

A Light Pulse Atom Interferometric Gravimeter in Dynamic Environments

By

Aaron Baumgart

B.S., University of Massachusetts Lowell, 2017

M.Sc., Brown University, 2021

Thesis

Submitted in partial fulfillment of the requirements for the Degree of
Doctor of Philosophy in the Department of Physics at Brown University

PROVIDENCE, RHODE ISLAND

October 2025

© Copyright 2025 Aaron Baumgart, All rights reserved. The author hereby grants to Brown University and The Charles Stark Draper Laboratory, Inc. permission to reproduce and to distribute publicly paper and electronic copies of this thesis document in whole or in part in any medium now known or hereafter created.

This dissertation by Aaron Baumgart is accepted in its present form by the Department of Physics as satisfying the dissertation requirement for the degree of Doctor of Philosophy.

Date _____
Vesna Mitrović, Advisor

Date _____
Richard Stoner, Advisor

Recommended to the Graduate Council

Date _____
Aaron Hui, Reader

Date _____
Kemp Plumb, Reader

Approved by the Graduate Council

Date _____
Thomas A. Lewis, Dean of the Graduate School

Curriculum Vitae

EDUCATION

Brown University Sc.M. in Physics, Ph.D. in Physics (Expecting) — GPA: 3.67/4.00	Providence, RI 2019–2025
Massachusetts Institute of Technology 8.311 Electromagnetism; 8.321 Quantum Theory I — GPA: 4.5/5.0	Cambridge, MA 2018–2019
University of Massachusetts Lowell B.S. in Physics — GPA: 3.85/4.00 (summa cum laude)	Lowell, MA 2013–2017

EXPERIENCE

Draper Laboratory Scholar Atomic/Molecular Optics — Development of atom gravimeter using light pulse atom interferometry.	Cambridge, MA 2019–2025
Draper Laboratory Employee Electro/Atom-Optic Engineer — Measurement and test design of photonic devices. — Testing and data analysis for systems of high precision instruments.	Cambridge, MA 2017–2019
Draper Laboratory Intern Mechanical Engineering & Optical Physics Internship — Testing of optical devices and mechanical design of experimental setups.	Cambridge, MA Summers 2013–2017
University of Massachusetts Lowell Research Assistant in Nuclear Instrumentation — Design, construction, and operation of array of germanium detectors for gamma spectroscopy. — Analog data collection and digital data analysis of radioactive samples.	Lowell, MA 2014–2017
University of Massachusetts Lowell Teaching Assistant in Nuclear Instrumentation Lab	Lowell, MA 2015–2016

SKILLS

- **Programming:** Proficient in MATLAB, Mathematica, Python, LATEX.
- **Mechanical Design:** Proficient in Solidworks, ANSYS. Experience working with machine shop and 3D printing.
- **Lab Experience:** Hands on experience with fiber and free space optics, mechanical test design, atomic physics.

INTERESTS

My research interests include atomic/molecular optics and it's many applications to cutting edge research and technology. I particularly enjoy experimental work that demands the integration of many areas of expertise.

Acknowledgments

I would like to begin by thanking my advisor at Draper, Rick Stoner, for the time and dedication he provided throughout my time as a graduate student. In the past six years, I never once felt I was left without support – a privilege I have come to realize is exceedingly rare. More than continuously driving me to grow into the physicist I am today, Rick taught me that our line of work will always be exciting and rewarding. Spending weeks running an experiment at sea was an unforgettable experience, and I can't wait for what the future holds.

The opportunity to fully achieve my academic goals would never have been possible without the support of my advisor at Brown, Vesna Mitrovic. Her continued support in the department ensured I was able to pursue my work to the fullest whenever difficulties arose. I would also like to thank my committee members, Aaron Hui and Kemp Plumb, for taking the time to read and review my work. I greatly appreciate the feedback and discussion of the project.

I would be remiss if I did not thank the rest of my team at Draper who also provided continuous support throughout my time as a student. Notably Susannah, who somehow spent more time than she was able to discussing physics and the fun experiments we ran. Additionally, our core physics team included Russ and Steve – I could not have asked for a better group to be a part of. I look forward to seeing where the project goes from here.

Pursuing a graduate degree is more than learning physics or running an experiment,

but also the journey it takes to grow as a person along the way. I would not be who I am today without my friends who were always there for me – Akshay, Brandon, Calvin, Eric, Isabelle, Jae, Janie, Khing, Nicole, Tim, and far too many more to list. My time at Brown was foundational in discovering not only how to be a scientist, but also a person, and would not have been the same without each and every one I had the pleasure of getting to know. Thank you all for the good times and the emotional support over the years.

Finally, I would like to thank my family for their unwavering support throughout my life that led me to where I am today. I would never have had the chance to make it through an entire graduate degree without the continuous encouragement of my mother, father, brother, and stepmother. At last we can say we have a doctor in the family (and perhaps multiple before too long).

Thank you all.

Contents

Acknowledgments	v
1 Introduction	1
1.1 Inertial Navigation	1
1.2 Basic Gravimeter Principle	4
1.3 Light Pulse Atom Interferometry	6
1.4 Field Mobility	10
1.5 Thesis Contributions	12
2 Theory	15
2.1 Stimulated Raman Transitions	15
2.2 Interferometry Sequence for Inertial Sensing	20
2.3 Laser Cooling and Trapping	27
2.4 Cosensor Operation	32
3 Apparatus	37
3.1 Sense Head	39
3.1.1 Optical Alignment	49
3.2 Photonics Subsystem	50
3.2.1 AOD box	55
3.3 Electronics	57

3.3.1	RF Control & Modulation	60
3.4	Experimental Sequence and Data Acquisition	66
3.4.1	MOT	67
3.4.2	Polarization Gradient Cooling	68
3.4.3	State Preparation	68
3.4.4	Interferometry Sequence	69
3.4.5	Detection	69
4	Sensor Performance & Characterization	71
4.1	Basic Atom Optimizations	71
4.2	Atom-Cosensor Input Axis Alignment	78
4.3	Vibration Suppression	86
4.4	Measurement Stability	89
4.4.1	Short-Term Precision	99
4.5	Absolute Accuracy	100
4.6	Optics Degradation	107
5	Field Deployment	117
5.1	Mobile Gravity Measurement	117
5.1.1	Building-Wide Gravity Gradient Measurement	118
5.1.2	White Mountains Excursion	121
5.2	At-Sea Deployments	127
5.2.1	First Deployment – Gulf of Mexico	127
5.2.2	Second Deployment – Hawaii	133
6	Conclusion	139
6.1	Future Work	141

List of Figures

1.1	Energy level diagram of the Raman interaction.	7
1.2	The three-pulse LPAI sequence.	8
2.1	The Raman interaction and energy level diagram.	16
2.2	The three-pulse interferometer sequence.	21
2.3	Sensitivity function of the three-pulse LPAI sequence.	24
2.4	Frequency response of the gravimeter.	25
2.5	Bloch sphere depiction of a Ramsey two-pulse sequence.	27
2.6	Laser cooling diagram.	29
2.7	Quadrupole trapping scheme.	30
2.8	The ABCD phase dither.	36
3.1	Block diagram of the overall system.	38
3.2	Picture and solid model of the overall system.	39
3.3	Picture and solid model of the sense head.	40
3.4	Picture and solid model of the collimator plate.	41
3.5	Picture and solid model of the MOT chicane.	43
3.6	Picture and diagram of the Raman racetrack optics.	44
3.7	Vacuum components in the sense head.	47
3.8	Block diagram of the laser system.	51
3.9	Rubidium spectroscopy for laser frequency locking.	52

3.10	An example EOM modulation scheme.	53
3.11	Picture of the AOD box.	56
3.12	Block diagram of the electronics.	58
3.13	AOD RF chain block diagram.	61
3.14	EOM RF chain block diagram.	62
3.15	Level diagram and transitions.	65
3.16	The experiment cycle and light pulses.	66
4.1	Demonstration of Rabi oscillation.	73
4.2	Raman pulse shapes.	74
4.3	Raman detuning scan.	75
4.4	Phase scan of the LPAI sequence.	77
4.5	Cosensor-gravimeter IA alignment.	79
4.6	Tilt test for IA alignment.	81
4.7	Tilt test projections determine IA alignment.	82
4.8	Tilt test bias shift.	86
4.9	Measured transmissibility of the shock mounts.	87
4.10	Example of vibration testing at 18 Hz drive.	88
4.11	A typical Allan deviation in the lab.	91
4.12	Demonstration of $\bar{\phi}$ rejection.	93
4.13	Month long stability measurement.	94
4.14	AC Stark scan using fourth pulse.	96
4.15	Stability comparison at Hawaii excursion versus recent performance.	98
4.16	Absolute accuracy over the course of the experiment.	105
4.17	Recent finding of DI Raman transition.	105
4.18	Relative atom number over the duration of the experiment.	107
4.19	Optical mode out of the ν_2 fiber.	109
4.20	Cloudy surfaces of internal optics.	113

4.21	Degraded HR surfaces of vacuum optics.	114
4.22	Coated HR surfaces of replacement vacuum optics.	116
5.1	The system on the mobile cart.	118
5.2	Allan deviations representative of building-gradient test.	119
5.3	Building-wide gradient test result.	120
5.4	Map of the land-based gravity measurement excursion.	122
5.5	Support structure for the system inside the van.	123
5.6	White mountain excursion results.	125
5.7	White mountain excursion residuals.	126
5.8	The system inside the CONEX box.	128
5.9	The Point Sur docked in Gulfport, MS.	130
5.10	Inertial environment at sea aboard the Point Sur.	131
5.11	Allan deviations during the first at sea underway.	132
5.12	At-sea track taken by the Point Sur	133
5.13	Allan deviations at the time of the Gulf of Mexico vs. Hawaii excursions.	134
5.14	The HMNZS Aotearoa at sea.	135
5.15	At-sea track taken by the HMNZS Aotearoa	136
5.16	Inertial environment at sea aboard the HMNZS Aotearoa.	137

List of Tables

3.1	EOM RF modulation signals and their purposes.	63
4.1	Accuracy and phase shift error sources.	104

CHAPTER 1

Introduction

1.1 Inertial Navigation

Many land, sea, and air vessels such as surface ships, submarines, and airplanes need to keep track of their position and heading to be able to navigate effectively. Global positioning system (GPS) is often employed for such tracking, but is not always reliable. GPS relies on satellite communication to triangulate position, and, due to the resulting sub femto-watt signals at the Earth's surface, is subject to jamming and low update rates. Moreover, satellite communication in certain scenarios may be unavailable, such as submarine navigation below the surface of the water. To better track position and heading in situations where uninterrupted navigation accuracy is crucial, inertial navigation is often deployed. Inertial navigation relies on internal sensors to track changes in movement, and as such requires no communication with the outside world. Due to the direct coupling of the sensors to the vessel, small and fast changes in movement are readily detected at a resolution impossible for GPS to match.

An inertial navigation system (INS) is a complete set of algorithms that monitors sensor readings to determine position and heading over time [1]. Inertial navigation is a concept that has existed for a hundred years, and the demands for more accurate and reliable

measurements in increasingly small form-factors have increased with advances in sensor technologies [2, 3]. An INS depends on sensors that comprise the inertial measurement unit (IMU) to calculate movement over time. The most basic complete IMU is comprised of three orthogonal accelerometers that measure three-dimensional acceleration and three orthogonal gyroscopes that measure three-dimensional rotation rate. Position and heading is calculated by integrating over sensor readings continuously. Every real sensor has a measurement error, which manifests as a drift in the accurate determination of position and heading calculated by the INS. For example, white noise on an accelerometer is integrated to a random walk in velocity, and thus a position error. The longer the INS continues to run, the larger the drift in position and heading error is. More precise sensors mean less drift over time, so significant effort has been given to creating accurate sensors. The precision of sensors in an IMU is therefore the primary determination of how long an INS can accurately determine position and heading within a certain bound.

No matter how accurate the sensor, there will always be some drift in the INS solution, so a way to correct the error or determine exact position will eventually be required. Many methods exist to bound INS drift, a few of which will be discussed briefly here. At its core, every method of INS correction presented here works by constraining position and heading based on a different measurement. The most common method is to utilize GPS to periodically update location. Doing so has the advantages of the short-term accuracy and high bandwidth of the INS measurement while periodically removing drift via GPS determination of position. This of course requires satellite communication at some regular interval, which, for example, bounds the length of time at which a submarine could remain submerged. Many other situations exist in which it is important to navigate accurately in the absence of GPS. Indeed, a simple situation where GPS reception on a surface ship is interrupted could cause serious problems for onboard systems. Another method of navigation aid is to visually observe location based on surrounding features [4]. Celestial navigation [5] uses observed directions of astronomical bodies to determine

platform position, but requires clear view of the sky which is sometimes obscured by clouds.

All of the methods of correcting INS drift listed so far require externally sourced messages, or visual information of variable availability. In contrast, map-matching is a technique wherein a measured quantity that depends on position is compared to a map of that quantity (e.g., gravitation, geomagnetic fields). A magnetometer, for example, can be deployed to measure magnetic field, and variations in magnetic field can be compared to maps to narrow down possible platform locations within a region [6]. The INS then acts as the main source of navigation, determining exact position, while informing the general region in which the magnetometer measurement can be compared to look for drift in the position determination.

While magnetic map matching may be able to bound INS drift [7], it has a number of disadvantages. Interfering magnetic fields, either time-varying geomagnetic fields or fields produced by the platform being navigated, can obscure detection of the desired signals. Gravity likewise changes with geographic location due to local variations in mineral composition and crustal density at the scale of multiple mgal over kilometer-scale distances and cannot be shielded or corrupted. A gravimeter capable of operating on a moving platform to make highly accurate measurements of local gravity is therefore of great interest as a navigation aid for INS correction since it cannot be spoofed or jammed, is always available, and is totally passive [8, 9].

Atom-interferometry based “quantum gravimeters” have demonstrated great progress in recent years and offer a pathway towards a highly accurate and field-deployable sensor capable of aiding in navigation. Another application of the quantum gravimeter stems from its potential for long term accuracy. Classical sensors not only have errors constant in time (and thus amenable to calibration), but also error that can evolve over time (which in many cases is not amenable to calibration). If the error of sensors is changing, it can become difficult to track error bounds of an INS solution. The quantum gravimeter can

then be deployed as a calibration source for non-quantum classical sensors, and its mobile nature allows calibration in the field. The ability of the sensor to function accurately without need for outside correction for long periods of time is a primary advantage of the gravimeter over classical sensors. Stability and accuracy is one of the most compelling reasons to develop an atom gravimeter, and both concepts will be discussed in detail.

1.2 Basic Gravimeter Principle

The fundamental basis of gravimeter operation is rather simple – a proof mass is subject to a gravitational field, which subjects the proof mass to a force. In modern gravimeters, the force effect on the proof mass is typically measured in one of two ways. The first is to attach it to a supporting spring. The spring (or flexure) will tend to deflect proportionally to the magnitude of the applied force. Either the deflection can be measured to determine applied force, or a counter-force proportional to an applied field can be applied (and measured) to maintain the flexure at a fixed position. Either of these approaches will report measurements relative to a calibration reference. A spring- or flexure-based gravimeter is therefore referred to as a “relative” gravimeter. Relative gravimeters require periodic calibration against an absolute source to account for drifts in the sensor. Notably, the spring constant is subject to change over time due to effects such as relaxation of stresses in its material composition [10].

Another way to handle the proof mass is to simply let it free fall for a period of time. The proof mass is dropped from rest and its position is tracked as it falls. The proof mass is typically dropped in a vacuum to avoid air resistance effects and is measured by laser interferometry, which provides a highly accurate measure of position. Measurement of distance fallen, taken with an accurate measurement of the fall time, enables determination of average acceleration. A minimal measurement of average acceleration is yielded by a series of three discrete displacement measurements: if the second measurement is taken at a dwell time T after the first, and the third measurement another dwell time T after the

second, the average acceleration can be expressed as

$$\langle a \rangle = \frac{x_1 - 2x_2 + x_3}{T^2}. \quad (1.1)$$

A series of in-vacuum proof mass position measurements at known times will yield the same results regardless of the proof mass’s size or composition, and the resulting measurement can thus be regarded as absolute. The state of the art absolute gravimeter uses a falling corner cube as a proof mass [11, 12]. Accuracy down to single-digit- μg can be achieved by such devices [13], but requires a quiet and stable environment to produce a good measurement of gravity and a trained operator to deploy in the field. In recent years, atom gravimeters have demonstrated increasingly impressive performance, and have begun to rival the falling corner cube in accuracy [14, 15, 16].

The atom gravimeter operates with the same basic free-falling principle as the corner cube gravimeter. Instead of a macroscopic object, it uses individual atoms as the proof mass. Using an atom as a proof mass is advantageous over a macroscopic object because the physical properties of the atom do not evolve over time. The physical properties of the atom are well studied and characterized and are therefore a reliable reference with which to perform a measurement. The atom gravimeter is comprised of a vacuum chamber in which atoms (such as Rubidium-85) are captured then dropped and measured as they fall, which is achieved through laser beams that manipulate the atom’s state.

Conceptually, an absolute gravimeter is simply a device that measures a mass in free fall over a period of time to determine average specific force. Due to the equivalence principle, the effect of a gravitational field is indistinguishable from an acceleration. The gravimeter can therefore be considered a special type of accelerometer with a relatively low bandwidth and high accuracy for the purpose of measuring quasi-static gravitational fields. As stated, accelerations can couple into the measurement of gravitational field, and are therefore a challenge to a gravimeter operating in a dynamic field that may be subject

to platform motion and higher frequency vibration. Techniques employed to mitigate such effects will be discussed in sections 2.4 and 4.3.

1.3 Light Pulse Atom Interferometry

To measure the free fall of the atoms, a technique known as light pulse atom interferometry (LPAI) is used. Lasers manipulate the atom by separating the wavefunctions into a coherent superposition of states and interfering them after some time. The fundamental mechanism through which LPAI is implemented is based on the early works of Raman [17], Rabi [18], and Ramsey [19]. It was discovered by Rabi that oscillation of a two-state quantum system could be driven by an external magnetic field, and Ramsey presented a method of improving the resolution of a measurement of the oscillation by separating two pulses that together produced Rabi oscillations. Raman discovered a type of inelastic scattering that is used as the fundamental interaction with the atom's states.

In this experiment, a two-photon process known as stimulated Raman scattering [20] is used to create coherent superpositions of the two ground state hyperfine levels in the Rb-85 atom. Pictured in figure 1.1, the interaction involves the transition from a ground state $|g\rangle$ or excited state $|e\rangle$, both in the same hyperfine grouping, to an upper intermediary state $|i\rangle$. A beam of frequency equal to the $|g\rangle \rightarrow |i\rangle$ transition is introduced along with a counterpropagating beam of frequency equal to the $|e\rangle \rightarrow |i\rangle$ transition. Each beam will independently induce stimulated absorption and emission to and from the corresponding states, but spontaneous emission will also be present when the atom is in the upper state $|i\rangle$. To suppress spontaneous emission and prevent a single photon interaction by the individual beams, the frequencies of both beams are detuned by the same amount Δ . The individual beams will now largely be unable to act on the atom alone, and the dominant process will be a two-photon interaction in which the state is temporarily excited from $|g\rangle$ to a virtual level detuned by Δ from $|i\rangle$ and immediately transferred via stimulated emission to $|e\rangle$, or vice versa. While both beams are incident on the atom, the state will

oscillate back and forth between the two lower levels, known as Rabi oscillation. The atom can then be considered a quasi two-state system in which the Raman beams transfer the atom between the two states $|g\rangle$ and $|e\rangle$. A more detailed theory of the interaction is presented in section 2.1.




Figure 1.1: Energy level diagram of the Raman interaction. Relative separations are not drawn to scale.

The Raman interaction provides a convenient and highly controllable method with which to manipulate the quantum state of the atom. This allows a realistic implementation of matter-wave interferometry. While the knowledge of de-Broglie matter waves has existed for a century, successful interference of neutral atoms was only demonstrated for the first time in 1991 by four independent groups [21, 22, 23, 24]. While all four methods were successful in performing atom interferometry, the one that emerged as the most practical method of atom interferometry was the use of light pulses by Kasevich and Chu. A key distinction of LPAI is the fact that its operation does not require spatial resolution of interferometer output momentum states (as do aperture or grating-based approaches). Rather, LPAI exploits entanglement of momentum with distinct internal states [25], making it practical in vibration-heavy environments. In other words, the momentum transfer of the Raman interaction (e.g. $|g, 0\rangle \rightarrow |e, \hbar k_{eff}\rangle$) encodes spatial

information of the atom within the Raman fields on the internal wavefunction of the atom.

The following year, Kasevich and Chu demonstrated the use of LPAI to measure gravitational acceleration [26]. The Raman interaction is used to implement a three-pulse “ $\pi/2 - \pi - \pi/2$ ” sequence that forms the atom interferometer. The interferometer will be described in more detail in section 2.2. The atom begins initially falling in a prepared state, say $|g\rangle$, then a “ $\pi/2$ -pulse” transfers the atom into a coherent superposition of equal probabilities of $|g\rangle$ and $|e\rangle$. The atom is allowed to fall for some dwell time, T , then a “ π -pulse” flips the states (i.e. $|g\rangle \rightarrow |e\rangle$ or vice versa) and, after another dwell time T , a final $\pi/2$ -pulse interferes the wavepackets. This three-pulse sequence forms a Mach-Zehnder interferometer analogous to an optical Mach-Zehnder interferometer. In the LPAI case the optical beams act as splitters and mirrors for the atomic wavefunction, with the $\pi/2$ -pulses acting as atomic beam splitters, and the π -pulse acting as an atomic mirror.




Figure 1.2: The three-pulse LPAI sequence measures average acceleration by imprinting a position-dependent laser phase on the atom at each evenly-spaced pulse.

The Raman beams can be thought of in a simplistic sense as a ruler against which the position of the atom is measured. A Raman pulse imprints a spatially-dependent phase on the atom’s wavefunction that corresponds to the position of the atom along the beam.

The resulting interferometer phase at the end of the three-pulse sequence determines the probability of the atom to be in one state or the other and is a function of the imprinted laser phase along the track:

$$\Delta\phi = \phi_1 - 2\phi_2 + \phi_3 \quad (1.2)$$

$$= k_{eff} (z_1 - 2z_2 + z_3) \quad (1.3)$$

$$= k_{eff} T^2 \langle a \rangle, \quad (1.4)$$

where k_{eff} is the effective wavenumber of the Raman interaction and $\langle a \rangle$ is the average acceleration along the direction of the beams during the interferometer sequence. The LPAI measurement thus produces a highly sensitive measurement of specific force, the benefits of which will be discussed in the following section.

LPAI's sensitivity advantage is essentially the same as that of the falling corner cube gravimeter using laser Michelson interferometry to measure proof mass motion. The precision of the laser wavelength is the basis for the precision of the distance measurements. Laser wavelength precision is achieved through highly accurate control of laser frequency, which is stabilized by comparison to atomic/molecular resonances. Interferometric distance measurement thus exploits frequency control to measure distance via accurate knowledge of the speed of light.

While the application of this work is focused on gravity-aided navigation, it is worth briefly discussing a few of the other numerous applications of atom interferometry. One of the most successful to-date is the atomic clock, which uses what is effectively a two-pulse atom interferometer sequence to measure the clock frequency of the hyperfine structure of an atom such as cesium or rubidium. This forms an incredibly stable reference to which other clocks can be synchronized [27]. Atom-based clocks have surpassed all other time-keeping devices in accuracy [28], and continue to pave the way for ultra-precise time measurements [29].

