

[84] K. I. Lee, J. A. Kim, H. R. Noh and W. Jhe, *Single-beam atom trap in a pyramidal and conical hollow mirror*, Opt. Lett. **21** (15), 1177–1179 (1996).

[85] J. Arlt, O. Marago, S. Webster, S. Hopkins and C. Foot, *A pyramidal magneto-optical trap as a source of slow atoms*, Optics communications **157** (1-6), 303–309 (1998).

[86] L. Zhu et al., *A dielectric metasurface optical chip for the generation of cold atoms*, Science Advances **6** (31), (2020).

[87] J. P. McGilligan, P. F. Griffin, R. Elvin, S. J. Ingleby, E. Riis and A. S. Arnold, *Grating chips for quantum technologies*, Scientific reports **7** (1), 1–7 (2017).

[88] A. Rakholia, *High data-rate atom interferometry for measuring dynamic inertial conditions*. Ph.D. thesis, Sandia National Lab.(SNL-NM), Albuquerque, NM (United States) (2015).

[89] A. Dunning, R. Gregory, J. Bateman, N. Cooper, M. Himsworth, J. A. Jones and T. Freegarde, *Composite pulses for interferometry in a thermal cold atom cloud*, Physical Review A **90** (3), 033608 (2014).

[90] U. D. Rapol, A. Wasan and V. Natarajan, *Loading of a Rb magneto-optic trap from a getter source*, Physical Review A **64** (2), 023402 (2001).

[91] H. Kurisu, K. Ishizawa, S. Yamamoto, M. Hesaka and Y. Saito, *Application of titanium materials to vacuum chambers and*

components, Journal of Physics: Conference Series **100** (9), 092002 (2008).

[92] M. Holynski, *Creating a two dimensional cold mixture experiment.* Ph.D. thesis, University of Birmingham (2012).

[93] M. E. Rudd and J. Craig, *Optimum spacing of square and circular coil pairs*, Review of Scientific Instruments **39** (9), 1372–1374 (1968).

[94] F. Theron, O. Carraz, G. Renon, N. Zahzam, Y. Bidel, M. Cadoret and A. Bresson, *Narrow linewidth single laser source system for onboard atom interferometry*, Applied Physics B **118** (1), 1–5 (2015).

[95] G. Ritt, G. Cennini, C. Geckeler and M. Weitz, *Laser frequency offset locking using a side of filter technique*, Applied Physics B **79** (3), 363–365 (2004).

[96] D. McCarron, S. King and S. Cornish, *Modulation transfer spectroscopy in atomic rubidium*, Measurement science and technology **19** (10), 105601 (2008).

[97] A. B. Deb, A. Rakonjac and N. Kjærgaard, *Versatile laser system for experiments with cold atomic gases*, J. Opt. Soc. Am. B **29** (11), 3109–3113 (2012).

[98] G. W. Hoth, E. A. Donley and J. Kitching, *Atom number in magneto-optic traps with millimeter scale laser beams*, Optics letters **38** (5), 661–663 (2013).

[99] O. Carraz, R. Charrière, M. Cadoret, N. Zahzam, Y. Bidel and A. Bresson, *Phase shift in an atom interferometer induced by the additional laser lines of a Raman laser generated by modulation*, Phys. Rev. A **86** 033605 (2012).

[100] L. Zhu, Y.-H. Lien, A. Hinton, A. Niggebaum, C. Rammeloo, K. Bongs and M. Holynski, *Application of optical single-sideband laser in Raman atom interferometry*, Opt. Express **26** (6), 6542–6553 (2018).

[101] A. Louchet-Chauvet, T. Farah, Q. Bodart, A. Clairon, A. Landragin, S. Merlet and F. P. D. Santos, *The influence of transverse motion within an atomic gravimeter*, New Journal of Physics **13** (6), 065025 (2011).

[102] A. Fitzgibbon, M. Pilu and R. B. Fisher, *Direct least square fitting of ellipses*, IEEE Transactions on pattern analysis and machine intelligence **21** (5), 476–480 (1999).

[103] R. W. G. Moore, L. A. Lee, E. A. Findlay, L. Torralbo-Campo, G. D. Bruce and D. Cassettari, *Measurement of vacuum pressure with a magneto-optical trap: A pressure-rise method*, Review of Scientific Instruments **86** (9), 093108 (2015).

[104] T. Arpornthip, C. A. Sackett and K. J. Hughes, *Vacuum-pressure measurement using a magneto-optical trap*, Phys. Rev. A **85** 033420 (2012).

[105] T. G. Petelski, *Atom interferometers for precision gravity measurements*. Ph.D. thesis, Paris 6 (2005).

[106] D. S. Weiss, E. Riis, Y. Shevy, P. J. Ungar and S. Chu, *Optical molasses and multilevel atoms: experiment*, J. Opt. Soc. Am. B **6** (11), 2072–2083 (1989).