Gravimeters have uses beyond navigation aids, and are an important tool in the exploration of fundamental physics. The measurement of fundamental constants is important in experimentally probing theories of the universe. The least-accurately known constant is the gravitational constant G , and atom gravimeters offer accurate measurements of its value [30, 31, 32]. The gravimeter is also a natural probe of the equivalence principle, and can be used to compare the universality of free-fall of different atoms [33, 34, 35]. The highly sensitive nature of the measurement makes the atom gravimeter a good candidate for measurements of gravitational waves [36] and probes in the search for dark matter [37]. While LIGO has successfully detected gravitational waves [38], its bandwidth is limited to the 10s-100s of Hz range. LISA [39] is planned to be able to detect sub-Hz waves, which leaves an empty frequency band between the two. Atom gravimeters could probe in and beyond this range, effectively filling the gap in detection bandwidth. Planning and testing of such sensors are underway, with efforts in motion to construct a 100 m drop tower for gravitational wave detection [40, 41, 42]. Such gravimeters could also perform quite well in the vibration-free environment of space [36, 43].

Atom interferometry can also be used as other types of inertial sensors such as magnetometers [44, 45], gyroscopes, and gravity gradiometers. Atom-based gyroscopes and gradiometers are conceptually similar to the design of an atom gravimeter, but are typically more complex. Atom gyroscopes, for example, involve measuring the movement of two (or more) atom clouds to determine rotation about an axis [46, 47, 48]. Gravity gradiometers measure the difference in acceleration between two atom clouds to determine the gravity gradient along a defined axis [49, 50, 51].

1.4 Field Mobility

As mentioned in section 1.2, the current state-of-the-art absolute gravimeter, the falling corner cube, has a rather slow measurement rate and cannot operate in dynamic environments. It requires a substantial platform with setup and measurement time. The

macroscopic proof mass is dropped over ~ 10 cm and has a repetition period on the order of multiple seconds. Operation in dynamic environments is therefore a significant challenge that such devices are yet unable to overcome. Atom gravimeters began as lab-based sensors [52], requiring large benchtop setups for laser generation and control. Advances in photonics technology has greatly reduced the footprint of the required lasers and optics, and today the systems can be packaged into a much smaller volume than original lab based experiments, allowing for the potential of operation outside of the laboratory.

In the recent decade or two, increasing interest in atom gravimeters has been driven by the potential to package a highly accurate gravimeter into a smaller package capable of operating in dynamic environments [53, 54, 55]. Successful at-sea deployments of atom gravimeters have been demonstrated by a French team [56] and more recently by a Chinese team [57]. The French team has also demonstrated airborne operation of an atom gravimeter and its performance at or beyond that of classical sensors aboard the flight [58, 59, 60].

The sense head itself requires very little space, with the atom cloud being on the order of $100\text{ }\mu\text{m}$ in diameter and typical drop times allowing for less than a millimeter of vertical movement. The drop time of the measurement is highly variable, meaning its sensitivity can be adjusted depending on the environment. With no macroscopic moving parts, it is also theoretically robust against shock and vibration that may be encountered in the field.

Central to the ability of the atom gravimeter to function outside of a laboratory environment is the size and robustness of the sense head. The atoms must be enclosed in a vacuum chamber, and the optics that deliver optical beams to the interaction region must remain rigidly aligned when subjected to large movement and vibration. Most laboratory LPAI setups utilize benchtop optics that are bulky and, while relatively easy to adjust and align, are often quite sensitive to vibrations. The sense head in this gravimeter is thus constructed using custom optics mounts that rigidly attach to a central region inside the vacuum chamber. The compact nature of the optics assembly allows for rigid alignment of

optics in a relatively small form factor. The full sense head enclosed in magnetic shields is a cylinder 10 inches in diameter by 15 inches tall. Further details of the construction of the sense head is described in chapter 3. The sense head is installed on a baseplate that supports packaged laser/photronics components and monitor electronics. Space was also allocated for an IMU, and the overall structure encloses a footprint of 2X3 feet and 3 feet tall.

Manipulation of the atoms requires fine control over the optical system and a fairly complex set of electronics is needed to accompany the system. Given the early prototype nature of the gravimeter, much of the electronics were constructed out of commercial off-the-shelf (COTS) evaluation boards packaged into groups of functional electronics such as an RF generation box and FPGA control box. The splayed-out nature of the electronics design allowed for initial development work, but requires an additional 19 inch electronics support rack 3.5 feet tall with a base of 2X3 feet. These components take up a significant amount of volume and additional power, but altogether still fit in a form factor capable of being transported and deployed on a ship for at sea testing. Much of the initial engineering work was focused on robustness of the system, rather than size, weight, and power (SWAP). Current and future engineering efforts to reduce the SWAP of the system, especially the electronics, is described in the conclusion. Nevertheless, the current set of electronics was able to be successfully deployed on multiple field excursions, both at sea and on land.

1.5 Thesis Contributions

This thesis covers the development of an atom gravimeter capable of operation in the field. The engineering work to assemble the system and integrate it into a deployable sensor is presented, followed by its operation in the field. Multiple deployments are demonstrated, including two at sea excursions and a land trip via van. The techniques required to overcome the challenging dynamic environments are described and direct

measurements of spatially varying gravity are presented. In particular, the implementation and optimization of feedforward and feedback algorithms enable significantly increased gravimeter sensitivity when subject to dynamically varying specific force. My particular contributions are in the construction, testing, and operation of the gravimeter. This includes, but is not limited to, the development of the RF system described in section 3.3.1, the optimization of the atom interferometer described in section 4.1, and the deployment and operation of the system in the field described in chapter 5.

The other large focus of this thesis is the ongoing work to improve stability and accuracy, which generally involves the study of optical and atomic physical effects that can induce phase shifts in the atom interferometer. Various sources of phase shifts are characterized by experimental measurements of effects such as dependence of optical beam intensity. They are matched to theoretical simulations of effects such as AC Stark shifts, which, once understood, can be eliminated either by removing the mechanism through which it produces an LPAI phase shift or by algorithmically accounting for the effect with engineering solutions such as optical intensity locks and thermal control of various system components. I played a large role in the evaluation and analysis through experimental testing of many of these effects described in sections 4.4 and 4.5. The study of interferometer phase shifts is essential to the reliable performance of an atom gravimeter both in the field and in the lab.

Finally, significant effort has been focused towards the compactification of electronics, optics, and mechanical structure to enable transportation and operation in the field. The electronics are composed of splayed-out groups of commercial components packaged into chassis that are stored either on the baseplate with the gravimeter or in the accompanying electronics rack. Optics are packaged into two boxes that are installed on the baseplate by the gravimeter. The resulting system is mobile and self contained, capable of executing an experiment cycle that captures atoms and performs LPAI to ultimately report an accurate value of specific force at a reasonably fast repetition rate. The system, while

mobile, consumes a significantly larger volume than is necessary for a more integrated design. Efforts to compactify the electronics and the resulting improvements to SWAP and also signal cleanliness are discussed.

CHAPTER 2

Theory

At the basis of the experiment is the manipulation of the Rb-85 atom using lasers to induce transitions within the atom's hyperfine structure. Since Rubidium is an alkali metal, manipulation of the atom's state is achieved by controlling the single valence electron in the atom's outer shell by driving electromagnetic fields. This section details the theory of the interaction, notably the two-photon Raman transition and how it is used to produce a gravity-sensitive measurement through interference of the wavefunction.

2.1 Stimulated Raman Transitions

The Raman interaction is a two-photon process that can transfer the atom between two states through an intermediary state. The two lower states, labeled ground $|g\rangle$ and excited $|e\rangle$ are hyperfine levels in $5^2S_{1/2}$ separated by the clock frequency, $\omega_{HF} = 3035.73$ MHz. The intermediary state $|i\rangle$ is in $5^2P_{3/2}$, which is 780.2 nm (384.23 THz) above the lower states. The optical beams are set up in a counter-propagating “racetrack” configuration (see chapter 3), with one beam directed vertically up and the other vertically down. One beam is set to near the $|g\rangle \rightarrow |i\rangle$ frequency and the other is set to near the $|e\rangle \rightarrow |i\rangle$ frequency. Both frequencies are red-detuned by the single photon detuning Δ such that

neither light is exactly resonant but the frequency difference between the two beams is exactly equal to the difference between $|g\rangle$ and $|e\rangle$ – denoted as ω_{HF} . The natural linewidth of the single-photon interaction is $\Gamma = 6$ MHz, so at $\Delta = 500$ MHz $\approx 80 \Gamma$, spontaneous emission is highly suppressed. Given the minimal population of the intermediary state, the system can be viewed as a quasi two-state system, with the Raman interaction serving to transfer the state between $|g\rangle$ and $|e\rangle$.

The Raman process works by the atom absorbing a photon, say $|g\rangle \rightarrow |i\rangle$, and emitting the other via stimulated emission. The photon is emitted in the direction of the other beam, which means it experiences a net change in momentum equal to the sum of the magnitude of momentum of both photons (rather than the difference):

$$\hbar \mathbf{k}_{eff} = \hbar(\mathbf{k}_1 - \mathbf{k}_2) \quad (2.1)$$

$$\hbar |\mathbf{k}_{eff}| = \hbar(k_1 + k_2) \approx 2\hbar k_1. \quad (2.2)$$

For Rb-85, the recoil velocity imparted by the Raman interaction is $v_r = \hbar k_{eff}/m \approx 12$ mm/s.




Figure 2.1: (a) Energy level diagram of the Raman interaction. Relative separations are not drawn to scale. (b) The momentum transfer in the two-photon interaction. The photon emitted via stimulated emission is in the opposite direction from the absorbed one, imparting a two-photon momentum kick.

To describe the interaction more formally, we construct the basic fundamental theory of the interaction of light with the atom in a Raman pulse. This prescription closely follows the work of [61, 62]. The Raman interaction has been studied extensively and has been described formally in many other works [63, 64, 26]. The analysis starts with the two electromagnetic plane waves incident on the atom

$$\mathbf{E}_1(t) = \mathbf{E}_{1,0} \cos(\phi_1 + \phi_{1,0}) \quad (2.3)$$

$$\mathbf{E}_2(t) = \mathbf{E}_{2,0} \cos(\phi_2 + \phi_{2,0}). \quad (2.4)$$

The amplitude of the waves $\mathbf{E}_{1,0}$ and $\mathbf{E}_{2,0}$ may vary in time but are assumed to do so slowly compared to the frequency of the light. $\phi_{1,0}$ and $\phi_{2,0}$ are arbitrary constant phase offsets, and the time-dependent phase components are given by

$$\phi_1(t) = \int_0^t [\omega_1(t') - \mathbf{k}_1 \cdot \dot{\mathbf{z}}(t')] dt' \quad (2.5)$$

$$\phi_2(t) = \int_0^t [\omega_2(t') - \mathbf{k}_2 \cdot \dot{\mathbf{z}}(t')] dt' - \frac{\hbar t}{m} \mathbf{k}_2 \cdot \mathbf{k}_{eff}, \quad (2.6)$$

where the $\dot{\mathbf{z}}$ term captures the Doppler shift of the atom in $|g\rangle$. The term outside the integrand in equation 2.6 is from the recoil velocity of the atom in state $|e\rangle$ given the assumption that $\mathbf{E}_1$ exclusively couples the states $|g\rangle$ and $|i\rangle$ and $\mathbf{E}_2$ exclusively couples the states $|e\rangle$ and $|i\rangle$. It is assumed that there is no other coupling of states by the fields (e.g. $\mathbf{E}_1$ does not couple $|e\rangle$ and $|i\rangle$). The Schrodinger equation is $i\hbar \partial \Psi / \partial t = (H_0 + V) \Psi$, where the electromagnetic interaction is described by $V = q\mathbf{r} \cdot \mathbf{E}(t)$. The Schrodinger equation then takes the form

$$i\hbar \begin{bmatrix} \dot{\Psi}_e \\ \dot{\Psi}_i \\ \dot{\Psi}_g \end{bmatrix} = \frac{\hbar}{2} \begin{bmatrix} 2\omega_e & \Omega_2 e^{i(\phi_2 + \phi_{2,0})} & 0 \\ \Omega_2^* e^{-i(\phi_2 + \phi_{2,0})} & 2\omega_i & \Omega_1^* e^{-i(\phi_1 + \phi_{1,0})} \\ 0 & \Omega_1 e^{i(\phi_1 + \phi_{1,0})} & 2\omega_g \end{bmatrix} \begin{bmatrix} \Psi_e \\ \Psi_i \\ \Psi_g \end{bmatrix} \quad (2.7)$$

where $\Omega_1(t) = \langle i | q\mathbf{r} | g \rangle \cdot \mathbf{E}_{1,0} / \hbar$ and $\Omega_2(t) = \langle i | q\mathbf{r} | e \rangle \cdot \mathbf{E}_{2,0} / \hbar$ are the complex single-photon Rabi rates that govern the transition probability between $|i\rangle$ and $|g\rangle$ or $|e\rangle$ respectively. To attempt to solve the equation, first shift the energy origin by $-1(\omega_e + \omega_g)\hbar/2$. Then introduce the unitary transformation $\Psi = U\mathbf{c}$:

$$\begin{bmatrix} \Psi_e \\ \Psi_i \\ \Psi_g \end{bmatrix} = \begin{bmatrix} e^{-i(\phi_1 - \phi_2)/2} & 0 & 0 \\ 0 & e^{-i(\phi_1 + \phi_2)/2} & 0 \\ 0 & 0 & e^{i(\phi_1 - \phi_2)/2} \end{bmatrix} \begin{bmatrix} c_e \\ c_i \\ c_g \end{bmatrix}. \quad (2.8)$$

The laser phase offsets $\phi_{1,0}$ and $\phi_{2,0}$ can be absorbed into the complex Rabi rates. After applying the transformation, equation 2.8 then takes the form

$$\begin{bmatrix} \dot{c}_e \\ \dot{c}_i \\ \dot{c}_g \end{bmatrix} = -\frac{i}{2} \begin{bmatrix} -\delta & \Omega_2 & 0 \\ \Omega_2^* & -2\Delta & \Omega_1^* \\ 0 & \Omega_1 & \delta \end{bmatrix} \begin{bmatrix} c_e \\ c_i \\ c_g \end{bmatrix} \quad (2.9)$$

where the two-photon detuning is defined as

$$\delta = (\dot{\phi}_1 - \dot{\phi}_2) - (\omega_e - \omega_g) \quad (2.10)$$

and the single-photon detuning is defined as

$$2\Delta = -2\omega_i + (\dot{\phi}_1 + \dot{\phi}_2) + (\omega_e + \omega_g). \quad (2.11)$$

The hyperfine splitting frequency between $|g\rangle$ and $|e\rangle$ is $\omega_{HF} = \omega_e - \omega_g$, and the time derivatives of $\dot{\phi}_1$ and $\dot{\phi}_2$ are determined from equations 2.5 and 2.6. The resulting

two-photon and single-photon detunings then take the form

$$\delta = (\omega_1 - \omega_2) - \left[\omega_{HF} + \frac{\hbar k_{eff}^2}{2m} + \mathbf{k}_{eff} \cdot \dot{\mathbf{z}} \right] \quad (2.12)$$

$$2\Delta = (\omega_1 + \omega_2) - \left[2\omega_i - (\omega_e + \omega_g) - \frac{\hbar k_{eff}^2}{2m} + (\mathbf{k}_1 + \mathbf{k}_2) \cdot \dot{\mathbf{z}} \right]. \quad (2.13)$$

Note that terms with k_1^2 and k_2^2 have been dropped from the above expressions by taking the rotating wave approximation. Also, the vector sum $(\mathbf{k}_1 + \mathbf{k}_2)$ is nearly zero, and the two-photon interaction is therefore largely unaffected by Doppler shift.

In the limit where $2\Delta \ll |\Omega_1||\Omega_2|$, the time derivative of the intermediary state $\dot{c}_i$ can be neglected. c_i can then be eliminated from the system of equations by substitution, which reduces the system to a two-state problem:

$$\begin{bmatrix} \dot{c}_e \\ \dot{c}_g \end{bmatrix} = -\frac{i}{2} \begin{bmatrix} -\delta + \frac{\Omega_2^2}{2\Delta} & \frac{\Omega_2\Omega_1^*}{2\Delta} \\ \frac{\Omega_1\Omega_2^*}{2\Delta} & \delta + \frac{\Omega_1^2}{2\Delta} \end{bmatrix} \begin{bmatrix} c_e \\ c_g \end{bmatrix} \quad (2.14)$$

Now shift the energy origin by $-(\Omega_2^2 + \Omega_1^2)/4\Delta$ and define the Rabi rate

$$\Omega_{eff} = \frac{\Omega_1\Omega_2^*}{2\Delta} = \frac{|\Omega_1||\Omega_2|}{2\Delta} e^{i\phi} \quad (2.15)$$

and the AC Stark shift

$$\delta_{AC} = \frac{|\Omega_2|^2 - |\Omega_1|^2}{4\Delta}. \quad (2.16)$$

Equation 2.14 then takes the form

$$\begin{bmatrix} \dot{c}_e \\ \dot{c}_g \end{bmatrix} = -\frac{i}{2} \begin{bmatrix} \delta_{AC} - \delta & \Omega_{eff} e^{i\phi} \\ \Omega_{eff} e^{-i\phi} & -(\delta_{AC} - \delta) \end{bmatrix} \begin{bmatrix} c_e \\ c_g \end{bmatrix}. \quad (2.17)$$

An approximate solution to equation 2.17 can be written in the form $\mathbf{c}(t) = e^{Ht}\mathbf{c}(0)$. The resulting matrix solution takes the form

$$\begin{bmatrix} c_e(t) \\ c_g(t) \end{bmatrix} = \begin{bmatrix} C^* & -iS^* \\ -iS & C \end{bmatrix} \begin{bmatrix} c_e(0) \\ c_g(0) \end{bmatrix} \quad (2.18)$$

where C and S are

$$C = \cos\left(\frac{\Omega_{gen}t}{2}\right) + i\frac{\delta_{AC} - \delta}{\Omega_{gen}} \sin\left(\frac{\Omega_{gen}t}{2}\right) \quad (2.19)$$

$$S = \frac{\Omega_{eff}}{\Omega_{gen}} \sin\left(\frac{\Omega_{gen}t}{2}\right) e^{i\phi}. \quad (2.20)$$

The generalized two-photon Rabi rate Ω_{gen} is defined as

$$\Omega_{gen} = \sqrt{\Omega_{eff}^2 + (\delta_{AC} - \delta)^2}. \quad (2.21)$$

The solution now demonstrates a quasi two-state system that oscillates sinusoidally between the two states in the presence of the Raman fields. Oscillation occurs at the generalized Rabi rate. The amplitude of this oscillation is maximized when $\delta_{AC} = \delta$ and the generalized Rabi rate is also minimized with this condition.

2.2 Interferometry Sequence for Inertial Sensing

The atom interferometer sequence is constructed by a series of three Raman pulses described by the above theory. From equations 2.19 and 2.20, it is evident that the state will oscillate between $|g\rangle$ and $|e\rangle$ in the presence of the Raman beams at a rate of Ω_{gen} . If the atom is initially prepared in one state, say $|g\rangle$, a “ $\pi/2$ pulse” of length $\tau_{\pi/2} = \pi/2\Omega_{gen}$ will place the wavefunction into an equal superposition of the two states. The portion of the wavefunction that represents state $|e\rangle$ will receive a momentum kick of $\hbar k_{eff}$, which spatially separates the two states (as can be seen in figure 2.2). The de Broglie wavelength of Rb-85 is determined by the velocity of the atom, which is roughly 30 mm/s at 10 μ K. The resulting coherence length is $\lambda_B = h/p = 10$ nm, and given a recoil velocity

of 12 mm/s, the wave packets corresponding to the two internal states are separated far beyond the coherence length. After a drop time T on the order of a few milliseconds, a “ π pulse” of length $\tau_{\pi/2} = \pi/\Omega_{gen}$ reflects the wave packets back towards each other. After another drop time T (identical to the first drop time), a final $\pi/2$ pulse then interferes the wave packets as they overlap, producing an interferometric phase. The phase imprinted on the wavefunction by a single Raman pulse contains a spatially-dependent term

$$\phi^{g,e} = \pm \mathbf{k}_{eff} \cdot \mathbf{z} \quad (2.22)$$

where g (e) denotes the transfer of the state $|g, 0\rangle \rightarrow |e, \hbar k_{eff}\rangle$ ($|e, 0\rangle \rightarrow |g, -\hbar k_{eff}\rangle$). Note the counterpropagating Raman beams impart a momentum transfer that carries a sign depending on whether the state is transferred from $|g\rangle \rightarrow |e\rangle$ or $|e\rangle \rightarrow |g\rangle$.




Figure 2.2: The three-pulse interferometer sequence. Average acceleration a has been subtracted from the position to illustrate the Mach-Zehnder interference scheme.

The three-pulse interferometer sequence forms a Mach-Zehnder interferometer, and the phase is determined by the path difference between the two arms:

$$\Delta\phi_g = (\phi_1^g + \phi_2^e + 0) - (0 + \phi_2^g + \phi_3^e) \quad (2.23)$$

$$= |\phi_1^g| - (|\phi_2^e| + |\phi_2^g|) + |\phi_3^e| \quad (2.24)$$

where $\phi^{g,e}$ is specified because of the directional dependence and spatial separation due to the momentum transfer of the Raman interaction. Equation 2.24 collects the terms from the interferometer and explicitly includes the sign of each term. With $T = 5$ ms, the spatial separation between the two wavepackets at the time of the second pulse is roughly 60 μm . Given equation 2.22, the phases represent the positions of the wavepackets at each point in equation 2.24, which results in the expression

$$\Delta\phi = k_{\text{eff}} [z_1^l - (z_2^u + z_2^l) + z_3^u] \quad (2.25)$$

assuming $\mathbf{k}_{\text{eff}}$ is aligned with $\hat{\mathbf{z}}$. To reconcile the upper (u) and lower (l) positions of the wavepackets at the second pulse, consider the classical trajectory of the atom. The deviation of the wavepacket from the classical position is due to a momentum kick $\hbar k_{\text{eff}}$. The lower position z_i^l matches that of the classical trajectory, while the upper position z_i^u deviates by $\hbar k_{\text{eff}} T/m$. These constant momentum terms can be seen to cancel by plugging $z_i^u = z_i^l + \hbar k_{\text{eff}} T/m$ in equation 2.25, resulting in the simple expression

$$\Delta\phi = k_{\text{eff}} (z_1 - 2z_2 + z_3) \quad (2.26)$$

where z_i is the classical position of the atom, or the position of the wavepackets without receiving a momentum kick. Equation 2.26 can readily be seen to provide a measure of average acceleration, which results in the overall measure of acceleration by the interferometer

$$\langle a \rangle = \frac{z_1 - 2z_2 + z_3}{T^2} = \frac{\Delta\phi}{k_{\text{eff}} T^2}. \quad (2.27)$$

The above theory of LPAI has been thoroughly described [65, 66, 67], and in reality there are many more phase shifts that arise in the three-pulse interferometer sequence [68, 69]. Some of these phase shifts will be discussed in the efforts to mitigate non-inertial or generally unintended sources of error in the LPAI sequence in section 4.5.

It should be emphasized that the interferometer interaction is between wavepackets derived from a single atom and the theory described above can be experimentally implemented with only a single atom [70]. To achieve practically useful signals, it is necessary to have many atoms in a cloud that can be simultaneously measured. To measure the phase of the interferometer, the state of the atom is detected at the end of the sequence. When measured, the wavefunction collapses into one of the two states ($|g\rangle$ or $|e\rangle$), so a single atom alone cannot provide a measure of the phase in a single measurement. Instead, many atoms are measured simultaneously, and the resulting detection of their states after the interferometer provides a statistical measurement of the probability of the atom to be in a certain state, thus measuring the phase of the interferometer. Despite the presence of multiple atoms, the interferometer acts independently on each atom and one does not influence any other. This is further emphasized by the fact that the atom cloud is in a very dilute gaseous form in this experiment despite the μK -scale mean temperature. Atom-atom interactions are small, higher-order effects on interferometer accuracy.

Thus far, the interferometer has been described as measuring the “average” acceleration. It is important to note that the acceleration measured by the gravimeter is a weighted average of the instantaneous acceleration during the interferometer sequence, which can be characterized by a sensitivity function [71]. The temporal sensitivity of the interferometer to instantaneous acceleration can be calculated by integrating over a sudden spike in acceleration $a(t) = \delta(t' - t)$ to determine the perturbation of the atom position at each Raman pulse:

$$\langle a \rangle = 0 - 2 \int_0^T \int_0^T \delta(t' - t) dt' + \int_0^{2T} \int_0^{2T} \delta(t' - t) dt' \quad (2.28)$$

where $t = 0$ is defined as the first pulse. The result is a triangular sensitivity function

centered on the second pulse (figure 2.3).

$$S_g(t) = \frac{1}{T} \begin{cases} t, & \text{if } 0 < t < T \\ 2T - t, & \text{if } T < t < 2T. \end{cases} \quad (2.29)$$

Intuitively this makes sense – a sudden acceleration at the beginning of the interferometer will produce close to a common mode shift and therefore have little effect on the phase. Similarly, a sudden acceleration at the end of the interferometer will produce very little shift in the final position. An acceleration closer to the center will produce the greatest phase shift between the pulses.




Figure 2.3: The sensitivity function of the three-pulse LPAI sequence describes the temporal response of the gravimeter to instantaneous specific force.

Another important consideration in the measurement of average specific force is the frequency response. The gravimeter has a flat response to low-frequency content, but sees a roll off starting around $f_c = 1/3T$ (around 70 Hz for $T = 5$ ms). The response decays as $1/f^2$ beyond this point, with periodic sensitivities at integer multiples of $f = 1/T$ (figure 2.4). The response is given by

$$F(\omega) \sim \frac{1}{\omega^2} \sin^2(\omega T/2). \quad (2.30)$$




Figure 2.4: Frequency response of the gravimeter is flat at low frequencies and rolls off as ω^{-2} with insensitivities at every $f = n/T$.

It is worth noting that the system could be used as a clock by locking the laser frequency to the oscillation of the atom state due to Raman interaction. The simplest method would be to apply a single π -pulse and keep the frequency locked such that maximum population transfer is achieved. The frequency of the laser can then be related to the hyperfine frequency ω_{HF} , which defines the reference clock cycle. This simple approach would produce fairly poor precision, since the noise of the fluorescence signal can be large compared to a frequency-induced population transfer shift at the detuning peak of the Raman interaction (see figure 4.3), and an accuracy of only about $200 \text{ kHz}/3000 \text{ MHz} = 70 \text{ ppm}$ could realistically be achieved. A smarter method is to create a Ramsey interferometer sequence in which two $\pi/2$ -pulses separated by a dwell time T are used to produce interference. This separation acts to increase the effective interrogation time and thus the sensitivity of the measurement. The Ramsey sequence accumulates a phase

$$\phi_{\text{Ramsey}} = \delta T, \quad (2.31)$$

where δ is the detuning from ω_{HF} . The measured phase can be locked by adjusting the detuning δ , which serves as the basis of the clock reference. The sensitivity then increases as a function of longer dwell time T , just as sensitivity of the three-pulse interferometer

increases as a function of longer drop times.

The Raman interaction and resulting interferometer sequence can be pictured on a Bloch sphere to provide an intuitive picture of the sequence [61, 72]. Start by considering the atom as the quasi two-state system as described above and representing it in terms of a density matrix

$$\rho(t) = |\Psi\rangle \langle \Psi| = \begin{bmatrix} \rho_{ee} & \rho_{eg} \\ \rho_{ge} & \rho_{gg} \end{bmatrix} = \begin{bmatrix} |c_e|^2(t) & c_e(t)c_g^*(t) \\ c_e^*(t)c_g(t) & |c_g|^2(t) \end{bmatrix}. \quad (2.32)$$

Project the density matrix onto the Pauli spin matrix vector by taking the trace of the product

$$\mathbf{P}(t) = \text{Tr}[\rho(t) \boldsymbol{\sigma}] \quad (2.33)$$

$$= \begin{bmatrix} 2\text{Re}[\rho_{ge}] \\ 2\text{Im}[\rho_{ge}] \\ \rho_{ee} - \rho_{gg} \end{bmatrix}. \quad (2.34)$$

This form represents the population in the ground versus excited states in P_z , while P_x and P_y represent the coherence between the two states. The equation of motion for the density matrix then takes the form

$$\frac{d}{dt}\mathbf{P}(t) + \mathbf{P}(t) \times \boldsymbol{\Omega} = 0, \quad (2.35)$$

where we define the drive vector as

$$\boldsymbol{\Omega} = \Omega_{gen} \begin{bmatrix} \sin \theta \cos \phi \\ \sin \theta \sin \phi \\ \cos \theta \end{bmatrix}. \quad (2.36)$$

Here, the polar angle θ has been defined as

$$\cos \theta = \frac{\delta_{AC} - \delta}{\Omega_{gen}}, \quad \sin \theta = \frac{|\Omega_{eff}|}{\Omega_{gen}}. \quad (2.37)$$

With the above definition, a Bloch sphere picture can be used to represent the Raman interaction. The presence of the Raman field will rotate the state about a vector on the horizontal plane, while the state will continuously precess about the z-axis. An analogy can be drawn between this representation and nuclear magnetic resonance (NMR), in which the Raman field acts as a weak tipping field while the state will experience Larmor precession about the vertical due to a strong magnetic field. The three-pulse atom interferometer sequence is then analogous to a spin-echo experiment in nuclear magnetic resonance [73, 61].




Figure 2.5: Credit to [62]. Bloch sphere depiction of a Ramsey two-pulse sequence. The atom initially in the ground state points up, a $\pi/2$ pulse tilts the state into a superposition of the two states in the xy-plane. The state precesses about the z-axis during the dwell time, and the second $\pi/2$ pulse then tilts the state about the x-axis. The projection onto the z-axis provides a measure of the phase accumulated during the dwell time.

2.3 Laser Cooling and Trapping

The first step in the experiment is to capture Rubidium atoms, which is done via laser cooling and trapping to form a magneto-optic trap (MOT). This technique takes advantage of lasers with very stable frequency control to manipulate the internal and external state

of the atom via optical pumping. The first neutral-atom MOT was experimentally realized in 1986 [74], and laser cooling has since become a rich field that has been extensively explored. The basic theory will be presented here, and more exhaustive descriptions can be found in [75, 76, 77].

Laser cooling exploits the fact that when an atom absorbs a photon, it receives a momentum kick along the beam, whereas the subsequent spontaneous emission photon is emitted in a random direction. The result is a time-averaged force along the direction of the laser. Laser cooling also takes advantage of the Doppler effect. When an atom is moving at a particular speed relative to the laser beam it will see a shift in the frequency of the light due to the Doppler effect. The laser is set to a frequency a few linewidths below the transition energy from the lower $F = 3$ state to the upper $F' = 4$ state. The atom then will be less likely to interact with the light unless it is moving against the beam within a range of speeds that Doppler-shift the light closer to resonance with the transition. In this case, it will absorb the red-detuned photons and spontaneously emit photons of full energy in a random direction. The atom receives a small net loss of momentum from the interaction. One such event is a very small net change in momentum (and kinetic energy), but the atom will be subject to many such interactions as it travels against the beam. As the atom absorbs many red-detuned photons and emits them as full-wavelength photons in random directions, the average effect over many such interactions will be a net loss of momentum along the direction of the beam, thus slowing the atom in that particular direction.

This effect as described with a single beam slows atoms moving at sufficient velocity against the direction of the light. To slow atoms moving in the opposite direction, introduce a counterpropagating beam of the same red-detuned frequency. The atoms will then be slowed in either direction along the shared beam axis. This can then be extended to three dimensions by introducing three orthogonal pairs of counterpropagating, red-detuned beams, and a region in space where all six beams overlap is created that will slow the




Figure 2.6: A red-detuned cooling beam acts to slow the motion of an atom over many interactions with the atom that impart small random decreases in net momentum.

atoms in all directions. The velocity distribution of atoms at a given temperature is described by the Maxwell-Boltzmann distribution, and while most atoms initially around room temperature are likely moving too fast to interact in the manner described, there is enough population in the slow-moving tail of the distribution to slow an appreciable number of atoms over time. This forms an optical molasses, where atoms are slowed to near zero velocity.

While the atoms in an optical molasses have been slowed, they are not confined in space. To trap the atoms, a quadrupole magnetic field is introduced that is centered on the overlap region of the six beams. The field is zero at the center and nonzero elsewhere. Where there is field present, the magnetic sublevels of the F states are split due to the Zeeman effect. This splitting of sublevels shifts the energy levels such that the difference between some levels (e.g. $|F = 3, m = 0\rangle \rightarrow |F' = 4, m = +1\rangle$) is smaller and therefore closer to resonance with the red-detuned beams. This allows the light to interact more readily with the atoms and the beams will push atoms that are not in the magnetic center region.

As described, any of the beams could push the atoms, which would serve to spatially

spread out the atoms. To prevent atoms getting pushed away from the center, the polarization of the beams is chosen such that the counterpropagating pairs are oppositely-circularly polarized (i.e. pairs are σ_+ and σ_-). When chosen correctly, optical transition selection rules will prevent the atoms from getting pushed away from the center. For example, the energy levels of atoms to the left of the center will be split such that $|F = 3, m = 0\rangle \rightarrow |F' = 4, m = +1\rangle$ is closer to resonance and the right-pointing σ_+ beam will push atoms towards the center, while to the left $|F = 3, m = 0\rangle \rightarrow |F' = 4, m = -1\rangle$ is closer to resonance and the left-pointing σ_- beam will push atoms towards the center (see figure 2.7). Extended to three dimensions, the beams will now serve to form a trap centered on the magnetic field zero. Combined with the optical cooling effect, a magneto-optic trap (MOT) is formed. Note that the term “MOT” is used to refer to both the magnetic and optical system used to trap atoms as well as to the atom cloud itself.




Figure 2.7: Spatial trapping of atoms is achieved through introduction of a quadrupole field that allows circularly-polarized beams to push atoms towards the field center.

As mentioned, the cooling process involves random spontaneous emission of photons. While the average velocity in laser cooling is zero $\langle v \rangle = 0$, the average speed, or mean squared velocity is nonzero $\langle v^2 \rangle \neq 0$. This limits the temperature to which the atoms can be cooled, defined by the Doppler limit T_D at which point the interaction equally cools

and heats the atoms

$$T_D = \frac{\hbar\gamma}{2k_B}, \quad (2.38)$$

where γ is the natural linewidth of the transition. For the Rb-85 $S_{1/2} \rightarrow P_{3/2}$ transition, the natural linewidth is 6 MHz and the Doppler limit is 145 μ K. A surprising observation in early laser cooling experiments was the realization of temperatures below the Doppler limit [78, 79]. This sub-Doppler cooling was eventually seen to be due to a number of effects depending on the exact conditions of the MOT construction. In this system, sub-Doppler cooling is realized by a mechanism called polarization gradient cooling (PGC). The detailed theory of polarization gradient cooling is beyond the scope of this work, but involves the distribution of population into rotating groups of sublevels by optical pumping. A more detailed theory of PGC can be found in [80]. For the purpose of this work, PGC enables the cooling of atoms below the Doppler limit. The temperature achievable by PGC is simply the recoil limit of the atom-photon interaction

$$T_r = \frac{\hbar^2 k^2}{k_B m}, \quad (2.39)$$

which for Rb-85 $S_{1/2} \rightarrow P_{3/2}$ is 370 nK.

It is worth noting that more sophisticated techniques of producing a MOT have been developed. Two-dimensional MOTs are often used to feed a three-dimensional MOT or to produce a slow fountain of atoms [81]. MOTs have also been produced using a single beam reflecting off the inside of a conical or pyramidal mirror [82]. The reflections off the mirror produce the inward-pointing beams necessary to trap atoms at the intersection. In more recent years, surface gratings similarly have been used to produce a MOT [83, 84, 85]. Referred to as GMOTs, such constructions achieve atom trapping in an impressively compact form factor using only a single laser. While such methods have great promise for compact atom sensing, the atoms are trapped near to the mirror or grating surface. This presents a challenge when interrogating the atom, as the Raman beam ends up

clipping the surface, thus creating wavefront distortions. Such distortions negatively interfere with the LPAI measurement in a way that would be difficult to account for. For this reason, a simpler 3D MOT construction was chosen in which the atom position is not tied to a surface. Although surface distortion effects made the implementation of GMOTs impractical, future developments may open the door for such constructions. Very compact, even chip scale GMOTs have been realized [86, 87] which is highly desirable when considering size reduction of the system.

2.4 Cosensor Operation

The gravimeter provides a very stable and accurate measure of specific force but operates at a slow repetition rate of 10s of Hz. While 10s of Hz is fast by LPAI standards, it is slow for implementation in a navigation solution. Moreover, knowledge of which fringe of the interferometer a measurement lies on is necessary to perform an absolute measurement since the atom interferometer interaction produces phases $\gg 2\pi$. If the specific force (or projection of specific force on the input axis) is varying enough to produce an interferometer phase shift of π , there will be an ambiguity of the measurement – also referred to as fringe hopping. Given the scale factor equation $\phi = k_{eff}T^2a$, the specific force must be kept within a fringe, defined as a_π

$$a_\pi = \pi/k_{eff}/T^2. \quad (2.40)$$

For $T = 5 \text{ ms}$, $a_\pi = 800 \text{ } \mu\text{g}$.

A ship can easily see 5 degrees of tilt corresponding to $\pm 4 \text{ mg}$ of variation. In addition, vibration at higher frequency than the bandwidth of the gravimeter may exist. To ensure reliable measurements in this kind of dynamic environment, a high-bandwidth MEMS accelerometer, referred to as the *cosensor*, is introduced. The cosensor is less accurate than the gravimeter and has a significant bias drift over time but provides a coarse

measurement of specific force along the gravimeter axis. The higher bandwidth also allows a measurement of higher-frequency inertial input to the gravimeter. This fast but inaccurate measurement of specific force is provided to the gravimeter via *feedforward*.

Feedforward takes the cosensor's measurement of specific force and converts it to an effective interferometer phase. This phase is then subtracted in real time from the LPAI sequence, leaving only a residual phase equal to the difference between the cosensor and atom gravimeter's measurement of acceleration. Feedforward is implemented by integrating the cosensor's measurement during the interferometer and applying an update to the lasers' frequency and phase just before each pulse based on the shift calculated by the integrated cosensor. The frequency shift Δf is due to the Doppler effect, and is calculated as

$$\Delta f = \frac{\Delta v}{c} f_{eff} = \frac{k_{eff}}{2\pi} \Delta v, \quad (2.41)$$

where f_{eff} (k_{eff}) is the effective frequency (wavenumber) of the two-photon $S_{1/2} \rightarrow P_{3/2}$ transition and Δv is the relative velocity of the atoms versus the sense head as measured by the cosensor. When an atom is subject to a constant gravity of 9.8 m/s^2 , its Doppler shift is roughly 25 kHz/ms . The frequency shift can cause a change in Rabi rate, so it is updated just before each real optical pulse in the interferometer. The phase can be considered an accumulation along the entire interferometer and can be updated only at the final pulse. The phase accumulated is

$$\phi = \phi_0 + 2\pi \int_{t_1}^{t_1+T} \nu(t) dx \quad (2.42)$$

$$= \phi_0 + 2\pi \int_{t_1}^{t_1+T} \vec{k}_{eff} \cdot \dot{\vec{z}} dx \quad (2.43)$$

$$= \phi_0 + k_{eff} (z_1 - 2z_2 + z_3), \quad (2.44)$$

where ϕ_0 is the target phase, t_1 is the beginning of the interferometer, and $\nu(t)$ is the true

hyperfine frequency of the atom in the laboratory frame. The positions z_i are estimated by the integrated cosensor to determine the phase shift to apply to the atoms. The integration of the cosensor is implemented in such a way that the triangular sensitivity function described in section 2.2 is inherently accounted for. Note this is a simplified expression, in reality there are other effects such as integration over pulses of non-infinitesimal length and allowance for time to update the laser frequency. Some of these effects will be discussed in more detail in later sections, but the above expression describes the main application of feedforward.

The above phase is subtracted from the laser phase, and if the cosensor and gravimeter are perfectly aligned and calibrated, the residual phase will be zero. Deviations of this residual phase signifies a disagreement between the two sensors. The gravimeter’s measurement is assumed to be significantly more stable and accurate than the cosensor and as such is taken to be “truth.”

The residual phase after feedforward is applied is interpreted as an inaccuracy in the cosensor, namely variations in the bias and scale factor of the cosensor. The cosensor outputs an analog voltage V that is linearly dependent on acceleration and is converted to a specific force via its scale factor s . The scale factor is not fundamentally known and is instead determined experimentally. Similarly, there may be a DC offset from zero, referred to as a bias b (in units of specific force) which is also experimentally determined.

$$g_{cosensor} = sV + b \tag{2.45}$$

The scale factor and bias may vary over time due to a number of unknown factors such as temperature and various other internal or external sensitivities. Since the gravimeter provides a reliable measurement, there is no need to attempt to model these variations. Instead, the residual phase measured after feedforward has been applied can be converted into a delta in specific force and used to generate a correction to the cosensor’s bias and

scale factor:

$$\delta\phi(t) \rightarrow \delta g(t) \sim \delta s(t) V + \delta b(t). \quad (2.46)$$

The implementation of the scale factor and bias correction is referred to as *feedback*. Feedback is comprised of two independent Kalman filters: a phase-tracking algorithm and a cosensor-correction algorithm. The phase-tracking Kalman filter observes the atom population ratio as measured by the fluorescence detector, and determines a residual phase (population ratio is the readout of the atom interferometer, see section 4.1 for definition). Changes in the population ratio are not due solely to phase shifts in the interferometer, but also changes in the midpoint and contrast. Multiple sources such as atom temperature and Raman beam intensity effect midpoint and contrast, and variations over time, if not accounted for, would cause inaccuracies in the phase determination. To this end, the Kalman filter tracks midpoint, contrast, and phase of the interferometer. The Kalman filter assumes the population ratio S takes the form

$$S = M + C/2 \cos(\phi_0 + \phi_r) \quad (2.47)$$

where ϕ_0 is an applied phase dither and ϕ_r is the residual LPAI phase. In normal operation, the phase ϕ_r is adjusted to stay constant by the feedback algorithm, which means the phase-tracking filter does not have a way to distinguish between shifts in midpoint, contrast, and phase. To address this issue, a four-point phase dither is applied to the interferometer, with a cycle of one of the “ABCD” points being applied per experiment. The ABCD dither applies a phase of $\pm\pi/2 \pm \pi/6$, which produces a square wave centered around the middle of the fringe. This allows the Kalman filter to disambiguate midpoint, contrast, and phase, while avoiding reaching the extremes of the interferometer at which point sensitivity to phase shifts is zero.

Note that feedforward-subtraction phase does not show up here explicitly because it is subtracted from the laser phase in real time, before the feedback algorithm receives




Figure 2.8: The ABCD phase dither moves the interferometer around a fringe to allow for unambiguous tracking of phase, midpoint, and contrast.

the signal. Also note that contrast is defined as the peak-to-peak amplitude of the interferometer. Changes in midpoint and contrast are assumed to be slow, so the filter gain is set low such that the ability of parameters to significantly change is on the order of 10 seconds. The phase is also tracked similarly slowly, which permits tracking of long-term drifts while ignoring short term variations such as shot-to-shot noise and vibrations.

The introduction of the cosensor through feedforward and feedback forms a hybrid system with the gravimeter that benefits from the advantages of both sensors. The gravimeter is able to provide highly sensitive and accurate measurements, while the cosensor provides fast measurements at a high dynamic range. Together, the hybrid system is able to operate in dynamic environments while producing very sensitive measurements of specific force.

CHAPTER 3

Apparatus

The gravimeter system is comprised of a large set of mechanical structures and electronics that enable manipulation of the atoms to measure gravity. At the core of the structure is the sense head, a vacuum chamber where the optics are delivered and the LPAI measurement occurs. The sense head is mounted by three legs to a 20X23 inch aluminum baseplate. Also atop the baseplate is the photonics system that generates light and delivers it to the sense head via a set of fibers routed through a support structure. Local electronics collect analog signals and digitize them to be sent to the control system.

The baseplate is supported by three 12 inch tall legs that have shock mounts to suppress vibration. The shock mounts serve to dampen high frequency vibration that could corrupt the low frequency specific force measurement and protect the system from large mechanical shocks while in motion. Below the legs is a support plate for rigging and mounting to whatever platform is carrying the system. A thermal enclosure constructed out of an 80/20 extruded aluminum frame and insulating foam panels isolate the system from environmental temperature fluctuations as well as providing a degree of protection from debris.

A 19 inch rack carries most of the electronics that control the experiment. The




Figure 3.1: Block diagram of the overall system.

electronics are comprised mostly of COTS components assembled together as a first prototype. The components are grouped into discrete boxes of general functions. The FPGA box houses the FPGA and supporting signal routing. The RF box houses chains of components that generate RF signals that modulate and control the laser system. The thermal box contains PID controllers that stabilize the temperature of various subassemblies in the system. All the electronics are supplied with 24 V DC power by a power converter box. An uninterrupted power supply (UPS) provides power to the power box as well as any auxiliary components throughout the system. All system power routes through the UPS, allowing for short-term operation when external 120 V AC power is not supplied. The UPS also serves to avoid ground loops that might result from use of multiple electrical outlets.

The overall system is grouped into two large assemblies – the sense head support structure with a rectangular base of 24X35 inches by 35 inches tall, and the electronics rack with a footprint of 23X32 inches and 42 inches tall. The electronics rack can be mounted either on top of the thermal enclosure frame for a minimal overall footprint, or placed to the side when constrained by height. A cable bundle runs between the FPGA

box and the proximity electronics, and two SMA cables run from the RF box to the photonics box. A power strip provides power from the UPS to the photonics box, and the UPS receives 120 V AC from an external source.




Figure 3.2: Picture and solid model of the overall system. The main structure houses the sense head, optics, and proximity electronics while the rack carries the main set of control and RF electronics.

3.1 Sense Head

The LPAI measurement occurs at the sense head, a vacuum chamber enclosed by a set of magnetic shields. A set of laser beam delivery optics receives light from the

photonics subsystem and launches it into the vacuum chamber. The vacuum chamber encloses optics that guide light to the atoms, a Rubidium atom source, and vacuum components responsible for maintaining vacuum over time. The sense head is cylindrical, the inner vacuum area being approximately 5 inch in diameter by 7 inch tall, and the outer shields that envelop the entire assembly being 10 inch in diameter and 15 inch tall. The chamber is built primarily out of paramagnetic titanium to avoid magnetic field distortions, and components internal to the vacuum are carefully chosen to avoid contamination or outgassing in the ultra-high vacuum (UHV, $\sim 10^{-9}$ torr).




Figure 3.3: Picture and solid model of the sense head. The model is colored in functional groups of optics (red), vacuum (green), cosensor (blue), and shields (yellow).