[107] T. M. Brzozowski, M. Maczynska, M. Zawada, J. Zachorowski and W. Gawlik, *Time-of-flight measurement of the temperature of cold*

atoms for short trap-probe beam distances, Journal of Optics B: Quantum and Semiclassical Optics **4** (1), 62–66 (2002).

[108] A. V. Rakholia, H. J. McGuinness and G. W. Biedermann, *Dual-Axis High-Data-Rate Atom Interferometer via Cold Ensemble Exchange*, Phys. Rev. Applied **2** 054012 (2014).

[109] M. Kasevich, D. S. Weiss, E. Riis, K. Moler, S. Kasapi and S. Chu, *Atomic velocity selection using stimulated Raman transitions*, Physical review letters **66** (18), 2297 (1991).

[110] G. D'Amico, G. Rosi, S. Zhan, L. Cacciapuoti, M. Fattori and G. Tino, *Canceling the gravity gradient phase shift in atom interferometry*, Physical review letters **119** (25), 253201 (2017).

[111] *IEEE Standard Definitions of Physical Quantities for Fundamental Frequency and Time Metrology—Random Instabilities*, IEEE Std Std 1139-2008 c1–35 (2009).

[112] *IEEE Standard Specification Format Guide and Test Procedure for Single-Axis Interferometric Fiber Optic Gyros*, IEEE Std 952-1997 1–84 (1998).

[113] N. El-Sheimy, H. Hou and X. Niu, *Analysis and Modeling of Inertial Sensors Using Allan Variance*, IEEE Transactions on Instrumentation and Measurement **57** (1), 140–149 (2008).

[114] A. Yariv, *Optical electronics*. Saunders College Publ., 4 Ed. (1991).

[115] *DEVICE SPECIFICATIONS,NI 6366,X Series Data Acquisition 2 MS/s/ch, 8 AI, 24 DIO, 2 AO* (2016).

[116] J. Vovrosh, L. Earl, H. Thomas, J. Winch, B. Stray, K. Ridley, M. Langlois, K. Bongs and M. Holynski, *Reduction of background scattered light in vacuum systems for cold atoms experiments*, AIP Advances **10** (10), 105125 (2020).

[117] M. Ben Dahan, E. Peik, J. Reichel, Y. Castin and C. Salomon, *Bloch Oscillations of Atoms in an Optical Potential*, Phys. Rev. Lett. **76** 4508–4511 (1996).

[118] M. Langlois, *Conception et réalisation d'un gradiomètre de gravité à atomes froids.* Theses, Université Pierre et Marie Curie - Paris VI (2017).

[119] M. Taborelli, *Cleaning and surface properties*, in *CERN Accelerator School, vacuum in accelerators, Platja d'Aro, Spain, 16-24 May 2006*, 321–340 (2006).

[120] Epoxt Technology, *EPO-TEK® H21D Technical Data Sheet* (2019). Rev. 8.

[121] Thorlabs, *Right Angle Prism Mirrors, Dielectric Coatings: Specs.* `https://www.thorlabs.com/newgrouppage9.cfm?objectgroup_id=7715`. Accessed: 2020-06-15.

Acknowledgements

There are many people to thank, as research is, by its very nature, collaborative. Many will go unnamed, but by reading this, you are to be thanked. Most of all I would like to thank my supervisor Mike for your support and encouragement in every aspect, I really couldn't have done it without your consistent help over the past four years. Additional thanks Kai for taking a chance on me in the first place, who knew a speculative email would have ended here.

The Gravity Imager project has had many members and sub-teams over its course, so I would like to thank Artur, Alex, Anthony, Andrew for guiding me through the first few years - at the time it seemed impossible to believe that an entire system could be understood. A particular thank you to Aisha, without whom I don't think I would have anywhere near the technical knowledge I do now. As the project evolved, the new team emerged. Jamie and Jonathan, your unwavering support and friendship over the last two years can never be repaid, I leave SADE in your capable hands. Special thanks to Steve and the Physics Workshop staff, who put up with our ridiculous time frames and designs.

All members of the interferometry team, past and present, which there are too many to name, are to be thanked for your help and for creating four years of memories. For your friendship and help during one of the worst

years of my life, I have to name my brothers in arms Rustin and Jonny, whose support always appeared in times of need.

To all my friends and family, near and far, I thank you for your patience and love. I know a lot has been sacrificed in order for me to get to this point, so I hope that we can begin to make up for some of the lost time. Mum, Dad, and Luke, thank you for your support in all forms over this difficult time.

Finally, to my dearest Hannah, thank you for your unwavering love and support. This hasn't been easy on either of us, but it has helped to build a strong foundation. As this thesis finishes, I look forward to our next chapter.