Outside the vacuum, light is delivered to the sense head from the photonics/AOD subsystem by three fibers, referred to as ν_1 , ν_2 , and ν_3 . The fibers are received by an assembly of free space delivery optics mounted to the bottom of the sense head, referred to as the collimator plate (figure 3.4). The collimator plate is responsible for collimating the input light, cleaning up the polarization, and steering the beams into the vacuum chamber. The three fibers each have a discrete set of optics that individually handle the light. The ν_1 fiber carries both Raman beams and one of the state prep beams, and gets collimated into

a 7 mm $1/e^2$ diameter Gaussian beam. The Raman beams are delivered with orthogonal polarization through the fiber and follow an identical path into the vacuum chamber. Phase shifts between the Raman beams that could be induced by different spatial paths are removed by keeping everything common mode until being separated inside the vacuum. One of the state prep beams follows the same path as the Raman beams and is delivered through the same window. The ν_3 beam carries the other state prep beam, which is collimated into a 7 mm $1/e^2$ diameter beam then launched into a separate vacuum window. Since the ν_3 beam carries only a single polarization, a clean up polarizer is introduced after the collimator to prevent polarization drift. The ν_2 beam carries light for MOT and PGC, which is collimated into a 7 mm $1/e^2$ diameter beam and delivered through a third vacuum window to be received by the MOT chicane optics internal to the vacuum. Similarly to ν_3 , the ν_2 beam carries only a single polarization and has a clean up polarizer at the beginning of the free space collimator optics train. Each beam path contains a half




Figure 3.4: The collimator plate delivers light beams from the three AOD output fibers. The ν_1 line (red) delivers Raman and state prep, ν_2 (green) delivers MOT, and ν_3 (blue) delivers the other state prep (and detect) beam. The picture (left) has blue tape over some of the pickoffs where intensity monitors are installed.

waveplate to rotate the input light. If the path has a clean up polarizer, the waveplate is after the polarizer, but before any further optics. Next is a pair of Risley prisms that can be used to fine tune the steering angle of the beam for alignment optimizations. A

translation sphere was also included in each path to adjust the translation of the beam, but was found to effect too small a movement to be useful. Each beam is guided by a series of optically bonded beam splitter cubes designed to reflect nearly all of the light. The 45 degree surface in the cube has nearly 100% reflectivity for light with polarization parallel to the surface and about 99% reflectivity for light with polarization perpendicular to the other mode. This allows the placement of photodetectors that measure the roughly 1% leakage light. Monitoring the light at the sense head just before injection into the vacuum provides the best achievable measure of optical power reaching the atoms. These intensity monitors are used for both monitoring and active servoing of each optical signal to account for optical intensity fluctuations.

At the center of the vacuum chamber is a ceramic Macor block to which is mounted a set of MOT optics and wound coils for the quadrupole MOT field, altogether referred to as the MOT chicane (figure 3.5). Mirrors in the MOT chicane guide a single input beam through a series of reflections that form three orthogonal beams that overlap in a center region. The beam is fully retroreflected such that it returns through the MOT chicane along the same path in which it entered, thus forming the three pairs of counterpropogating beams necessary for MOT capture. The mirrors are coated such that the polarization state of reflected light is not changed from that of the incident light and also to maximize reflection. A $\lambda/4$ waveplate at the beginning of the path sets the light to circular, and a $\lambda/4$ waveplate at the retro mirror flips the circularity such that each pair of beams is oppositely polarized, enabling trapping as described in section 2.3. The mirrors are installed in custom-built tilt mounts to adjust alignment.

The MOT coils are wound directly onto the Macor block. Each coil is 84 turns of 26 AWG wire with a mean coil diameter of 43 mm and separation of 26 mm. The four leads of the two coils are connected to a 9-pin D-sub vacuum feedthrough. The lead pairs are kept separate so the electronics outside of the chamber can flip the current to choose between Helmholtz and anti-Helmholtz mode. The transfer function for the coils




Figure 3.5: Solid model and photo of the MOT chicane. (a) Exploded view of the MOT optics with support structure. (b) Photo of the MOT chicane on benchtop support fixture. (c) Alternate exploded view without support structure.

is designed to be $60 \text{ mA}/(\text{G}/\text{cm})$ in anti-Helmholtz mode and $40 \text{ mA}/\text{G}$ in Helmholtz mode. The quadrupole gradient calculated is along the center axis defined by the coils, the gradient generated in orthogonal directions is half this value.

The MOT chicane is supported by a metal structure that mounts to the base of the vacuum chamber. The support structure also houses the Raman optics, which consists of the periscope and a corner cube (figure 3.6). The periscope is a custom optic made of optically bonded glass sections that receives light from two input windows and guides them to the atoms. The first input window delivers the two Raman beams and one of the state prep beams. One of the Raman beam polarizations is orthogonal and gets transmitted upward directly through the periscope, while the other Raman (and ν_1 state

prep) beam get reflected through the periscope and redirected down to the atoms by the corner cube (see figure 3.6). The other window delivers the second state prep beam, which gets transmitted through the periscope and up along the transmitted Raman beam. The construction is such that one Raman and one state prep beam get transmitted up and directly to the atoms, while the other Raman and state prep beams are sent past the atoms where they hit a corner cube. The corner cube reflects the beams around and back down to the atoms, thus setting up the pairs of counter-propagating Raman and state prep beams. The periscope and corner cube are carefully designed such that the beams are always parallel, even if the input beam is tilted. Tilt and translation of the corner cube produce displacements of the beam, but do not produce angular misalignments. The angular misalignment is therefore bounded by construction to roughly $30 \mu\text{rad}$.




Figure 3.6: The Raman “racetrack” delivers the beams in opposite directions, with one injected straight up to the atoms and the other being directed around the periscope and through the corner cube from above. In the picture on the left, the periscope, corner cube, and MOT region are obscured by other components.

The construction of the Raman beams described above form a “racetrack” design, where the two beams are incident on the atoms from opposite directions. The Raman beams

must be in opposite directions to be Doppler sensitive (DS) and measure acceleration. If the beams are incident from the same direction, the position-dependent term gets canceled and the Raman interaction is Doppler insensitive (DI). A more common way to produce oppositely incident DS beams is to use a retroreflection configuration in which both beams are initially incident in the same direction and both get reflected directly back. One result of this retro arrangement is two DS Raman beam pairs. A Doppler shift is required to make only one of the pairs in Raman resonance. Two detuned DI beam pairs are also produced. The racetrack design, although demanding a slightly larger space for the travel of the outer beam, removes the unwanted extra DS pair and the two DI pairs. These DI beams can act as parasitic interferometers that introduce unwanted phase shifts in the LPAI measurement. The second DS pair produces a two-photon light shift to the Raman resonance of the other pair, another undesirable phase shift (see section 4.5). The benefits of the racetrack design therefore outweigh the slight increase in volume.

Two windows are located on the side of the vacuum chamber for observation of the atoms. One window looks directly into the center of the MOT chicane where the atoms are loaded. This is the primary observation point for fluorescence measurements of atom number and population ratio after interferometry. With a drop time of up to 12 ms ($2 \cdot 5$ ms interferometry, 1 ms state prep, <1 ms detection), the atoms fall up to $700\text{ }\mu\text{m}$. This is well within the viewing area of the initial window. A collection lens is placed inside the vacuum to increase the capture efficiency of scattered light, and the focus is located just outside the window. A quad detector made up of a sectioned 1 cm diameter photodiode chip receives the fluorescence light outside the window and provides an amplified analog voltage for each quadrant. Movement of the atom cloud can then be tracked by the relative intensities across the quadrants of the detector. A lower window is located 50 mm below the upper window to allow for longer-drop measurements. A drop time of 100 ms and therefore up to $T = 50$ ms is possible with this lower window (up to 100X more sensitive than $T = 5$ ms). The same setup of a collection lens and quad detector is placed

at the lower window.

To inject Rubidium atoms into the vacuum, a nozzle on the side of the chamber is mounted via a Conflat flange and fed by a reservoir of atomic Rb-85. The reservoir is a copper tube with a flange on one side and sealed by a pinch-off at the other. A sealed glass ampule of atomic Rb-85 is placed inside the pinched-off tube and mounted vertically to an elbow. The elbow terminates at a nozzle with an opening of 1 mm, which mounts to the conflat on the side of the chamber. Being an alkali metal, Rubidium readily oxidizes in air, and therefore must be kept under vacuum at all times. Only once the whole chamber is evacuated, the Rubidium is released from the ampule by squeezing the copper tube until the glass is shattered. The rate of atoms ejected from the nozzle is determined by the pressure of gaseous Rubidium in the reservoir, which is controlled by changing the temperature and therefore the vapor pressure of the reservoir. The nozzle is heated above the temperature of the reservoir to avoid condensation of Rubidium on and around the nozzle. Atoms coming out of the nozzle form an atomic beam that is centered at the MOT chicane. The rate of atoms injected into the system can thus be controlled, and the MOT is loaded primarily from atoms in the beam, allowing the background pressure to be kept low.

The vacuum is maintained by three components that keep the vacuum pumped down to around 10 nanotorr. The first is an SAES Capacitor Z-100 non evaporable getter (NEG) that absorbs most molecules such as H_2 , N_2 , and O_2 at a rate of around 150 l/s. The NEG is passive, only needing to be activated at initial pump down of the chamber. The NEG cannot absorb some molecules, such as noble gasses and Rubidium. To remove noble gasses, an active ion pump (Modion C-0609QM) is introduced. The ion pump has a relatively slow pump rate of 2 l/s and is meant only to control small leaks of noble gasses. Given the constant influx of Rubidium from the atom source, a POCO graphite NEG is also added to absorb Rubidium more quickly. Removal of Rubidium prevents the background pressure from becoming too large, which can negatively impact the MOT and




Figure 3.7: Vacuum components in the sense head. (a) Solid model showing each vacuum-maintaining component and atom source. (b) The atom source is comprised of a Rb-85 reservoir and nozzle to emit inject the atoms into the main chamber. (c) The atom source with heaters wrapped in insulation to control emission rate.

LPAI sequence.

Pump down of the chamber is initially done using a turbo pump connected to a copper tube. The tube is permanently affixed to the chamber via a conflat flange. At initial pump down, a bakeout is performed to more quickly remove unwanted gases and water from the chamber. At this point, the NEG and ion pump are activated, and the tube is pinched off to form a vacuum-tight seal. The only remaining actively powered component needed to maintain vacuum is the ion pump. It draws relatively little power and can be kept active for days with a small pack off AA batteries. Vacuum pressure was found to hold steady at or below 10 nanotorr for a full five-year operational period, and showed no signs of degradation by the time the system was opened up for repair of internal optics (see section 4.6). The vacuum likely could have been maintained for many more years without issue.

Surrounding the vacuum chamber is a two-layer magnetic shield that dampens outside sources of magnetic fields. Field variations can have unwanted effects on the atoms, some of which will be discussed in sections 4.4 and 4.5. External magnetic field sources include

Earth’s magnetic field, nearby electronics, and various ferromagnetic materials such as metal fixturing. Both static and time-varying fields can be problematic, so it is beneficial to shield from outside sources as much as possible. The shields are made out of 0.050 inch thick Mu Metal, which has a high permeability of around 80,000 (about 100X larger than stainless steel). They are cylindrical, and stitch together at a top and middle seam. The shielding factor of each layer is designed to be 40-50X along the vertical axis for a combined shielding of 2000X.

Attached to the inside walls of the inner shield is a pair of printed flexible Kapton PCB that encompass the vacuum chamber. Wires run through the PCB to form three orthogonal pairs of bias coils. The Z axis coils have a transfer function of 2.07 G/A and the X/Y axis coils have a transfer function of 0.28 G/A. Capable of safely running at 1 A, the max output field is then 2 G along z and 0.3 G along x and y. The purpose of these coils is to provide trimming of residual fields that make it through the shields. The Z axis coil also acts as the bias coil to separate out magnetic sublevels during the interferometer sequence.

Over the course of three years of near-continuous operation, a loss in atom number was observed. In attempt to address this, a bias field during MOT was applied to shift the center of the quadrupole field to a location in which a larger number of atoms could be effectively loaded. As will be discussed in detail in section 4.6, this did not address the root cause of the loss of atom number, but was successful in achieving a temporary boost in atom number. As this loss continued, stronger bias fields were required, so two new Helmholtz coil pairs were installed on 3D-printed mounts inside the shield. With a transfer function of 1 A/G, they are capable of driving stronger fields than the X/Y axis PCB trim coils.

Located atop the sense head in between the inner and outer shields is the cosensor triad. Three identical single-axis accelerometers (Colibrys SF1600SN.A) are mounted orthogonally to a metal block which is connected to the sense head via three short posts.

The Colibrys accelerometer was chosen due to its relatively low noise ($\sim 3 \mu\text{g}/\sqrt{\text{Hz}}$) and high bandwidth (1.5 kHz). The IA of the z-axis cosensor is nominally aligned with the atom IA. The cosensor is mounted as directly as possible to the atom sense head so they are as mechanically linked as possible. Mechanical coupling and IA alignment is critical for feedforward/feedback performance, as will be discussed in section 4.2.

3.1.1 Optical Alignment

Optical alignment of the sense head takes two stages, the first being initial alignment of the vacuum optics. These components need to be initially cleaned then carefully handled during assembly to avoid contamination of the ultra-high vacuum. The MOT chicane is assembled on a benchtop fixture, where laser light can be inserted for initial alignment. MOT alignment is carefully adjusted as mirrors are installed one-by-one. The light is centered through each opening of the MOT chicane until it passes centered through the output where the retro mirror is located. The retro mirror is then installed and the alignment is fine tuned to ensure maximum centering and counter-propagation of each beam. The beam is also broadened to 8 mm to ensure minimal clipping of the light is occurring through the chicane. The laser then is fully retroreflected back to its input. Waveplates are also adjusted to ensure maximum circularity of the light. The waveplates initially chosen were found to have a poor circularity and relatively high temperature dependence, so compound quarter wave plates were constructed out of a multi-order quarter and half waveplate. With the proper relative orientation of the two waveplates, this effectively cancelled out these effects, resulting in a vacuum-compatible compound quarter waveplate capable of producing reliably greater than 89 degree circular light.

With the MOT chicane internally aligned, it is then installed into the vacuum structure, where light can now be injected by the collimator plate through the vacuum window. The alignment of the input light is adjusted to ensure it matches the benchtop setup and thus produces non-clipping, reotroreflected beams. The Raman beam path can now be aligned

using light through the ν_1 collimator into the periscope. The input beam straight through the periscope is aligned to be as centered as possible through the MOT chicane, ensuring the beam will be centered on the atom cloud. The other Raman beam travels around the periscope and through the corner cube before being sent down to the atoms. The corner cube is adjusted to ensure maximum overlap and angle alignment of the Raman beams. Finally, the ν_3 state prep beam is aligned to overlap with the Raman beams (and therefore also the state prep light that is inserted through ν_1).

Alignment of the beams can be benchmarked by looking at positions of light returning through the vacuum windows and hitting the input collimators. The fully retroreflected MOT beam should enter and return to the same position through ν_2 , while the Raman and state prep beams entering from ν_1 (ν_3) returns to the opposite ν_3 (ν_1). After the alignment is complete, the vacuum chamber is sealed and the collimator plate is removed to allow for pump down and bakeout of the chamber. After vacuum has been established, the collimator plate is reinstalled. Dowel pins and a cam lock ensure highly repeatable replacement of the collimator plate. Alignment of the beams is then checked via their return positions at the collimators.

3.2 Photonics Subsystem

The photonics subsystem is responsible for generating the laser light necessary for manipulation of the atoms at the sense head. All components are COTS and fiber coupled, connected together either by FC/APC connectors or fused fiber splices. A single laser source is used, which greatly simplifies the front end optics and decreases SWAP. All components are packaged in a 12X12X6 inch box that sits next to the sense head, referred to as the *photonics box*. A cable bundle provides power, communication, and RF control. A single optical fiber delivers light to the AOD box, which ultimately delivers light to the sense head.

Laser generation begins with a single 1560 nm laser amplified by an erbium-doped fiber amplifier (EDFA). 1560 nm light is a convenient wavelength for optics technology because of its predominance in telecommunication industry and the resulting maturity of many devices. The EDFA delivers approximately 1000 mW of 1560 nm light into a periodically poled lithium niobate crystal (PPLN), which frequency-doubles the light to about 500 mW of 780 nm light. One of the great practical benefits of working with Rubidium 85 is that the wavelength of light equivalent to the $S_{1/2} \rightarrow P_{3/2}$ transition is 780 nm, which as described is readily achieved from mature COTS technologies. Significant heat is generated in these three front-end devices, so they are each independently temperature controlled by thermo-electric coolers (TECs) to ensure stable operation and prevent thermal damage.




Figure 3.8: Block diagram of the laser system. 780 nm light is generated using COTS fiber-coupled components.

A 1% fiber tap feeds 780 nm light to a laser control module (LCM) custom-designed by Vescent Technologies that frequency-locks the laser light. The LCM contains a Rb-85 cell which it utilizes to perform saturated absorption spectroscopy. The frequency is chosen to be 500 MHz red-detuned from the $F = 3$ to $F' = 3'/4'$ crossover feature (figure 3.9). This sets the laser frequency to a point at which no optical pumping will occur without further modulation. The laser frequency is coarsely tuned by adjusting the die temperature, then finely servo'd by adjusting the pump current. After identifying the desired spectroscopy

feature, the LCM automatically servos laser current to keep the frequency locked to the feature.




Figure 3.9: Rubidium spectroscopy is used to lock the laser frequency to the $3'/4'$ crossover, detuned by 500 MHz. (left) A large frequency sweep of the absorption spectrum. (right) Zoomed in error (filtered derivative) signal showing the $3'/4'$ crossover to which the LCM locks the laser.

After spectroscopy locking, the light passes through an electro-optic modulator (EOM). The EOM modulates the frequency of the light via RF generated by a custom-built RF generation system. The EOM provides fast $< 1 \mu\text{s}$ tuning of the frequency that makes up the difference between the locked laser frequency and the level needed to drive particular transitions within the hyperfine structure of the atom. A single modulation frequency will generate a high and low sideband, but only the high sideband will be on resonance with a transition. The EOM in this system has a damage threshold of 300 mW, which is the effective limit of total optical output of the laser system. The light enters a final 1% fiber tap used for optical monitoring and diagnostics, then is sent to the AOD box.

All components described above are fiber coupled. Most connections are made by fused fiber splice. While using FC/APC mates may be easier and more convenient, they are more subject to misalignment and therefore polarization variation than a direct splice. Slight misalignments may also cause some of the optical power to leak into the cladding of the connector, which at $>100 \text{ mW}$ can generate a lot of localized heat and melt the




Figure 3.10: An example EOM modulation scheme. The carrier (black) is modulated by two RF frequencies. A high and low band is generated for each (blue, green), along with mixing between the two (cyan).

connector at and around the core. A spliced fiber is capable of handling much more than 1000 mW of optical power, with typically only around 0.3 dB loss. Polarization maintaining fiber is used to preserve the orientation of light through the fiber, and a splice can easily achieve rotational alignment that maintains >20 dB polarization extinction ratio (PER).

Components are mounted to the interior walls of the photonics box. The thin metal walls act as a minor heat sink for components that generate less heat. Heat sinks for the laser and EDFA are mounted on the corresponding outside location of the component to allow greater flow of cooler air. Fiber between components and splices and FC/APC connectors are managed in spools also mounted to the inside of the box. This format allows for convenient access to individual components, but leaves a lot of unused empty space.

The laser system has performed reliably over more than three years of near-continuous operation with the exception of the PPLN (see next paragraph). The system has proven robust to the environmental threats, notably vibration and shock. The box is mounted on vibration-dampening rubber pads to decrease direct vibration, which has proven sufficient to protect the system over multiple field tests. The greatest challenge to reliable operation of the laser system is thermal fluctuations. While not a significant risk for damage, thermal

fluctuations generate optical power fluctuations. Large enough power variations affect the spectroscopy signal and can cause the LCM to lose its lock feature. Fast fluctuations in the range of 5-10°C inside the photonics box tend to cause the LCM to lose lock. With the box closed and the larger system thermally shielded, fluctuations at this level are rarely seen, even outside in direct sunlight. The lock system has thus proven quite reliable in nearly all reasonable field environments.

Most of the components have demonstrated excellent reliability over the life of the system, with a couple notable exceptions. The first was a failure in the LCM. What was believed to be a mechanical break in the coupling of the detector to the spectroscopy cell produced large signal variations, thus making the LCM lose lock very regularly. The first LCM operated for around 2 years, and after replacement by an identical part has operated reliably ever since. Second, the PPLN was found to degrade after about half a year of operation. Optical power out started to steadily decline, a trend that was allowed to continue until optical pumping of the atoms was significantly impacted by a lack of power. The PPLN was replaced with an identical model, and the same issue arose after roughly another six months. The source of this degradation was never definitively determined, but the most likely explanation was a drift in the alignment of the fiber coupling to the chip inside the package. Another possibility was optical darkening of the crystal, but this is less likely at less than 1000 mW of optical power. The PPLN was replaced with a comparable one from a different vendor, and the output power has subsequently remained steady for greater than a year.

Fiber outside the photonics box is routed through a series of 3D printed structures that support and protect the fiber. A single fiber delivers light from the photonics box to the AOD box, and three fibers deliver light from the AOD box to the sense head. The fibers are fully wrapped and stowed in these structures, with almost no exposed gaps. The fiber structures serve to mechanically protect from any direct damage and also to add a layer of thermal isolation from the environment. Thermal fluctuations in the polarization

maintaining fibers can produce polarization drift between modes. Polarization drifts are largely removed by polarizers located at fiber outputs, which then results in power fluctuations. Thermal isolation thus helps to decrease variation in optical power.

3.2.1 AOD box

The AOD box acts as a fast single input multiple output light switch, steering input light into three fibers that ultimately deliver the light to the sense head. At the core of the AOD box is a free-space acousto-optic deflector (AOD), which deflects light traveling through a crystal via the acousto-optic effect. The deflection angle is a function of an RF signal applied across the crystal. The amplitude of deflected light is a function of the strength of the RF signal. The photonics box delivers about 200 mW of frequency-locked 780 nm light through a single polarization maintaining fiber. To avoid unwanted polarization-dependent effects from leakage light in the off axis of the fiber, a cleanup polarizer is placed at the output of a collimator. The light then travels through the AOD, which deflects the light if an RF signal is applied. Four discrete signals are generated depending on the part of the experiment cycle – Raman P (35 MHz), state prep (44 MHz), MOT/PGC (53 MHz), or Raman S (62 MHz). The RF signals are 9 MHz apart to allow sufficient separation of deflected light to avoid accidental coupling of light into other channels. The RF signal pulses generated by the RF chain were initially found to have unwanted frequency content that produced unintended coupling of light into other channels, so the RF generation was carefully considered to ensure pure frequencies drive the AOD (see section 4.1). Four pickoff mirrors are placed in series to catch only one of the four deflected beams, while the remaining light passes straight past to a beam dump. A photodiode at the beam dump provides monitoring of overall optical power in the AOD box.

Chains of optics in custom-built mounts then manipulate the beams and guide them into three fiber-coupled collimators, labeled ν_1 , ν_2 , and ν_3 . Pairs of Risley prisms for each




Figure 3.11: The AOD box guides the main beam into three delivery fibers to deliver fast optical pulses to the sense head.

of the beams are placed after the pickoffs to allow for fine adjustment of the beam steering. The MOT/PGC (53 MHz) beam gets coupled into the ν_2 fiber collimator. The Raman S and P (35 and 62 MHz) beams both get coupled into the ν_1 fiber collimator, but the polarization of Raman P gets rotated by 90 degrees to keep it separated from the Raman S light. This distinction of Raman beams is needed because they are both modulated by the EOM before separation, and it is necessary to avoid ambiguity of the direction of the effective $\mathbf{k}$ vector at the atoms. Placed in the Raman P path is a roof prism on a slidable stage. This allows adjustment of the path length relative to the other Raman beam, which can be used to optimize Rabi rate in the presence of competing transitions (in particular, a 3.5 GHz-detuning Raman transition, see section 3.3.1). The state prep (44 MHz) beam gets split into two beams by a half waveplate then polarizing beam splitter cube. The waveplate is rotated such that roughly 10 percent of the light gets reflected

where it then couples into the ν_3 fiber. The remaining 90 percent of light passes through to a 90/10 non-polarizing beam splitter cube, which reflects 10 percent of the light into the ν_1 fiber along with the Raman beams (while the remaining 90 percent is dumped). The half waveplate can then be rotated to adjust the relative state prep power carried by the two fibers and thus the balance of state prep light in the counterpropagating beams at the atoms. Most of the light is dumped because the state prep (and detect) beams only ever drive resonant transitions, which requires relatively little optical power.

The input fiber is a pigtailed collimator spliced directly to the photonics box output fiber. The three AOD output fibers are mounted into the 3D printed fiber support structure and spliced to FC/APC connectors with anti-reflection coatings that insert light into collimators at the sense head.

With free space optics coupling into fiber collimators, robustness of the alignment was a big concern. Slight misalignments of the beams can result in dramatic loss of fiber coupling, so all optics need to remain tightly locked down. Similar to the photonics box, rubber damping pads are placed under the box where it is mounted to the baseplate. Despite these concerns, the AOD box alignment proved highly robust to shock and vibration, demonstrating no discernible change in fiber coupling over multiple shipping and field deployments. A slight temperature dependence of the alignments was seen, but is rather low over many degrees Celsius. A heater on the back of the box keeps the temperature locked, preventing significant thermal variation.

3.3 Electronics

The main control system automates the experiment by commanding the photonics system and monitoring resulting signals, primarily from the sense head. At its core is a field programmable gate array (FPGA) that sends out and receives carefully timed signals to the rest of the electronics. An ARM processor, referred to as the PROC, communicates

with the FPGA to act as a human interface to provide settings and returning data. The RF control system receives signals from the FPGA and is responsible for manipulating the laser light that ultimately reaches the atoms. A set of analog to digital converters (ADCs) provide signals from optical monitors and the three-axis cosensor triad.




Figure 3.12: Block diagram of the electronics. The left-hand group is located in the electronics rack, and the right-hand group is located at the sense head structure.

The *FPGA box* contains the set of hardware that implements the experiment and routes communication and signals to the necessary hardware. At its core is the FPGA, which executes the instructions that comprise the experiment cycle. The FPGA was modified to reference a chip scale atom clock (CSAC) for stable and accurate timing of signals. The digital signals to and from the FPGA are routed through sets of aggregator cards, which transmit the signals over HDMI to reliably send and receive over longer distances (multiple feet). The first set sends signals to the RF box, which manipulates the photonics system and generates specific light pulses at the sense head. The second set of aggregators sends and receives signals to and from stacks of ADCs located near the sense head. The ADCs receive analog signals from intensity monitors and quad boards. These signals provide measures of the optical output and atom fluorescence, which ultimately reports the interferometric measurement. The cosensor signals are also digitized through

this stack of ADCs and sent to the FPGA. A gyroscope triad located on the baseplate measures rotations of the system and generates a digital output sent directly to the FPGA through the aggregators.

Temperature control is implemented primarily through sets of resistive heaters regulated by individual PID controllers. The sense head contains a set of heaters pasted on the outer surface of the sense head. Care was taken to distribute heaters in attempt to provide an even distribution of heat, thus avoiding large thermal gradients across the system that could potentially affect optical alignments. Another heater pair is placed on the inner shield to provide an additional layer of thermal regulation. The baseplate is heated by a large circular heater centered on the underside of the plate to provide evenly distributed heat, especially at the three feet of the sense head. The cosensor triad has its own heater to provide local temperature stability because its bias is highly sensitive to temperature. Three heater groups on the Rubidium atom source provide heat to generate vapor pressure behind the nozzle. A thermal gradient hottest at the nozzle ensures atoms do not condense on and clog the nozzle. The AOD box also contains a heater on the back plate to prevent temperature-dependent alignment variations. All of the PID controllers are packaged into the so-called *thermal box*, along with a set of DAQs that record temperature. The DAQs receive roughly twenty thermistors and thermocouples distributed across the sense head and other components in the system to provide logs of temperature during operation.

Additional auxiliary monitors and controls were added to provide further monitoring of the system. The LCM is not integrated into the main control system and is communicated with directly via USB. Commands are sent to lock the laser initially, then a set of monitors are reported regularly to indicate behavior of the lock. Tracking of the LCM provides insight into its behavior and stability under adverse conditions. The temperature DAQs are similarly recorded over USB to provide tracking of temperatures variations of the system.

The final auxiliary monitor is a pair of inclinometers installed on the baseplate.

These single-axis inclinometers have excellent resolution and bias stability, providing an accurate measure of angle near zero of down to single-digit millidegrees ($\sim 20 \mu\text{rad}$). The inclinometers can thus be calibrated to report zero tilt when the gravimeter is pointing directly up, then used to make sure the system is as level as possible. This provides a way to accurately measure gravity in absence of an IMU. The inclinometers can be used to record tilt angles but have a temperature-dependent and nonlinear scale factor that makes them less accurate away from zero. Their bandwidth is on the order of 0.5 Hz, so they can only be used in static situations.

3.3.1 RF Control & Modulation

Section 3.2 describes how light is generated and delivered to the sense head. The light is locked at a frequency below which any $S_{1/2} \rightarrow P_{3/2}$ transition can occur. To induce specific transitions, the light is manipulated via the *RF box*, a set of components that generate RF signals that modulate the light. Modulation occurs at two devices – the EOM and the AOD. The RF box then contains two separate RF chains, the EOM RF chain that manipulates the frequency of the light, and the AOD RF chain that guides light into the desired channels. The RF box is controlled in real time by the FPGA, which sends signals over HDMI cables from the FPGA box via aggregator cards.

The source of the AOD chain is a direct digital synthesizer (DDS) that provides two channels of 35-62 MHz RF at 1 dBm. Each channel has a separate branch of the RF chain that begins with a variable voltage attenuator (VVA) that attenuates the RF output based on a drive voltage. Next is a pair of switches that are toggled on or off in unison for maximum output isolation when off. The two branches are combined and sent together through a series of two amplifiers, which together provide up to 45 dB of amplification for a maximum of 30 dBm output power. The combined, amplified signal then goes through a 30-88 MHz bandpass filter to remove higher order harmonics and is sent over an SMA cable to the AOD.




Figure 3.13: The AOD RF chain provides up to 30 dBm of 35-62 MHz RF pulses that control the AOD. The chain along with the AOD box effectively acts as a light switch that delivers optical pulses to the sense head.

The two switch pairs receive digital signals from the FPGA to toggle either channel on or off, while the VVAs receive variable analog voltages from a DAC that is controlled by the FPGA. The first DDS channel is the primary channel and is variable, outputting either 35 (Raman P), 44 (state prep/detect), or 53 (MOT/PGC) MHz. The second channel only outputs 62 MHz, and is used only to generate Raman S light pulses during interferometry. The AOD RF chain thus enables fast control over when certain light is sent to the sense head and at what optical power.

Once implemented, it was discovered that the switches chosen (Analog Devices HMC849ALP4CE) were designed more for higher-GHz signals and had a low-frequency rolloff that greatly attenuated the ~ 50 MHz inputs. Additionally, the switches had a sharp turn-on time that generated a non-negligible distortion of the output generated by the AOD when subject to sharp pulses. This manifested as a significant amount of light distributed across deflected angles, some of which coupled into the other fibers. The result is unwanted light reaching the atoms through the other fiber outputs, thus impacting

performance of the gravimeter. To address this issue, the RF switches were replaced with Mini-Circuits HSWA2-30DR+ switches that had a slower turn-on time. The transmission of light by the AOD box therefore has a response time of roughly $1\text{ }\mu\text{s}$, but does not skew the output light into undesirable channels.

Light can be delivered to the atoms as described, but is still detuned from any transition by the 500 MHz set by the laser lock. To make up the difference in frequency, the light is modulated by the EOM, which is controlled by the EOM RF chain. The chain begins with four dual channel signal generators (Windfreak SynthHD) that output fixed RF signals in the 500 to 3000 MHz range at -10 to +15 dBm. Each channel is responsible for driving a different transition (MOT cycling, state prep, repump, etc.) and is selected by a series of switches controlled by the FPGA. The channels can be grouped loosely into two sections, one for primary/pumping transitions and the other for repump transitions.




Figure 3.14: The EOM RF chain modulates light pulses with up to 20 dBm of 500-3500 MHz of RF. This modulation is used as fine control of the light frequency that enables specific interactions with the hyperfine structure of the atoms.

The repump signal is either 3422 or 3431 MHz and drives the $F = 2 \rightarrow F' = 3$ transition. There are two different frequencies for repump because it modulates the light

going through either the 53 MHz ν_2 or 44 MHz ν_1/ν_3 channel depending on the part of the experiment sequence. The frequency gets selected by a switch then attenuated by a VVA. The signal is then amplified by 20 dB for a total output of roughly 0 to 10 dBm. The other group of switches selects between six RF inputs – two for cooling/trapping, two for Raman, and one each for state prep and detection. The selected frequency gets attenuated by a VVA then amplified by 44 dB for up to 20 dBm of total RF power. The natural linewidth of the $S_{1/2} \rightarrow P_{3/2}$ transition is 6 MHz, so accuracy and stability of the RF frequency need to be on the order of a single MHz to avoid variations in pumping. This is easily achievable with modern RF hardware.

Signal	Transition	Frequency	Optical Channel
MOT pump	$F = 3 \rightarrow F' = 4$ (-3Γ)	489 MHz	ν_2 (53 MHz)
PGC pump	$F = 3 \rightarrow F' = 4$ (-14Γ)	420 MHz	ν_2 (53 MHz)
MOT/PGC repump	$F = 2 \rightarrow F' = 3$	3422 MHz	ν_2 (53 MHz)
State prep	$F = 3 \rightarrow F' = 3$	396 MHz	ν_1/ν_3 (44 MHz)
Detection	$F = 3 \rightarrow F' = 4$	516 MHz	ν_1/ν_3 (44 MHz)
SP/detect repump	$F = 2 \rightarrow F' = 3$	3431 MHz	ν_1/ν_3 (44 MHz)
Raman +k	$F = 2 \rightarrow F = 3$	3026.73 MHz	ν_1 (35/62 MHz)
Raman -k	$F = 2 \rightarrow F = 3$	3080.73 MHz	ν_1 (35/62 MHz)

Table 3.1: EOM RF modulation signals and their purposes.

The Raman RF signal gets modulated before amplification by a custom-made single sideband mixer (SSB). The SSB generates a primary single output frequency generated as the sum of the main RF Raman signal from the signal generator and an 18 MHz signal generated by a third separate channel of the DDS. The DDS is controlled by the FPGA and allows fast fine tuned control of the frequency and phase of the output signal. This DDS control is the mechanism by which feedforward is implemented – the frequency and phase of each Raman pulse can be adjusted based on measurement of specific force provided by the cosensor (see section 2.4). Unlike the other signals that perform optical pumping, the Raman transition requires a very fine-tuned selection of the RF frequency that defines the two-photon detuning. Frequency accuracy on the scale of 10s of kHz is easily achieved with the DDS modulation scheme.

It is worth noting that the symmetric sidebands generated by the EOM produces not only the intended resonant frequency, but an additional off-resonant sideband. The RF modulation scheme was carefully designed such that none of the off-resonant sidebands unintentionally coupled with another transition within the hyperfine structure of the atom. The unwanted sideband can therefore be neglected as producing any interaction, though the Raman interaction deserves special consideration. The counterpropagating Raman beams are both modulated by the EOM, which produces the relevant two-photon resonance with one carrier frequency and one sideband. The symmetric modulation also produces a 3.5 GHz detuned Raman pair, which is far enough detuned to be neglected. Nonetheless, an interferometer can in principle be formed through this highly detuned interaction, and therefore may be worth consideration as a source of error in the future.

The two selected RF signals (pumping and repumping) then go through another set of switches with low pass filters that cut out higher order harmonics. The output can be selected to be just the main output (such as for Raman), or both the main and repump signals combined through an RF coupler (such as for cooling and state prep). The final output is then sent to the EOM by an SMA cable. An RF detector is placed at the output of a 1 percent RF pickoff located near the EOM to provide a monitor of the RF signal received by the EOM.

The amplitude of the modulation relative to the carrier (unmodulated) frequency is adjusted by the VVAs and typically chosen to be the same amplitude as the carrier for pumping signals, and roughly 1/10th of the carrier for repump, since very little repump is required to shift the population back into the desired state. The Raman sideband amplitude is tuned to cancel AC Stark shifts, as will be discussed in section 4.4. Note that the RF modulation of the EOM generates both a high and low sideband, but only the high sideband is near any hyperfine transitions. The low sideband and carrier are hundreds of MHz detuned from any transitions, preventing any unwanted interactions with the atoms. When both pumping and repumping are present multiple sidebands are

generated, but the frequencies are carefully chosen to avoid accidental pumping of the atoms.




Figure 3.15: Level diagram of Rb-85 hyperfine structure and the transitions for each part of the experiment.

The final component in the RF box is the coil driver board that generates current to supply to the bias and MOT coils. It is commanded by digital communication from the FPGA, and is separate from the rest of the RF box systems. The coil driver board receives isolated 5 V power to supply the MOT coil current to avoid electrical pickup with the rest of the electronics. It is designed to provide fast turn on and turn off times of magnetic fields through the experiment cycle, but was found to have a significant time constant. The design of the circuitry is being considered for redesign, but is sufficient with some settling time to provide the fields needed to manipulate the atoms.

3.4 Experimental Sequence and Data Acquisition

The broad outline of the experiment sequence is catching and preparing the atoms, dropping the atoms and LPAI, then detection. Each step is optimized to provide the largest and cleanest measurement of specific force for each cycle. The procedure is then repeated continuously to produce periodic measurements of gravity. The entire cycle typically takes ~ 25 ms (~ 40 Hz) under optimal conditions, but has historically been slowed to ~ 25 Hz when MOT loading was poor.




Figure 3.16: The experiment cycle consists of a sequence of events that catch and prepare, then measure the atoms. The red trace represents relative optical intensity of each event. Times are exaggerated for illustration.

The experiment is controlled by the FPGA, which runs a compiled firmware designed in Simulink. The compiled FPGA firmware defines the sequence and timing details of the cycle – commanding RF and light pulses, reading ADCs and other sensors, implementing feedforward, etc. An ARM processor (“PROC”) local to the FPGA board runs a real-time Simulink front end that acts as an interface to the FPGA program, allowing changes to timing setpoints and control of the RF system. Data is sent from the FPGA to the PROC, where data is collected and saved for analysis. The feedback Kalman filter is implemented on the PROC, where cosensor and LPAI signals are used to correct sensor error. A Windows laptop executes the PROC Simulink program and provides a human interface with buttons and plotted graphs of data output. The PROC and FPGA are local to a COTS FPGA board and are capable of executing the experiment independently, but require communication with the Windows laptop to provide data output and human

control of the system.

3.4.1 MOT

The theory of MOT loading is described in section 2.3; this section will discuss practical implementation. At the beginning of the sequence the quadrupole gradient field is turned on. $F = 3 \rightarrow F' = 4$ cooling light red-detuned by approximately 3Γ is directed through the three dimensional MOT chicane along with a smaller amount of resonant $F = 2 \rightarrow F' = 3$ for repump. The detuned cooling light generates the magneto-optic trap while the repump light prevents atoms from getting stuck in the $F = 2$ dark state where the cooling light would be unable to interact with them. The detuning of the cooling light affects capture efficiency, and the optimal detuning depends on a number of factors such as overall optical power. A loss in atom number was observed over the long-term operation of the gravimeter (see section 4.6), and decreasing the detuning to closer to 1.5Γ was found to produce a larger relative atom number due to the suboptimal loading conditions. This issue has since been addressed and the cooling light again operates near 3Γ . The period allocated for MOT loading is variable depending on loading efficiency and desired signal size. MOT load is the longest part of the sequence and typically lasts between 10 ms and 25 ms to produce a sufficient atom number.

Further attempts to improve loading performance led to implementation of additional steps in the MOT loading sequence. The first was to address potential beam misalignment by adjusting the loading center by introducing additional bias coils that changed the force-balanced region of the MOT. If there was a shift in the optimal loading location due to beam misalignments, the field center could be adjusted to match this offset. This means the atom cloud could be released off-center from the Raman beams, which would decrease Raman population transfer and thus interferometer contrast. A period at the end of MOT was added in which the quadrupole field remained on while the bias fields were turned off. This shifts the field center back to the Raman beam center, thus allowing

MOT loading at whatever location is most optimal while pushing the atoms back to the center of the Raman beams for the interferometry sequence. The second addition was to allow a period at the end of MOT to increase detuning by switching to the PGC frequency. The atom cloud becomes less spatially uniform as laser frequency approaches resonance, so increasing the detuning can help clean up the overall MOT distribution at the end of loading.

3.4.2 Polarization Gradient Cooling

To achieve temperatures below the Doppler limit (equation 2.38), which for Rb-85 is about $145\text{ }\mu\text{K}$, a two-step period of polarization gradient cooling (PGC) is implemented. After the MOT load, the light and field are turned off, allowing the atoms to begin free-falling. After sufficient time for the field to settle to zero, the cooling light is turned on but is now detuned by $7\text{--}15\text{ }\Gamma$. After about one millisecond in this first stage, the optical power is decreased by roughly 4X for a second stage. The full-power PGC beam allows a smooth transition from MOT loading and initially cools the atoms more quickly but has diminishing returns as heating from recoil begins to balance out the PGC cooling. Decreasing the power then allows the cooling to continue while decreasing recoil effects. Temperatures as low as single digit μK can be achieved when properly optimized. The atoms can begin to fall after MOT, and with the capture field turned off during PGC, the movement is no longer strictly confined. PGC continues to act as a slowing force, however, so very little movement of the atoms occurs until PGC light is turned off. The end of PGC is therefore considered the beginning of the free-fall period.

3.4.3 State Preparation

After PGC the atoms are in free fall, and need to be prepared in an initial state before interferometry. Resonant $F = 3 \rightarrow F' = 3$ cycling light is applied along with $F = 2 \rightarrow F' = 3$ repump. During state preparation, a horizontal bias field breaks the

degeneracy of the magnetic sublevels, which distinguishes the forbidden $|F = 3, m_F = 0\rangle \rightarrow |F' = 3, m_F = 0\rangle$ transition for π - polarized light. With the polarization of the state preparation light aligned with the horizontal magnetic field, the cycling light will excite the atoms unless they happen to end up in the $|F = 3, m_F = 0\rangle$ state, where they are dark to the photons and will remain in that sublevel. State preparation only takes about 40 μ s of cycling light to achieve around 90 percent of atoms in the $|F = 3, m_F = 0\rangle$ state. The interaction heats up the atoms, so as little time and optical power as possible is used while still achieving sufficient quality of population preparation. The atom temperature in this system is typically 15-20 μ K after state prep.

3.4.4 Interferometry Sequence

With atoms prepared and freely falling, the $\pi/2 - \pi - \pi/2$ pulse sequence is applied. A vertical bias field is turned on just before the sequence to split the magnetic sublevels. At 1 Gauss with a Zeeman shift of 470 kHz/G, any residual atoms in $m_F \neq 0$ sublevels are effectively detuned from the interaction, thus avoiding coupling of these states into the interferometer. Feedforward is applied before each pulse as described in section 2.4 to account for Doppler shift of mean movement of the cloud and the resulting phase shift. The interferometer sequence typically allows for 0.5 ms of field settling time at the beginning, then lasts another $2T$, typically at $T = 5$ ms, though anything between $T = 0.2$ ms to $T = 50$ ms is possible.

3.4.5 Detection

Immediately after the final Raman pulse, detection occurs. The bias field remains on to avoid field settling effects during detection. Detection measures atoms in the $F = 3$ state by applying a resonant $F = 3 \rightarrow F' = 4$ pulse. While the cycling light is on, the atoms will fluoresce, which is measured by the quad detector. A $F = 2 \rightarrow F' = 3$ repump pulse then excites all atoms into the $F = 3$ state, followed by another $F = 3 \rightarrow F' = 4$

detect pulse. The population ratio of atoms in $F = 3$ after interferometry normalized by the total number of atoms then equates to a measurement of the final interferometer phase and therefore specific force.

Under normal operation, the experiment cycle is repeated continuously and the MOT is immediately turned back on after detection. With normal settings ($T = 5$ ms), the atoms fall less than a millimeter and are still illuminated by all of the 7 mm diameter MOT beams when turned back on, thus allowing 97-99% of the atoms in the atom cloud to be recaptured. Recapture is essential to achieving a fast repetition rate of the experiment cycle, allowing a MOT load time on the order of 10 ms, as opposed to a loading time without recapture on the order of 1000 ms.

CHAPTER 4

Sensor Performance & Characterization

4.1 Basic Atom Optimizations

This section describes the basic measurements and optimizations performed to produce the best possible LPAI measurement. In broad terms the goal is to maximize the atom signal for best possible signal to noise ratio (SNR). This is achieved by maximizing atom number and by maximizing the fraction of atoms in the desired initial state. Optimizing the timing then ensures best possible Raman transfer and largest interferometer contrast.

The first step is to maximize atom number in the shortest reasonable time. MOT loading is discussed in detail in sections 2.3 and 3.4, so it will only briefly be discussed here. The total drop time is kept relatively short, so the atom cloud can largely be recaptured at the beginning of the following experiment cycle. With $T = 5$ ms the total drop time is roughly 12 ms, during which the atoms will fall about 0.7 mm in Earth's gravity. The MOT beams have a $1/e^2$ diameter of 7 mm, so the atoms never move far from the intensity center of the beams and are easily recaptured. The detuning of the MOT cycling light was found to have a large effect on the capture rate and is adjusted to

experimentally produce the largest atom number.

PGC is similarly experimentally adjusted to produce the coldest atom cloud. The temperature is fairly insensitive to detuning frequency, and was found to have little effect over > 10 MHz. The amplitude of light during the second PGC step proved to have the strongest effect on cloud temperature. Even without the MOT field, PGC can act to continue capturing atoms to some degree, so a tradeoff exists between atom number and temperature. Temperature is minimized at a fairly low PGC optical intensity, while atom number is increased at higher optical intensity. Since the primary goal of PGC is to produce cold atoms and MOT load time can be increased to increase atom number (to a point), optical intensity during PGC is optimized for coldest possible temperature.

The Raman interaction provides a convenient and sensitive probe of the atom state and can be used to evaluate the initial preparation of the atoms. A single Raman pulse is applied after state prep and the resulting population ratio measured at detection equates to the relative population transferred to the other state:

$$\text{population ratio} = \frac{\text{detect pulse 1 } (F = 3 \text{ only})}{\text{detect pulse 3 (total atom fluorescence)}}. \quad (4.1)$$

The Rabi oscillation can be characterized by varying the single Raman pulse length τ . By sweeping over a range of times the oscillation can be mapped out (figure 4.1). The real $\pi/2$ and π pulse lengths can be identified as the length at which half maximum and maximum transfer respectively is achieved. The oscillation can be fit to determine Rabi rate Ω and spontaneous emission characteristics. The expected pulse lengths for the $\pi/2$ and π pulses would simply be $\tau_{\pi/2} = \pi/2\Omega$ and $\tau_{\pi} = \pi\Omega$, but the measured Rabi oscillation deviates from a simple sinusoid due to spontaneous emission and the temperature of the atom cloud that creates a distribution of individual atom Rabi rates:

$$S(t) = 1 - Me^{-t/T_d} \left[1 + \cos(\Omega t/2) + \frac{2}{\Omega T_d} \sin(\Omega t/2) \right], \quad (4.2)$$




Figure 4.1: A Rabi scan demonstrates Rabi oscillation by varying the Raman pulse width. The oscillation is damped by spontaneous emission and distribution of Rabi rates due to atom velocity distribution.

where M is the maximum population transfer, Ω is the Rabi rate, and T_d is the spontaneous emission decay period. The model given by equation 4.2 includes effects of spontaneous emission and the spread of atom velocities [88]. It may not perfectly represent the measured trace due to spread of Rabi rate and potential parasitic interferometers that may be present (see the end of section 4.5).

The oscillation is damped by these effects and can be seen to decay after a few full oscillations ($t \sim 5\tau_\pi$). Decreasing cloud temperature therefore improves Rabi rate. The real pulse lengths are determined by direct inspection of the measured transfer amplitude because they deviate from nominal values. This is due to the fact that the optical pulse waveform changes shape with pulse duration because of the slow bandwidth of the switches in the AOD RF chain.

The Raman pulses were intentionally made to be slow because sharp turn on of square pulses created a frequency distribution that leaked light into multiple output channels of the AOD box. As an unintended effect, the slow rise time made the pulses asymmetric, which must be characterized carefully to avoid asymmetry of the interferometer sequence




Figure 4.2: Select time traces of Raman pulses. Real optical pulses are shorter than commanded RF pulses (lighter traces) and are delayed due to response time of the AOD. They are asymmetric due to slow response of AOD RF chain switches. Real pulse lengths and centers are defined by integration over windows over the pulse (shown as dark dashed lines).

(see section 4.5).

The state of the atoms can also be probed by varying the two-photon detuning of a single Raman pulse (with other Raman pulses turned off) and tracing out population transfer as a function of frequency detuning. This “detuning scan” (figure 4.3) provides several useful measures of the atom cloud. The main feature is a broadened peak centered at the clock frequency of the hyperfine states. This peak represents atoms in the ground state being transferred by the Raman pulse into the excited state. The transfer depth is a function of the pulse length (determined by Rabi scans) and optical power, and is limited by a number of factors such as temperature distribution and spontaneous emission.

The transfer peak is broadened due to Doppler shift of the atoms and is therefore a measure of the mean velocity of the atoms. The population transfer as a function of




Figure 4.3: A detuning scan measures response of the atoms to varied two-photon frequency detuning. This provides a useful way to measure the atoms' states, especially atom temperature and m-level distribution.

frequency $S(f)$ approximately takes the shape of a Gaussian

$$S(f) = 1 - M \exp\left(-\frac{(f - \mu)^2}{2\sigma^2}\right) \quad (4.3)$$

where M is the max transfer of population, μ is the center frequency of the transfer, and σ is the width of the peak. If the center frequency of the interaction with a stationary atom has been correctly determined, μ is the Doppler shift due to mean velocity of the cloud. With $\sigma = \Delta f = \frac{k_{eff}}{2\pi} \langle v \rangle$, the effective cloud temperature can be calculated as

$$\frac{1}{2}k_B T = \frac{1}{2}m_{Rb} \langle v \rangle^2 \rightarrow T = \frac{m_{Rb}}{k_B} \frac{4\pi^2}{k_{eff}^2} \sigma^2. \quad (4.4)$$

The width of the peak is also broadened by optical intensity, so low-power Raman beams must be used to produce an accurate measure of detuning width and therefore temperature.

Small humps can be seen to either side of the main peak, which represent atoms in magnetically sensitive hyperfine sublevels. These sublevels are separated in energy due to the z-axis magnetic bias field (along the direction of the Raman beams) applied

during the Raman sequence. Given the separation in energy, they do not interact with the on-resonance two-photon Raman interaction and therefore do not contribute to the interferometer sequence. It is still desirable to minimize atoms in the magnetically sensitive sublevels, which is achieved by optimizing state preparation. The total signal size is directly proportional to the number of atoms in the $m = 0$ sublevel since only these atoms interact with the Raman beams to produce interference. This is true to first order, but any other unexpected interaction with atoms in other sublevels, such as spontaneous emission or interaction with unintended frequencies in the Raman beams serves only to corrupt the measurement with unwanted signals. It was found that the best practical optical/RF optimization placed $\geq 90\%$ of the atoms in $m = 0$. Further increasing the optical/RF power has diminishing returns in which the atoms are heated by repeated pumping while fewer atoms are placed in $m = 0$. State prep efficiency is sensitive to the alignment of the bias field with respect to the pump beams' polarization. If the field is not orthogonal to the pump beams' propagation direction, their polarizations cannot be parallel to the bias field. The pumping is then less effective and requires more optical power to achieve the same state preparation. The alignment was characterized by applying trimming fields to tilt the effective bias field in three-dimensional space to find the vector at which state prep can be performed with as little as possible optical power.

As mentioned, cloud temperature is an important consideration in the operation of the gravimeter. Hotter temperatures means slower and more distributed Rabi rates, which decreases overall contrast of the interferometer. This is due to the larger cloud being further spread across the Raman beams intensity. It also means the cloud will grow faster as it is in freefall, which limits the length of time in which a meaningful interferometric measurement can be realized. Additionally, atoms that fall or spread out beyond the MOT beams cannot be recaptured, which decreases the signal size that can be achieved at a given overall repetition rate. Temperature is therefore minimized as much as possible by tuning PGC and by minimizing state prep power. Temperatures as cold as 6 μK have

been achieved with the system, but overall better interferometer contrast and system SNR are obtained at higher temperatures (largely due to state prep heating).

The optimizations of temperature, state preparation, and Raman pulses described above produce the best possible interferometer contrast. The interferometer fringes can be scanned by artificially introducing a phase shift in the Raman pulse. This scans out a sinusoidal interference pattern given by

$$S = M + C/2 \cos(\phi), \quad (4.5)$$

where M is the midpoint, C is the interferometer contrast, and ϕ is the interferometer phase. The contrast is defined as peak-to-peak of the sinusoid and represents the effective signal size on the population ratio.




Figure 4.4: Phase scan of the LPAI sequence. Sweeping the phase of the RF modulation of the relevant sideband demonstrates the interferometer response to varied phase. This acts as an experimental simulation of varied gravity. Real response to higher-frequency inertial input can be seen as fluctuations on the sinusoid.

4.2 Atom-Cosensor Input Axis Alignment

An important consideration in the construction of the hybrid cosensor + gravimeter system is the alignment of the input axes of the two sensors. Misalignments negatively impact performance in dynamic environments. This primarily manifests in the behavior of the feedback bias correction, as will be described in this section. The following prescription analyzes the behavior of feedback when an input axis (IA) misalignment exists and describes a method to account for it. The real misalignment is experimentally determined and a software-based correction is introduced that takes advantage of the off-axis MEMS accelerometers.

To investigate the effects of IA misalignment, begin with the simple case of a constant inertial input $g\hat{z}$ defined as the vertical direction. The gravimeter is tilted at an angle θ from the vertical and is assumed to be perfectly accurate, such that its measurement of specific force is simply

$$a_{atom}(t) = a_a(\theta) = \vec{g} \cdot \hat{z}_a = g \cos \theta. \quad (4.6)$$

For simplicity, assume that all rotation and misalignment are within the xz plane. The i -axis cosensor is defined by its exact scale factor s_i and bias b_i . Assume the x and z cosensors are orthogonal to each other and the z-axis cosensor is misaligned from the gravimeter IA by an angle ξ (figure 4.5). The cosensors then measure specific force given by

$$a_z(\theta, \xi) = \vec{g} \cdot \hat{z}_c = g \cos(\theta + \xi) = g \cos \theta \cos \xi - g \sin \theta \sin \xi \quad (4.7)$$

$$a_x(\theta, \xi) = \vec{g} \cdot \hat{x}_c = g \sin(\theta + \xi) = g \sin \theta \cos \xi + g \cos \theta \sin \xi. \quad (4.8)$$

Tilt angle is arbitrary and may vary, so it is desirable to eliminate explicit dependence on θ by expressing a_a in terms of a_z and a_x :

$$a_a(\theta) = g \cos \theta = a_z(\theta, \xi) \cos \xi + a_x(\theta, \xi) \sin \xi, \quad (4.9)$$




Figure 4.5: 2D model of cosensor misalignment from the gravimeter. The static misalignment between gravimeter and z-axis cosensor is ξ .

or in terms of cosensor raw signal V_i , scale factor s_i , and bias b_i ($i = x, z$),

$$a_a = [s_z V_z(\theta, \xi) + b_z] \cos \xi + [s_x V_x(\theta, \xi) + b_x] \sin \xi \quad (4.10)$$

$$= s_z V_z(\theta, \xi) \cos \xi + s_x V_x(\theta, \xi) \sin \xi + b_z \cos \xi + b_x \sin \xi. \quad (4.11)$$

The above describes the physical relationship of the sensors, but how does this translate into operation of the system? Feedback assumes the z-axis cosensor and gravimeter are perfectly aligned and attempts to apply a bias correction. The assumed relationship of the sensors is simply

$$a_a = s_z V_z + b_z, \quad (4.12)$$

which is clearly not true if a misalignment exists. Instead of b_z , the bias will be the bias correction generated by feedback b_{FB} . Feedback will adjust b_{FB} such that equation 4.12 equates to the real expression given by equation 4.11

$$a_a = s_z V_z + b_{FB} = s_z V_z + [s_z V_z(\cos \xi - 1) + s_x V_x \sin \xi + b_z \cos \xi + b_x \sin \xi] \quad (4.13)$$

$$b_{FB} = s_z V_z(\theta, \xi)(\cos \xi - 1) + s_x V_x(\theta, \xi) \sin \xi + b_z \cos \xi + b_x \sin \xi. \quad (4.14)$$

Note that the correct form is obtained for zero misalignment ($\xi = 0$). From equation 4.14, it is clear that the bias correction b_{FB} carries a significant dependence on the tilt angle θ

and misalignment angle ξ . If the system is static or slowly tilting, the bias correction b_{FB} will match the given expression and the reported specific force will match the gravimeter's measurement, a_a . If θ varies faster than feedback can react, the reported specific force will not be correct, and a time-dependent bias correction will lag behind the delta in specific force.

If ξ is known, an angle correction of the form given by equation 4.11 can be applied to include the off-axis cosensor. Feedback then assumes a relation of

$$a_a = s_z V_z \cos \xi + s_x V_x \sin \xi + b_{FB} \quad (4.15)$$

$$b_{FB} = b_z \cos \xi + b_x \sin \xi. \quad (4.16)$$

Now the bias correction no longer depends on θ , so it will remain constant as a function of tilt.

The misalignment angle can be determined experimentally by performing a series of static tilt tests. The entire sense head is tilted by set angles about an axis in the XY plane. The system is allowed to settle for a period of time and a measurement of raw cosensor and feedback corrected cosensor is obtained at each tilt angle. The raw cosensor is simply the cosensor with a fixed nominal bias applied, while the feedback corrected cosensor has the dynamically updating bias correction applied. As discussed, the feedback corrected cosensor signal when not under motion and subject only to a non-varying specific force will reflect the gravimeter measurement as defined by its IA (not the cosensor's IA). A full set of tilt measurements will result in a pair of cosine curves, one for the raw cosensor and one for the gravimeter-corrected signal (figure 4.6). Zero degrees tilt is defined as the point at which the gravimeter is pointing directly up (as defined by the local gravity vector). The feedback corrected signal will then simply have a tilt angle dependence of

$a_a = g \cos \theta$, while the cosensor will measure

$$a_z = g \cos(\theta + \xi_i) + \delta b, \quad (4.17)$$

where ξ_i is the projection of the total misalignment ξ along the tilt axis and δb is DC bias error. The DC bias error can simply be subtracted out of the data since only a relative measurement of cosensor versus gravimeter is of interest here.




Figure 4.6: A tilt test performed with static tilts about the x-axis of the baseplate. The angle is defined by the gravimeter measurement and the shift of the cosensor measurement is used to calculate ξ_α . A measure of the z-axis cosensor scale factor error $\delta s_z = 1.76\%$ can also be deduced.

Measurements at each static tilt were performed for about 15 minutes to allow sufficient averaging of noise. Time-dependent variations in the cosensor's DC bias will distort the curve, but this variation is on the order of 1 mg. Static tilts are performed in pairs of $+/- \theta_i$ rather than in order from one extreme to the other to avoid time-dependent skewing of the curve. The gravimeter is assumed to perfectly measure specific force along its input axis, so the tilt angle is defined by equation 4.6.

The tilt tests measure projection of the IA misalignment onto the tilt axis, but a

three-dimensional alignment is desired. To measure the total misalignment, the tilt test is performed in two orthogonal axes within the XY plane. The tilt direction is defined by the rotation vector, $\hat{\alpha}$ and $\hat{\beta}$, which is aligned to the inclinometers and the rectangular base of the gravimeter support structure. The two tilt tests then measure IA misalignment ξ_α and ξ_β projected onto the βz and αz planes respectively.




Figure 4.7: Orthogonal tilt tests determine projection of the absolute cosensor axis $\hat{a}$ onto the plane of the tilt test ($\hat{a}_\alpha$ and $\hat{a}_\beta$). The direction of the tilt ϕ_{tilt} is determined by the projection of these vectors onto the xy plane.

From the geometric relations of the tilt angles projected into the $\alpha\beta$ (xy) plane seen in figure 4.7, the total IA misalignment can be calculated as

$$\sin^2 \xi = \sin^2 \xi_\alpha + \sin^2 \xi_\beta \quad (4.18)$$

$$\xi^2 \approx \xi_\alpha^2 + \xi_\beta^2 \quad (\text{small angle}), \quad (4.19)$$

and the azimuthal direction of the tilt ϕ_{tilt} is calculated as

$$\tan \phi_{tilt} = \frac{\sin \xi_\alpha}{\sin \xi_\beta} \approx \frac{\xi_\alpha}{\xi_\beta}. \quad (4.20)$$

The misalignment angle can now be factored into the feedback algorithm per equation 4.16. If the misalignment was accurately determined, the bias correction will no longer be tilt-dependent.

At this point it is worthwhile to explore the effects of inaccuracies in the determination of misalignment. Define the real misalignment as ξ and the determined misalignment ξ' different by $\delta\xi$, such that $\xi' = \xi + \delta\xi$. Feedback will now assume the form

$$a'_a = a_z \cos(\xi + \delta\xi) + a_x \sin(\xi + \delta\xi) \quad (4.21)$$

$$= a_z(\cos \xi \cos \delta\xi - \sin \xi \sin \delta\xi) + a_x(\sin \xi \cos \delta\xi + \cos \xi \sin \delta\xi) \quad (4.22)$$

$$= (a_z \cos \xi + a_x \sin \xi) \cos \delta\xi - (a_z \sin \xi - a_x \cos \xi) \sin \delta\xi \quad (4.23)$$

$$= a_a \cos \delta\xi - (a_z \sin \xi - a_x \cos \xi) \sin \delta\xi \quad (4.24)$$

where a_a is the real gravimeter value. Now assume the biases are near zero, $b_z = b_x = 0$, and the scale factors are 1, $s_z = s_x = 1 \text{ g/V}$. Then the cosensor measurements will take the form

$$a_z = V_z \approx g \cos(\theta + \xi) \quad (4.25)$$

$$a_x = V_z \approx g \sin(\theta + \xi). \quad (4.26)$$

Plugging into the expressions for a'_a and a_a give

$$a_a = g \cos(\theta + \xi) \cos \xi + g \sin(\theta + \xi) \sin \xi + b_{FB} \quad (4.27)$$

$$= g \cos \theta + b_{FB} \quad (4.28)$$

and

$$a'_a = a_a \cos \delta\xi - g[\sin \xi \cos(\theta + \xi) - \cos \xi \sin(\theta + \xi)] \sin \delta\xi + b'_{FB} \quad (4.29)$$

$$= g \cos \theta \cos \delta\xi + g \sin \theta \sin \delta\xi + b'_{FB} \quad (4.30)$$

$$= g \cos(\theta - \delta\xi) + b'_{FB} \quad (4.31)$$

where b_{FB} and b'_{FB} are the biases determined by feedback in each case. Since the gravimeter

is held as truth by feedback, the biases are controlled such that $a'_a = a_a$. Now we can determine the bias shift generated by alignment error:

$$\delta b = b' - b = g \cos \theta - g \cos(\theta - \delta \xi), \quad (4.32)$$

and with small angles

$$\delta b \approx g \left(1 - \frac{\theta^2}{2} \right) - g \left(1 - \frac{\theta^2}{2} + \theta \delta \xi - \frac{\delta \xi^2}{2} \right) \quad (4.33)$$

$$\approx -g \theta \delta \xi. \quad (4.34)$$

So the bias shift due to tilts when the axes are misaligned is proportional to the tilt angle θ and the error in the programmed alignment $\delta \xi$ to first order. How does this compare to the case in which no alignment correction is applied? This is simply the case where $\delta \xi = -\xi$. Then the expression for a'_a reduces simply to $a'_a = a_z + b'$, as expected. And the bias shift is $\delta b \approx g \theta \xi$. Numerically, at a tilt of $\theta = 50 \text{ mrad} = 3^\circ$, an uncorrected misalignment of $\xi = 3 \text{ mrad}$ results in a tilt-induced bias shift of $150 \text{ } \mu\text{g}$. If a correction is applied to within $\delta \xi = 0.1 \text{ mrad}$, the bias shift is within $5 \text{ } \mu\text{g}$. While this level of bias shift is readily apparent in stability results, it has little effect on long term performance. In most situations, a zero-mean average tilt angle over periods of time greater than a minute can be assumed. This oscillation will therefore average away given sufficient measurement time. What is of most importance is that the oscillation is significantly less than a fringe. At $T = 5 \text{ ms}$, half a fringe is $a_\pi = 800 \text{ } \mu\text{g}$, so oscillation of the corrected tilt angle is a fraction of a percent of a fringe.

Up until now it has been assumed that the cosensor scale factors are exactly correct and unvarying. The importance of an accurate scale factor determination can be explored by introducing a fractional scale factor error δs , such that the reported specific force of a cosensor is

$$a'_i = s_i (1 + \delta s_i) V_i + b_i. \quad (4.35)$$

Start with only an error in the z-axis cosensor δs_z . Equation 4.13 becomes

$$a'_a = s_z (1 + \delta s_z) V_z \cos \xi + s_x V_x \sin \xi + b'_{FB}. \quad (4.36)$$

Assume ξ is perfectly known, as terms involving $\delta \xi$ will be much smaller than those involving ξ . Again, invoke feedback to assert that $a_a = a'_a$ by adjusting b_{FB} and make the same assumptions about the behavior of cosensor measurements given by equations 4.25 and 4.26 to obtain an expression for bias shift.

$$\delta b = g \delta s_z \frac{\theta^2}{2} + g \delta s_z \frac{\xi^2}{2} \quad (4.37)$$

$$\approx g \delta s_z \frac{\theta^2}{2}. \quad (4.38)$$

The ξ^2 term can be ignored to first order since misalignment is on the order of a milliradian, and is significantly smaller than the scale of tilt angles.

The same procedure can be followed for an x-axis scale factor error to obtain the expression

$$\delta b = s_x \delta s_x V_x \sin \xi \quad (4.39)$$

$$\approx g \delta s_x \theta \xi. \quad (4.40)$$

At a tilt angle of $\theta = 50 \text{ mrad} = 3^\circ$, a misalignment of $\xi = 3 \text{ mrad}$, and a scale factor error of $\delta_i = 1\%$, the x-axis term contributes a bias shift of $\delta b = 1.5 \mu\text{g}$ while the z-axis term contributes $\delta b = 12 \mu\text{g}$.

Since all these inaccuracy terms are small, they can be assumed to be independent to first order and collected into a single error expression

$$\frac{\delta b}{g} \propto \theta \delta \xi + \delta s_z \frac{\theta^2}{2} + \delta s_x \theta \xi. \quad (4.41)$$

The dependence of bias to errors given by equation 4.41 provide another way to evaluate bias shift through tilt tests. By plotting bias shift versus tilt angle, the IA misalignment ξ can be evaluated (figure 4.8). After correcting for IA misalignment, remnant angle-dependent bias shifts can be seen, with the linear term being a combination of off-axis scale factor error and residual misalignment error, and the quadratic term being due to z-axis scale factor error.




Figure 4.8: The bias shift as a function of tilt angle provides a measure of error terms given in equation 4.41. The large slope in blue is from the tilt test shown in figure 4.6 and the small slope in red is after applying the misalignment correction. The nonlinearity is due to the quadratic effect of δs_z .

4.3 Vibration Suppression

A critical component to consider when operating a sensor in the field is how to handle vibration. The goal is to measure local gravity, but any other specific force will couple into the measurement. Since gravity is slowly varying due to location (on the order of 10s of seconds), it is desirable to remove any high-frequency acceleration that may exist due to coupling through the platform. Vibration suppression in the system takes two forms – passive (hardware-based) and active (software-based). Passive vibration isolation

is achieved by the inclusion of rubber shock mounts installed in the three legs of the support structure. The shock mounts, LORD 200XP-90, have a mount resonance under optimal load of roughly 20 Hz (see figure 4.9), meaning any higher-frequency vibration is suppressed with a roll-off beginning at the resonance frequency. Any inertial input at or below the resonance will transmit through the shock mounts and will therefore be measured by the gravimeter.




Figure 4.9: Measured transmissibility of the shock mounts. A resonance is observed near 20 Hz above which the suppression improves quickly as a function of frequency. At single-digit Hz, the “thumper” does not generate much vibration, so the result is unreliable. In reality, the shock mounts allow full transmission below the mount resonance.

Low-frequency acceleration due to platform motion will still be measured, and any such specific force along the gravimeter input axis will be conflated with the measure of gravity. A simple approach to handle platform motion is to take repeated measurements and average them together. If motion such as rocking or surging occurs at a specific frequency, it will average down over time with a slope of t^{-1} . For example, fairly aggressive ship motion of period 5 seconds producing a heave of 0.2 g will average down to 2 mg after 500 seconds (roughly 1 mg over 10 minutes).

A more effective way to handle platform motion is to use an IMU to inform tilt angle and acceleration. With short term information about tilt and heave, the gravimeter’s specific force measurement can be used to calculate gravity along the vertical direction.

Despite passive vibration suppression, significant higher-frequency vibration may still couple through to the gravimeter. With $T = 5$ ms, a half interferometer fringe is $a_\pi = 800 \mu\text{g}$. Any vibration at this scale or larger will produce a fringe hop (see section 2.4), creating ambiguity in the measurement of $2\pi n \rightarrow n \cdot 1.6 \text{ mg}$. Active vibration suppression is thus implemented through feedforward, as introduced in section 2.4. Taking advantage of the MEMS cosensor’s high bandwidth of 1500 Hz, vibration can be subtracted from the interferometric phase measurement during the three-pulse sequence.

Vibration suppression is then evaluated as the ability of the cosensor to measure high-frequency acceleration and apply a correction to the atom measurement. Feedforward performance is tested by applying a known vibration to the platform and measuring residual phase variation in the interferometer. A “thumper” was constructed to generate vibration by taking two loudspeakers and mounting them to the baseplate. The speakers are driven in phase at various frequencies and amplitudes. The frequency content in the residual atom phase variation can be mapped out to determine feedforward vibration suppression.




Figure 4.10: Example of vibration testing at 18 Hz drive. (Left) The thumper generates many resonances which are measured at a 39 Hz repetition rate. (Right) The peaks can be unfolded to calculate a suppression factor at each resonance. The suppression of the shock mounts and feedforward are measured to determine total suppression as a function of drive frequency. The resonances above 110 Hz are below the noise floor of the atom measurement, so they are estimated as constant above this level.

Total vibration suppression is a combination of passive shock mount absorption and active feedforward cancellation. The shock mounts are ineffective at frequencies near or below the mount resonance but have a good high-frequency suppression factor. Conversely, the cosensor does a good job at handling low-frequency content but is limited by its bandwidth of 1500 Hz. The combination therefore is effective in handling any vibration threats to the system. Improvements in vibration suppression enable more sensitive measurements of gravity in dynamic environments and is thus of high interest. One avenue of improvement is to improve the shock mount system by redesigning the support structure and using softer shock mounts. This effectively lowers the mount resonance and therefore the frequency at which suppression rolloff occurs.

Another avenue of improvement to vibration suppression is through feedforward, either through improvement to the actual cosensor or development of a better feedforward algorithm. A higher-bandwidth cosensor enables the measure and cancellation of higher frequency content, but real sensors usually have a tradeoff of bandwidth versus measurement noise. Cosensor noise is also a factor in the application of feedforward, since any noise generated in the cosensor measurement gets imprinted onto the interferometer phase. One approach to dealing with the limited cosensor bandwidth is to introduce a projected phase correction, in which the measurement of the cosensor is extrapolated forward by a characteristic amount of time to account for measurement lag due to bandwidth. This does not fundamentally allow higher-frequency measurement, but can account for bandwidth-induced delay in the signal. Such improvements are a focus of current and ongoing work to improve field performance of the gravimeter.

4.4 Measurement Stability

One of the key characteristics of the atom gravimeter that makes it an appealing sensor is its long term stability. Given the fundamental basis of the measurement discussed in section 2.2 that relies on unchanging properties of the atom, the atom gravimeter

has great potential to produce a reliable measurement of gravity that does not change over long periods of time. In reality there are other sources of phase variation in the interferometer that can produce instability of the intended measurement. Significant effort has been made in suppressing or accounting for these phase shifts to produce the most stable measurement over long periods of time.

The primary analysis tool used to evaluate stability is called an Allan deviation (or Allan variance) [89]. The Allan deviation quantifies how accurate $\sigma(\tau)$ a measurement is when taken continuously over a period of time τ . A typical Allan deviation curve can be broken into three general sections: short-term noise, long-term drift, and a medium-term flicker floor between the short and long term. The short-term measurement typically sees an improvement (decreasing σ) as a function of increasing measurement time τ due to averaging over noise. For example, white noise has a log-log slope of $-1/2$. Long term drift can be seen as a worsening of the precision (increasing σ) as a function of increasing τ due to instabilities that may cause the signal to fluctuate over long periods. For example, a random walk has a log-log slope of $+1/2$ and a linear drift has a log-log slope of $+1$. In between the short term improvement in σ and long term drift, there is typically a minimum bucket at which averaging over the signal has improved the precision but before long term drifts start to degrade the result. More details on the Allan deviation calculation can be found in [89, 90, 62].

Figure 4.11 is an Allan deviation representative of typical cosensor behavior (red trace) and hybrid atom gravimeter correction (blue trace). Short term noise on the cosensor averages down over time, starting at a shot-to-shot white noise of around $100\text{ }\mu\text{g}$. Random noise averages down with a log-log slope of $-1/2$, and is typically characterized by the 1-second value of $4\text{ }\mu\text{g}/\sqrt{\text{Hz}}$. After about 10 seconds, the cosensor reaches its noise floor and bias drift begins to limit the precision of the signal. The signal continues to drift over time, and the cosensor never reaches a precision at or below $1\text{ }\mu\text{g}$.

At roughly 60 seconds, the gravimeter bias correction begins to take effect. The




Figure 4.11: A typical Allan deviation of an overnight run in the lab. The hand off by feedback of the cosensor to the gravimeter can be seen around 60 s, at which point the signal continues to average down until a noise floor of 200 ng.

gravimeter has no effect on the cosensor signal at shorter periods of time because of the slow response of the Kalman filter. This is intentional to prevent the coupling of short-term noise of the gravimeter onto the cosensor, which has better short-term noise characteristics. The signal continues to average down until around 1 hour, at which point the gravimeter's noise floor of about 200 ng is reached. Factors that limit the noise floor to this level are a subject of continued investigation.

An important tool in suppressing phase shifts is k-reversal (or state reversal). The input axis is defined by the direction of the effective k-vector, and the resulting phase due to position at a Raman pulse is $\phi_g \sim \vec{k}_{eff} \cdot \vec{r}$. The sign of the k-vector is determined by the direction of the photon momentum transfer to the atoms at the first beamsplitter pulse, which can be chosen based on modulation of the two Raman beams. The carrier frequency of one beam and the 3 GHz sideband of the other is set such that the frequency difference is ω_{HF} . Due to the modulation of the AOD, the opposite beams' sideband/carrier pair is detuned by 54 MHz, so interactions of this pair are insignificant. By adjusting the modulation frequency, the resonant pair can be switched (e.g. the upward carrier and

downward sideband are resonant, versus upward sideband and downward carrier). This reverses the direction of the momentum transfer, thus flipping the sign of the momentum-sensitive term. The position-dependent phase then carries a sign based on the direction of the k-vector, $\phi_g \sim \pm k_{eff}z$. Many phase shifts in the interferometer from non-position dependent sources are common mode, meaning they do not depend on the direction of the k-vector. Common mode phase $\bar{\phi}$ and phase difference $\delta\phi$ can then be defined as

$$\bar{\phi} = \frac{1}{2}(\phi_{+k} + \phi_{-k}) \sim \text{common mode phase shifts} \quad (4.42)$$

$$\delta\phi = \phi_{+k} - \phi_{-k} \sim k_{eff}g + \text{direction dependent phase shifts.} \quad (4.43)$$

The k-vector flip is achieved in the experiment by simply changing the EOM modulation frequency. Since the Raman beams are deflected into different paths through the AOD box, the center frequency of the two beams differs by 27 MHz. The RF modulation is set such that the difference between one beam's center frequency and the other's sideband is exactly the clock frequency, 3035 MHz. As seen in section 4.1, the width of the transfer peak is on the order of 100 kHz, so although both beams are modulated, the difference in frequency of the unintended center frequency and sideband is well beyond the frequency range in which Raman interaction can occur. Now by changing the modulation frequency by 54 MHz, the interaction can only occur in the other sideband and center frequency. Since each beam is incident on the atoms in opposite directions, this swapping of the effective sideband flips the direction of k_{eff} and thus the sign of the specific force measurement.

In normal operation, k-reversal is implemented by flipping the k-vector every other ABCD cycle. In other words, the four-point ABCD dither is run with $+k$, then flipped to $-k$ and runs the four-point ABCD dither, then back again. The state-tracking Kalman filter tracks the two phases separately (ϕ_{+k} and ϕ_{-k}) and updates the corresponding one at the appropriate experiment cycle. The Kalman filter then outputs $\bar{\phi}$ and $\delta\phi$ as defined above which are handled separately. Since $\bar{\phi}$ contains unwanted phase shifts, it is held at




Figure 4.12: Demonstration of $\bar{\phi}$ rejection by measuring corrected cosensor across a manually-induced $\bar{\phi}$ shift. (Left) The original KF algorithm induces a shift of 2ug/rad. (Right) the new KF algorithm rejects $\bar{\phi}$ at better than 30ug/rad.

zero by a simple integrating feedback loop. $\delta\phi$ is then sent to the bias correction Kalman filter, which converts it to an effective residual gravity and generates a bias correction.

Figure 4.12 demonstrates the effect of k-reversal – a common mode shift of 10 degrees (180 mrad) is induced, which is readily measured in the response of $\bar{\phi}$. At $T = 5$ ms the scale factor is 250 ng/mrad, so this would be a shift in measured gravity of around 50 μg . The original feedback algorithm did not independently track the two phases nor lock $\bar{\phi}$ to zero, so a significant shift of $\delta\phi$ (and therefore specific force measurement) is observed. After updating feedback to correctly handle $\bar{\phi}$ and $\delta\phi$, no discernible response can be seen in the feedback-corrected cosensor output down to at least the 100 ng level. Common mode phase shifts are therefore suppressed by at least a factor of 500 (as can be measured, likely even better). $\bar{\phi}$ variation in the lab is observed at the scale of one or two degrees over many hours to days, which with $T = 5$ ms and a suppression of at least 500X bounds such effects to < 20 ng.

To demonstrate longer term stability, a month-long measurement was performed in the lab in which the gravimeter was run continuously without interruption. The resulting Allan deviation (figure 4.13) demonstrated a stability bounded to 150 ng over the 28-day

period and suggests the system could be expected to continue performing as such over even longer periods. One hindrance in the test was a loss of atom number due to degradation




Figure 4.13: Month long stability measurement. A stability floor of 150 ng over days to weeks is realized. A correction for the loss of atom number over the course of the test is applied to achieve the ultimate stability.

of the MOT optics (see section 4.6). This degradation has since been addressed, but unfortunately led to a long-term trend in the gravimeter output during the month-long test. A post correction based on a correlated intensity monitor was used to subtract the drift due to atom number loss. This trend is likely an artifact due to trending signal size on top of a background that led to an incorrect determination of the normalized population ratio. Such an effect does not represent an inherent instability in the gravimeter, but is a practical implementation issue that is under consideration. One effort to address this issue is the implementation of a background subtraction algorithm that periodically turns off the MOT light (roughly once every few hours) to let the atoms disappear and measure the background quad detector signal, which may be a combination of fluorescence of hot atoms and light scattered off the edges of the MOT chicane. By periodically measuring background, long term trends in signal to background ratio can be accounted for. One mechanism that produces such shifts was observed in the quad detector’s dark current,

which was found to change by as much as 100% due to 10 °C temperature shifts.

An important consideration in the stability of the gravimeter is the stability in the timing of the Raman sequence. The FPGA implements a reliably repeatable timing sequence, but depends on a reference clock to produce the real-world outputs. Variations in this clock can therefore produce variations in the timing, which is most important in the determination of the interferometer dwell time T . Crystal oscillators have excellent short term periodicity but can suffer from long term drift, so variations in T can occur over long (≥ 1 hour) periods. Variations due to such a drift in T cannot be known exactly and therefore produce a scale factor shift, which manifests as a shift in predicted gravity:

$$\phi_g = k_{eff} g T^2 = k_{eff} (g + \delta g) (T + \delta T)^2 \rightarrow \frac{\delta g}{g} \sim 2 \frac{\delta T}{T}. \quad (4.44)$$

A fractional change in the timing is directly proportional to a fractional change in measured gravity, so clock stability is essential to measurement stability. The FPGA reference clock was changed from a small PCB-mount oscillator to a chip scale atomic clock (CSAC) to produce a 100 MHz reference with a fractional stability floor of $\sim 10^{-11}$ over a day or longer.

Another potential source of drift is laser frequency variations. Measurements of the LCM lock found a stability of no worse than 10-20 MHz. Since this drift is common mode, it does not affect the detuning of the Raman pulse. Instead it manifests only as a variation in expected scale factor through k_{eff} . At a wavelength of 780.2 nm, this represents a variation of around 25-50 parts per billion. It is assumed that frequency generators are stable to the kHz range, so are not significant contributors to overall frequency fluctuation.

A large source of variation in the interferometer phase is the AC Stark effect. The Raman beams produce an intensity-dependent shift in the energy levels of the atom, which is seen as an effective change in detuning of the light from the atoms. This frequency shift

results in a phase shift

$$\phi_{AC} = \frac{\Delta_{AC}^{(3)}}{\Omega_{eff}^{(3)}} - \frac{\Delta_{AC}^{(1)}}{\Omega_{eff}^{(1)}} \quad (4.45)$$

where $\Delta_{AC}^{(i)}$ is the frequency shift of the i -th pulse [52]. The AC Stark detuning is proportional to optical intensities of the two Raman beams as $\Delta_{AC}^{(i)} \sim I_A + I_B$, while the effective Rabi rate is proportional to the beams as $\Delta_{AC}^{(i)} \sim \sqrt{I_A I_B}$. Without proper care, variations in optical power can induce AC Stark phase shifts in the order of a few degrees, or ~ 100 mrad. For $T = 5$ ms, this is a variation in measured gravity of 10s of μg .

A few tricks are implemented to effectively remove AC Stark phase shifts. First, the optical power of the beams are locked via PID loops using intensity monitors at the sense head. This prevents large variations in optical power that would induce unstable Stark shifts. Secondly, ϕ_{AC} is a common mode effect, so k-reversal subtracts it from the gravity-sensitive $\delta\phi$. Finally, adjusting RF sideband power can cause the collective shifts from all sidebands in a beam to cancel. The RF power is then locked via an RF monitor




Figure 4.14: AC Stark scan using fourth pulse. The phase shift due to toggling each Raman beam of the detuned fourth pulse on/off is measured at various RF powers. The RF power is chosen at which there is no phase shift, effectively removing phase sensitivities due to optical power variations. The AC Stark is optimized separately for both +k and -k.

to prevent fluctuations. With single-beam compensation, intensity locks, and k-reversal,

AC Stark effect is entirely removed as a contributor to stability (and accuracy error). Figure 4.14 demonstrates the characterization of AC Stark shift using a fourth Raman pulse. The pulse is detuned such that no coherent interactions occur and is turned on and off to determine AC Stark phase shift as a function of sideband power. The point at which the phase shift is zero is the point at which AC Stark compensation is optimized, thereby preventing optical power fluctuations from producing phase shifts in the interferometer.

It is worth noting that “solid earth” tides produce real oscillation in local gravity which may need to be accounted for. The gravitational pull of the Sun and Moon minutely warp the Earth’s crust as the Earth rotates. This warping produces variations in effective distance of the surface from the planet’s center of mass and therefore changes in gravity at the surface. The tides occur at a period of 12 hours and reach an amplitude of ~ 50 ng. Given the gravimeter’s stability floor of 150 ng, land tides are not an issue until better stability is achieved, at which point they will need to be factored into the expected value of gravity.

The stability contributors discussed in this section were investigated over the past six years of the project. Many of them had not yet been addressed by the time of the at-sea deployments presented in chapter 5. Figure 4.15 demonstrates the substantial improvement to stability achieved in the period following the field demonstrations. One of the causes of the worse performance was the choice to operate at a smaller T to account for ship motion, which significantly decreased the sensitivity of the gravimeter. A particular source of instability was that the CSAC had not yet been replaced, so clock drift could have been significant. Additionally, k-reversal had not been implemented, and the initial implementation of feedback (discussed previously) produced a roughly 1/10X bleedthrough of common mode effects. To prove this was not a fundamental limit to the system’s performance in dynamics, further land-based excursions discussed in chapter 5 demonstrate significantly improved performance in dynamic environments following the aforementioned upgrades.




Figure 4.15: Stability comparison at the time of the Hawaii excursion versus more recent performance. The SNR of the gravimeter is significantly improved as well as its long term stability floor.

Further investigations into sources of instability are continuing into the current day with the goal of achieving best possible long term stability. Such investigations involve potential residual effects of phase shifts that may somehow couple into $\delta\phi$, either through fundamental effects or through implementation errors. Central to such possibilities is the accurate application of k-reversal, since any errors in the calculation in feedback provide a means for the various common-mode mechanisms to couple into the specific force measurement. The artificial $\bar{\phi}$ shift test demonstrated earlier in this section significantly bounds direct coupling. Inaccuracy in the phase-tracking Kalman filter could introduce discrepancies in the calculated interferometer phase which could allow common mode phase shifts to couple into $\delta\phi$ in a way that would not be observable in the artificial $\bar{\phi}$ shift test. A better way to determine real response to such effects would be to directly shift the physical parameter and observe the resulting phase shift. For example, the Raman bias field amplitude could be varied to observe the resulting phase shift and look for residual bleedthrough onto $\delta\phi$ (and therefore the inertial measurement). Such tests can be performed generally for any given control parameter of the system and serve as an

experimental investigation into sources of bias shifts or instabilities. The ability of the feedback algorithm to measure the cosensor and compare it to the atom gravimeter can also be generally questioned at very small values. In other words, is there a limit to how small a fluctuation can be measured and effectively controlled. The shot-to-shot noise of the cosensor is around 40 mg, so fluctuations at or below around 5 ppm are of interest. Such investigations into sources of instability are on-going.

As discussed in the introduction, one of the primary benefits of the atom gravimeter is its long term stability. Stability of gravity measurements not only over days or weeks, but months or even years is of interest. Given the continuous development and improvement of the sensor, it is difficult to characterize such long term stability. The absolute measurement has been shown to vary no more than $\sim 10 \mu\text{g}$ over many months of periodic measurements. It is likely that the true stability is even better, with many of the changes being due to phase shifts introduced in various upgrades or changes to the operating conditions of the system.

4.4.1 Short-Term Precision

While long-term stability is of great importance to the gravimeter, short-term precision is also of interest. The better the short-term noise, the more quickly an accurate measurement can be achieved. Being able to quickly produce an accurate measurement effectively improves the sensitivity of the gravimeter while in motion. For example, when comparing measurements to a gravity anomaly map while moving with some speed, each measurement must be made with sufficient precision before the position changes to a location with a significantly different value of gravity.

Short-term noise is dominated by the cosensor, since MEMS sensors can achieve fast repetition rate measurements at relatively low shot-to-shot noise. The gravimeter is rather noisy with a 1-sec precision of around $50 \mu\text{g}/\sqrt{\text{Hz}}$, compared to the cosensor's $4 \mu\text{g}/\sqrt{\text{Hz}}$. The SNR of the gravimeter is unlikely to improve enough to bring the short-term noise

below that of the cosensor, but it matters for the performance at medium-length time scales. Figure 4.15 demonstrates an improvement to SNR from 2022 to 2024, which shifts the medium term trend from top-right towards bottom-left. The “hand-off” of the hybrid specific force measurement from the cosensor to the gravimeter must occur after the gravimeter’s precision reaches below the cosensor, which can be seen as the separation of the two Allan deviations (e.g. in figure 4.11). Allowing the feedback correction to act too fast introduces a period of decreased precision of the hybrid signal as seen in the blue trace in figure 4.15. The gravimeter SNR also determines how quickly the ultimate noise floor is achieved, as it sets the short-term $(-1/2)$ trend below the cosensor’s $\sim 1\ \mu\text{g}$ noise floor.

4.5 Absolute Accuracy

The previous section describes stability and precision, but does not address absolute accuracy of the measurement. It is possible to have a large inaccuracy in the measurement while remaining very stable. Indeed, a bias on the order of 10s of μg is observed in lab measurements. Accuracy is determined by comparing the measured value at the noise floor of the gravimeter to local gravity surveys performed in the building. Gravity surveys were independently performed by the National Geospatial-Intelligence Agency (NGA) at nearby locations at the same elevation in the building in 2012 and reported a value of $980387.07 \pm 0.04\ \text{mgal}$ ($-283.41 \pm 0.04\ \mu\text{g}$ from $1\ \text{g} \equiv 9.806\ 65\ \text{m/s}^2$) [91]. While a precision/stability of as good as $200\ \text{ng}$ can be achieved, a much larger bias is present. Since the hybrid feedback signal ultimately reports the value of gravity as determined by the gravimeter, the inaccuracy is attributed to the gravimeter, and is thought to be produced by a number of systematic errors in the interferometer sequence.

A possible source of phase shift is magnetic field effects. The magnetic shields attenuate external sources of fields, but some residual field could potentially penetrate through to the atoms. More significantly, the intentionally-driven fields such as MOT quadrupole,

state prep bias, and Raman bias fields have transients that induce eddy currents through the internal vacuum structure. These transients result in exponential field settling times on the order of 1 ms. The atoms are prepared in the magnetically insensitive $m = 0$ ground state, which suppresses the linear Zeeman effect, but the quadratic Zeeman effect still produces shifts. (Note that fields are on the order of 1 G, so Zeeman shifts are in the weak field regime). The atoms serve as a highly sensitive probe of the magnetic field, and as such can be used to map out the shape and magnitude of the magnetic field by shifting the interferometer in time. [92] provides the relation of phase shift due to magnetic fields. Although a useful method of characterizing the field in the experiment, these phase shifts are undesirable when performing measurements of acceleration. A constant phase shift across the LPAI sequence inherently gets subtracted, so the only phase shifts of importance are those that are varying during the interferometer. Measurements have demonstrated phase shifts due to field on the order of 1 rad. A simple solution to the transient field issue would be to wait a few ms until the field has settled, but this would sacrifice repetition rate. Instead, k-reversal is used to subtract out the effect. Since the phase shift is independent of $\mathbf{k}_{eff}$, it gets canceled by the k-reversal sequence. Tests have demonstrated subtraction of the roughly 1 rad phase shift to below 1 mrad, the resolution of the measurement, or a factor of 1000X or better suppression. Future engineering work may attempt to reduce the field transient by modifying the structure to prevent eddy currents and to implement a field driving scheme that accounts for eddy currents by an optimized overdriven current scheme.

The FPGA reference clock error described in the previous section not only produces instability in the measurement due to clock drift, but can also produce an absolute bias shift. An offset in the real dwell time from the expected dwell time produces a scale factor error which manifests as a shift in the value of gravity reported from the interferometer phase. The FPGA implements timing by converting real time values to number of clock cycles. With a 50 MHz reference clock, a clock cycle is 20 ns. Errors in the rounding

and precision of timing values being converted into clock cycles was found to produce a single clock cycle offset in the commanded dwell time. The resulting bias shift is given by equation 4.44, which for $T = 5 \text{ ms}$ is $8 \mu\text{g}$. By more carefully rounding and converting timing values, this offset was corrected.

A similar type of scale factor offset exists due to the real pulse lengths of the Raman pulses. The theory discussed to this point considers the Raman pulses to be infinitesimal, but in reality they are on the order of $\tau = 5 \mu\text{s}$, or 0.1% of the dwell time. This finite pulse effect modifies the effective length of the interferometer which results in a phase shift

$$\phi_g = k_{\text{eff}} g \left[\left(\frac{2}{\pi} - \frac{3}{4} \right) \tau^2 + T^2 \right]. \quad (4.46)$$

Finite pulse effects are noted in [67, 93] and are arrived at by performing a sensitivity function analysis [94], which is a more careful and formal accounting of the gravimeter sensitivity pictured in section 2.2. It is worth noting that both [67] and [93] arrive at a modification to the scale factor that has a linear τT term, but careful analysis of the definitions of interferometer timing and frequency chirp implementation in feedforward lead to the form given by equation 4.46. A significant part of the large bias shift seen in the first year of operation (see figure 4.16) can be attributed to an incorrect implementation of the finite pulse correction that included a linear term, which produced a phase shift on the order of $1000 \mu\text{g}$. The quadratic term produces a correction to the atom scale factor that produces a shift on the order of 100 ng in this experiment.

The above analysis of finite pulse effect considers the pulses as square waves, but as seen in section 4.1, the pulses are rounded and have an asymmetric shape. The shape must be considered when defining the pulse width, and an integral definition is chosen

$$\tau_{\text{real}} = \frac{1}{\max(V)} \int_{t_1}^{t_2} V(t) \, dt, \quad (4.47)$$

where $V(t)$ is the time trace of the pulse and t_1 and t_2 define a window around the

pulse. The pulse center is similarly defined as the time at which the integral is half the total value. The real pulse sizes and positions are calculated by recording time traces at various commanded τ_{set} and calculating the real pulse shape and center. These values are implemented by creating a lookup table that relates the set pulse to the actual pulse in the FPGA software to ensure the real pulse shapes are being accounted for in the timing.

The AC Stark effect described in the previous section also can produce a bias error. While fluctuations in beam intensity induce variations in the phase, the absolute shift would contribute to the reported gravity. As mentioned, AC Stark is compensated for by tuning RF and optical power. Additionally, k-reversal subtracts contributions to the gravity sensitive term. Contributions from AC Stark are therefore theoretically zero, and experimental tests have bounded the effect to below 200 ng (the noise floor of the measurement).

A special case of the AC Stark effect known as the two photon light shift (TPLS) arises from the presence of off-resonant Raman field pairs. An analogy of this effect to the Bloch-Siegert shift in NMR [95, 96] can be drawn, which provides an intuitive understanding of the TPLS as a shift of the rotating wave approximation in the presence of a strongly driven system [97]. Unlike the AC Stark effect described in section 4.4, the TPLS carries a sign, which means it does not get canceled by k-reversal [98]. Fortunately, the racetrack design of the Raman beams prevents this effect. In [98] and other studies of this effect, the commonly used retro-reflection construction of the beams produces off-resonant light in both directions. Since these beams are not present in the racetrack implementation, there is no mechanism through which the TLPS can arise and it is therefore suppressed. Some residual effect may be present due to polarization leakage that allows for a small amount of cross-coupled light, but this effect is at the sub-20 μg level.

The above sources of bias shift are many of the mechanisms considered thus far. Absolute accuracy is experimentally tracked over the history of the gravimeter's life by repeated overnight stationary measurements of gravity in the lab. Figure 4.16 plots

Error source	Estimated size	Description
AC Stark	$\sim 20 \mu\text{g} \rightarrow 0$	Compensated
Clock fencepost	$8 \mu\text{g} \rightarrow 0$	Fixed
Finite pulse effect	$< 50 \text{ ng} \rightarrow 0$	Accounted for
Raman beam angular misalignment	$< 1 \text{ ng}$	See section 3.1
Two-photon light shift	$< 500 \text{ ng}$	Suppressed by racetrack
Gravimeter/cosensor alignment	0	Suppressed by feedback
Bias field	0	Suppressed by k-reversal
Wavefront curvature	?? ng	Under investigation
Parasitic DI interferometer	$\sim 10 \mu\text{g}$	Under investigation
Total (current)	$-7 \mu\text{g}$	Measured at $T = 5 \text{ ms}$

Table 4.1: Some of the main sources of gravimeter bias offsets considered. Estimated magnitudes are based on testing and calculations at $T = 5 \text{ ms}$. Note that a value of “0” in this table implies that the effect has been theoretically removed, though the ability to measure residual effects are limited to the gravimeter’s noise floor of 150 ng.

absolute gravimeter shift observed during laboratory experiments from July 2022 to April 2025. This plot is not representative of long term bias stability, but rather reflects continued reduction in absolute error resulting from system improvements implemented over that time period. The initial 1 mg offset was due to an incorrect determination of finite pulse effects, which introduced an erroneous τT term to the scale factor. After correcting this offset, the bias sources listed above became subjects of investigation, and were eliminated or addressed as they were understood. The current bias offset measured by the gravimeter is roughly $-7 \mu\text{g}$ from the surveyed value.

Further investigations into absolute accuracy are underway, given the requirement of gravimeter accuracy for its effective use as a navigation aid. The search for fundamental phase shifts in the interferometer are of interest for an understanding of the physical effects that may influence the measurement, such as wavefront curvature and other light shifts. Wavefront curvature of the Raman beams is of particular interest and is currently being studied – ripples in the Raman beam cause spatially-dependent Rabi rates that can affect Raman transfer of the atoms and manifest in phase shifts of the interferometer. Initial calculations suggest the effect is small in this system due to the carefully designed




Figure 4.16: Evolution of absolute accuracy over the course of the experiment. The recent 1.5 years see a bias typically within $20 \mu\text{g}$ of the expected value.

optical path, but these results are not yet finalized. Algorithmic/engineering effects such as the handling of bias fields and timing accuracy are also important considerations, as they can confound investigations into more fundamental effects.

Recent investigations have identified a likely source of interferometer instability and inaccuracy. A parasitic interferometer is present due to cross-coupling of the polarized Raman beams. This cross-coupling creates two pairs of beams traveling in the same direction that are Doppler insensitive (DI). The DI transition can be seen by taking detuning scans at different pulse lengths, as seen in figure 4.17. The cross-coupling is




Figure 4.17: The DI Raman transition measured by performing detuning scans at different pulse lengths.

on the order of 1%, which is enough to manifest as a phase shift when measuring the intended main Doppler sensitive interferometer. Although the investigation is still in the preliminary stages, recent tests suggest it could be a large source of the remaining inaccuracy of the LPAI gravity measurement, as well as a source of instability.

Another tool with potential benefits for accuracy is the lower detection window. Dropping the atoms long enough to detect them at the lower window would allow a significantly longer interrogation time and thus a significantly larger scale factor. At a distance of 50 mm below the initial MOT location the atoms would free fall for around 100 ms, allowing for a dwell time of up to $T = 50$ ms, or a scale factor of 2.5 ng/mrad (as opposed to 250 ng/mrad for $T = 5$ ms). This highlights another great benefit of the quantum gravimeter – the sensitivity can be adjusted simply by changing the timing of the sequence and requires no adjustments to the hardware, allowing for easy and immediate changes of the sensitivity. A more mature gravimeter could foreseeably dynamically adjust its scale factor based on the dynamic environment, allowing for best possible sensitivity given the vibration and motion of the platform.

Efforts to use this lower window have thus been minimal as much of the work to this point has been focused on the engineering and physics efforts at shorter dwell times. Detection was demonstrated briefly at the lower window, with atom fluorescence and Raman population transfer observed. Further efforts to use this lower window are planned, with the goal of demonstrating successful interferometry and the resulting increased sensitivity to gravity. A practical challenge is the repetition rate at which the measurement at longer drop times can be achieved. The current MOT loading time of 10-20 ms highly depends on recapture of the atom cloud from the previous experiment cycle, as loading a full MOT directly from the source beam takes on the order 1-5 seconds. The atoms fall out of the MOT beam region after approximately 30 ms, so longer drop times do not allow for recapture. Such long drop times therefore only allow for a repetition rate on the order of single digit seconds, which greatly decreases short-term precision.

Finally, an improved cosensor may be necessary to achieve feedforward of the precision required for long drop times/large scale factor.

4.6 Optics Degradation

After two years of operation, a downward trend in the atom number was observed (see figure 4.18). Several root causes were investigated and multiple “fixes” were implemented. These fixes, while able to temporarily improve the atom number, did not arrest the downward trend of the atom number. Eventually, the source of the degradation was




Figure 4.18: Relative atom number over the duration of the experiment as measured by the fluorescence detector. Loss of atom number over the course of many months is observed over multiple periods. The sudden increases of atom number is due to the various fixes discussed in this section, with the final vacuum repair occurring at the end of 2024. The apparent decrease in atom number around March 2025 is not due to degradation of the system, but rather due to an increase in experiment repetition rate at the expense of MOT loading time.

determined to be internal to the vacuum chamber. The optical power transmit through the MOT chicane and returning through the retro reflected path was observed to be only 14% of the input power. After beginning to observe atom number loss, this optical attenuation took nearly a year to identify as the root cause because there was no direct monitor

of optical power returning from the vacuum chamber. Desire to keep the gravimeter operational in spite of dwindling atom number delayed further investigation until the signal was too low to produce a discernible measurement above the noise level.

A few basic ingredients are required to generate a MOT in the system:

- Detuned cycling RF (frequency and amplitude)
- On-resonant repump RF (frequency and amplitude)
- Optical power
- Optical alignment
- Quadrupole field
- Rubidium atoms

Both amplitude of the RF and the correct frequency are needed. To generate the MOT, detuned $F = 3 \rightarrow F' = 4$ cycling light is required to create an optical molasses. Simultaneous on-resonance $F = 2 \rightarrow F' = 3$ repump light prevents atoms from remaining in the dark state for the cycling light. Since repump is on-resonance, it requires proportionally significantly less light to provide sufficient repumping. The RF signal is monitored next to the EOM that modulates the light. Investigations into the RF chain did not yield any issues. The frequency and amplitude remained stable over long periods of time and increasing/decreasing RF amplitude did not improve MOT capture. It was found that, as the atom number decayed, some portion could be recovered by decreasing the detuning of the cooling light from 3Γ to about 1.5Γ . Over time, the atom number continued to decay and this frequency tuning could not account for the gradually decreasing MOT capture.

Both RF signals get modulated on the MOT light that is injected from a single fiber into the fully-retroreflected MOT chicane. Optical power is monitored at the sense head just before insertion into the vacuum chamber. As described in section 3.2, several issues with degradation of the laser generation existed but were successfully addressed. These degradations did decrease atom number, but once addressed did not account for the longer-term downward trend. Upon replacing the failing PPLN, an improvement in

atom number was observed compared to before replacement, but never to the level of further back in time. Eventually, the fiber-collimator input to the MOT chicane at the sense head was investigated and a melted fiber core was observed. The melted fiber tip affected the mode of light distribution out of the fiber and the overall optical power out of the collimator (figure 4.19). The fiber tip was polished and the resulting mode was improved to nominally good performance. Optical power out of the collimator was 35% of optical power out of the fiber before polishing and 75% after polishing. Again, a clear improvement in atom number was observed, but continued to decay after the fiber polish.




Figure 4.19: Optical mode out of the ν_2 fiber before (left) and after (right) repolishing the FC/APC tip. The repolished tip produced a cleaner mode that is better collimated at the sense head.

At this point it was hypothesized that the internal optical alignment of the MOT chicane could be drifting over time, but there was no way to measure alignment without breaking vacuum to investigate the internal optics. The slow and smooth downward trend of the MOT decay over years meant that a sudden change in optical alignment was highly unlikely. The alignment was therefore quite robust against common sources of optical misalignments such as vibrations or shocks to the system. This is further evidenced by its performance across multiple field deployments in which the shock and vibration threat was significantly larger than that of the lab. Loss of atom number did not demonstrate

any reaction to field deployments, instead decaying slowly over the time regardless of environment. The mechanism of misalignment therefore would be limited to slow drift effects such as creep and relaxation of mirror tilts due to thermal cycling of the sense head.

Another potential system failure is in the generation of the quadrupole field required for MOT capture. The quadrupole field is generated by a pair of anti-Helmholtz coils attached to the MOT chicane inside the vacuum chamber. The field is driven by a Coil Driver Board, which is meant to provide a steady and equal current to both coils. The coils are driven independently by the board but the circuit is constructed such that the current is always the same. Changes in coil resistance therefore should not effect the field generated, for example those due to changes in temperature. Indeed, the coils running at about 1 A generate a significant amount of heat inside the vacuum. Without convection to cool them, the temperature of the MOT chicane appears to significantly couple to driving of the quadrupole coils. At room temperature a single coil measures $1.8\ \Omega$ and when the system is running with a 60% MOT duty cycle, the resistance increases to $2.25\ \Omega$. Copper has an electrical conductivity thermal coefficient of $0.393\ \%/^{\circ}\text{C}$, so the $0.45\ \Omega = 25\ \%$ increase in resistance corresponds with a temperature delta of roughly $65\ ^{\circ}\text{C}$ above room temperature.

The coil driver board was found to be limited in its combined MOT coil output to 2.65 W. After factoring in cable resistance, the total combined MOT coils resistance at $85\ ^{\circ}\text{C}$ was measured at $5.9\ \Omega$, or a maximum current of 0.67 A. The anti-Helmholtz coils are designed with an axial transfer function of $59\ \text{mA}/(\text{G}/\text{cm})$, meaning 0.6 A generates a 10 G/cm axial gradient. This should be sufficient for MOT loading, but increasing the gradient could improve capture efficiency if there is not enough cycling light. This suggests MOT field drive could be improved, but does not offer a mechanism of long term atom loss. Resistance measurements made recently match resistance measurements made at the initial construction of the sense head, suggesting there has been no change

over time in coil performance. A turn-to-turn short in a coil could create a change in field, and if it happened in only one could change the quadrupole field center. Resistance measurements confirm no such shorts exist, and moreover, a short would almost certainly be an abrupt change rather than a long term drift in field.

In attempt to improve MOT loading, a set of bias coils were introduced along with the quadrupole field to shift the force balance region of the MOT. The hypothesis was a misalignment in MOT beams could leave the central region imbalanced, thus decreasing the ability to capture atoms. By shifting the field center, the atoms could be loaded where the beams were more balanced. This was found to have great effect, and MOT loading was dramatically improved by the introduction of the bias fields. A fairly large ~ 1 G field was needed to find optimal loading, and the magnitude of this bias field needed to be increased over time as atom number decay continued to be observed. Eventually, the bias coils could not account for atom loss, even when greatly increasing field strength. This suggested that the optical power imbalance was continuing to trend over time. It should be noted that the hypothesis of beam misalignment was ultimately incorrect, despite the temporary success of applying MOT bias coils.

Another potential mechanism in loss of atoms in the MOT is a loss in output of Rubidium atoms by the atom source. This could be due to clogging of the nozzle or evaporation of the limited atoms in the reservoir. The reservoir initially contained about 500 mg of atomic Rb-85, and flux out of the nozzle was estimated to be around 10^{13} atoms/s. Assuming uninterrupted emission from the source at this flux, the predicted lifetime of the source is therefore on the order of $10^{22}/10^{13} \sim 10$ years. It is therefore highly unlikely that all the atoms evaporated away. Additionally, the pressure in the reservoir and therefore beam flux would be expected to be relatively constant as the source evaporated until the end where nearly all the liquid atoms disappeared. Given the slow trend of MOT number decay, source depletion was unlikely. The other source-related issue could be clogging of the nozzle. If the wall temperature of the nozzle was colder

than the surroundings, it could be possible for Rb atoms to partially clog the output. A test was conducted in which the reservoir temperature was increased while keeping the nozzle temperature constant. If the nozzle was clogged, reservoir temperature should have little effect on atom flux because the atoms would still be trapped behind the nozzle. The test found that reservoir temperature did increase atom flux, and it was thus concluded that the nozzle was most likely not clogged. This was later confirmed visually when the vacuum system was opened for repair.

Given the exhaustive investigations of RF, optical power, magnetic field, and atom source output described above, the only remaining culprit was the integrity of the optics inside the vacuum. A measurement of light returning through the retroreflected MOT chicane was finally made by standing off the collimator plate from the vacuum chamber to measure input and output through a beam splitter cube. The input light was irised and the return was found to be aligned with the input. This means optical alignment of the internal vacuum components did not significantly change. Accounting for the beam splitter ratio, the return light was found to be only 14% of the input light, implying that the optics inside the vacuum were absorbing 86% of the light traveling through the MOT chicane. This level of attenuation also means the power balance between input and output MOT beams in each direction were highly unequal, which would significantly decrease the ability to load a MOT.

Finally, the vacuum chamber was opened to directly observe the optics. Cloudiness was immediately observed on high-reflective surfaces of a couple of the optics (figure 4.20). The optics were removed from the vacuum structure for individual measurement, and it was found that many of the high-reflective (HR) surfaces displayed non-uniformities indicative of coating degradations (figure 4.21). The corner cube in the Raman beam path in particular displayed significant non-uniformity in a region near a MOT coil wire that ran close to the HR surface. Many of the optics had anti-reflective coatings, all of which showed no signs of degradation. Non-coated glass surfaces similarly looked unchanged.




Figure 4.20: Cloudy surfaces of internal optics observed upon vacuum break. (Left) The Raman corner cube with cloudiness on the surface concentrated near the MOT coil wires. (Right) The MOT roof prism with even cloudiness across HR surface.

Optical measurements indicated significant loss of the reflection on each of the degraded HR surfaces. Reflectance of each degraded surface was found to be 60-70%, while spare optics not subjected to the vacuum system environment measured very close to 100% reflectance. Attempts to clean the HR surfaces were met with limited success. Applications of methanol, isopropyl alcohol, and acetone removed some of the apparent nonuniformities, but was largely unable to return the optics to its original state. In some cases reflectivity was improved, but in others, it remained largely degraded.

The inability to clean the HR surfaces suggest that they were not being coated by something unwanted inside the vacuum chamber, but instead the coatings were chemically altered. The HR coatings are a stack of thinly applied layers, the topmost of which are either Ta_2O_5 or Nb_2O_5 . Similarly, the AR coatings are a stack of layers, but the topmost is SiO_2 . Chemical analysis was performed on samples of the coatings as well as various surfaces inside the vacuum chamber. Fourier Transform Infrared Spectroscopy (FTIR) looked for organics but found very little evidence on most of the samples. This ruled out outgassing of polymers present in the vacuum, such as polyimide coatings of coil wires and PEEK spacers used in the construction of the optics. Energy-Dispersive X-ray spectroscopy (EDAX) was performed to look for specific elements in the samples. The EDAX measurements revealed evidence of Rubidium on most of the surfaces throughout



Figure 4.21: Degraded HR surfaces of the Raman corner cube (left) and two of the MOT mirrors (right). The surface degradation is illuminated by shining white light on the surface at a particular angle of incidence.

the chamber (HR, AR, non-optical, etc.). Given the prevalence of Rb everywhere and the fact that the HR surfaces could not simply be cleaned, it was concluded that the elemental Rb ejected from the atom source was reacting with the HR coatings hence degrading their reflectance.

Due to the slow nature of this interaction over the past few years, we have not been able to directly replicate or confirm this mechanism as the source of surface degradation, but it nevertheless remains the most plausible hypothesis. The most likely form of reaction is the formation of Rubidium (Tantalum) Niobate, RbNbO_3 (RbTaO_3). Nb_2O_5 and Ta_2O_5 are chemically similar, so it is reasonable that, if one reaction is possible, both may be. Rubidium and Lithium are both alkali metals, so RbNbO_3 and LiNbO_3 should be chemically similar. Lithium Niobate is a well studied material and there is a vast collection of literature available that describes many ways in which LiNbO_3 can be created [99, 100, 101]. One such process is atomic layer deposition (ALD), in which a Niobium Oxide substrate such as Nb_2O_5 , is heated in a vacuum and atomic Lithium is deposited via a matter beam. A parallel to this process of LiNbO_3 generation can be drawn to

our system, in which Nb_2O_5 or Ta_2O_5 surfaces receive atomic Rb to create RbNbO_3 and RbTaO_3 . The process of LiNbO_3 generation is typically done at much higher temperatures, so it is reasonable to assume that the process may occur, but significantly slower, at only slightly elevated temperatures of below 100°C .

Solutions to the surface degradation issue follow two general paths. The first is to remove sources of unwanted elements inside the vacuum. Organics and general outgassing of plastics are known to cause issues in holding an ultra high vacuum. Given the general lack of organics found by the FTIR and that Rubidium is the likely culprit for surface degradation, it is not feasible to try to prevent contaminants from reaching the HR surfaces. Instead, the solution is to coat the HR surfaces with a protective outer layer that prevents contaminants from interacting with the Nb_2O_5 or Ta_2O_5 layer. Alumina (Al_2O_3) is known to be nonreactive and is applied to optical surfaces in other situations for protection [102, 103, 104]. By depositing a 10 nm thick layer on the HR surface, the HR coating can be protective without significantly degrading the reflectivity. Alumina deposition was performed on a fresh set of optics via atomic layer deposition (ALD). ALD creates dense layers, the thickness of which can be controlled as accurately as a single atom layer [105]. The denseness is beneficial to prevent permeation of Rb through the protective layer. At 10 nm thick, a decrease in reflectivity of about 2.5% was expected, and indeed after the ALD deposition, a reflectivity of 97.5-98.5% was measured on each HR surface. With 13 reflective surfaces in the fully retroreflected MOT path (from input to return), the expected compound loss is $(1 - 0.98^{13}) \cdot 100\% \approx 25\%$.

After reassembly of the MOT chicane, the compound loss was measured to be 26%, in close agreement with the predicted loss. This was deemed acceptable and the coated optics were installed in the system. MOT loading was demonstrated to be significantly improved, with lower source temperatures, lower MOT bias fields, and less overall load time required to produce a comparatively large atom number. This allowed the overall experiment cycle repetition rate to be increased from 25 Hz to 40 Hz, which improves




Figure 4.22: New replacement optics with alumina coating on HR surfaces do not exhibit the same surface defects seen in old vacuum optics.

short and medium term gravity measurements. After roughly four months, no sign of degradation or MOT number decay has been observed. Longer term observation of atom number will be required to determine if the optics coatings provide a truly long term solution or if further development will be required to produce long-term robust optics inside the vacuum.

CHAPTER 5

Field Deployment

The deployments described in this section are not presented in chronological order. Due to timelines of funding cycles, the arguably more complex at-sea demonstrations were executed in 2022, shortly after construction of the system had completed. Significant development and improvement of the system occurred in the following years, and performance of the system matured closer to the levels demonstrated in chapter 4 by the time the land-based excursions were performed.

5.1 Mobile Gravity Measurement

Development of the gravimeter occurs primarily in a fixed location in the lab, but its intended purpose is to be able to provide accurate measurements of gravity in the field. To that end, the system was built into a mobile platform that can be transported around the building. This served two primary purposes. The first is to demonstrate its ability to operate outside of the lab. This requires collecting all supporting electronics and fixturing into a compact format that can be carted together while continuing to function without issue. The second purpose is to demonstrate accurate measurements of varying gravity.

5.1.1 Building-Wide Gravity Gradient Measurement

At the surface of the Earth the gravity field gradient is around 3080 Eotvos (roughly $0.1 \mu\text{g}/\text{ft}$), so the approximately 4 meter difference between a single story of the building is around $1 \mu\text{g}$ of difference in Earth's gravity field. This is easily distinguishable by the gravimeter given a long enough measurement interval. For example, the gravimeter can typically reach a noise floor of 200 ng after an hour, so two points separated by a few feet of elevation is discernible by comparing an hour long measurement at each point.




Figure 5.1: The system on the mobile cart for building-wide survey.

To enable easy transportation of the gravimeter and its support systems, a custom aluminum cart was constructed (figure 5.1). The frame was built out of 2 inch thick 80/20 extruded aluminum struts, and had a footprint just large enough to house the main support structure alongside the electronics rack. At 3.5X5 feet, the cart fit through the doors of and into the building's freight elevator. The main structure and electronics rack were bolted into a thin base plate and 80/20 support beams a couple feet up provided

further mechanical constraint. The cart rested upon 4 hard rubber wheels, and 3 legs each contained a screw-adjustable metal foot. The feet could be retracted to allow for wheeling of the cart, and could be extended to provide fine-tune adjustments of tilt angle of the cart.

As previously stated, variations in gravity from floor to floor are on the order of a single- μg , so sub- μg resolution is necessary to provide discernible measurements. The gravimeter reaches a precision of 1 μg after averaging for about 300 seconds. It typically reaches a noise floor of 200 ng after about 7000 seconds, or 2 hours. In order to achieve best




Figure 5.2: Allan deviations representative of each stationary measurement. The Allan deviations of a two hour run cannot directly report the accuracy after two hours, so longer in-lab tests are used to demonstrate the validity of projecting the error to the full 7200 second dwell. Each measurement is similarly evaluated to report error bars.

possible results, measurements at each floor were taken for at least 2 hours. Measurements were performed multiple times on some floors and at a couple different locations across some floors to demonstrate repeatability of the measurements. Each measurement was performed at a stationary location, and the leveling feet were used to ensure the gravimeter was pointing up as defined by the gravity vector. Two orthogonal inclinometers in the xy plane were calibrated to the gravimeter and used to measure the tilt angle to an accuracy of at least 10 mdeg (175 prad). This corresponds to a maximum tilt-induced decrease in

specific force of 15 ng, which is well below the level of desired measurement and precision of the gravimeter. The Allan deviations in figure 5.2 are a representative sample of the precision of measurements at select locations. Variations in the Allan deviations are due to different inertial environments across the building, notably larger vibrations at higher floors. With a dwell time of at least two hours, all measurements were precise to better than 300 ng.




Figure 5.3: Building-wide gradient test result. The resulting gradients calculated from both pre- and post- power cycle are close to the expected gradient of 3080 Eotvos.

The measurements of gravity at each location were plotted and the linear slope was calculated to determine the local gravity gradient (figure 5.3). Given the length of each measurement, the test was performed over a few consecutive days. Part way through the campaign, the system was inadvertently power cycled while being transported to a new measurement location due to a failing UPS battery. Power was quickly restored and operation resumed, but a shift in the absolute measurement appeared. The goal of the test was primarily to demonstrate the relative gravity as a function of elevation, so this absolute shift did not corrupt the result. In fact, a large bias shift on the order of 1 mg was erroneously believed to exist at the time due to an incorrect determination of phase shift due to finite pulse width (see section 4.5). This phase shift has since been corrected,

but the absolute bias error is ignored for the purposes of this result. The linear fit of the data before and after the power cycle yield a measurement of linear gravity gradient of -3120 ± 390 Eotvos and -3290 ± 320 Eotvos respectively. The error bars of the linear fit are within the bounds of the expected value of -3080 Eotvos.

5.1.2 White Mountains Excursion

To further demonstrate gravimeter performance in the field in a more challenging and varying environment, the system was installed into a van and driven along a mountainous road to measure variations in local gravity. Measurements of local gravity were taken at specific locations along the travel route to provide accurate and stable measurements. A track almost entirely along highways was chosen to avoid uneven roads that would pose a threat to the system in transit. The route chosen was primarily along US interstate 93, which runs from the greater Boston area up into the White Mountains of New Hampshire (figure 5.4). The highway has many convenient locations in the Franconia Notch area at which the van can be parked long enough to perform stationary measurements, with elevations as high as 1800 feet. The Kancamagus Highway (NH-12) deviates from I-93 at the town of Lincoln, NH, and has convenient pull-over locations up to elevations of 2800 feet. Assuming a constant linear gravity gradient of 3080 Eotvos, the expected change in local gravity due to elevation change is up to $280 \mu\text{g}$, which is readily measurable by the gravimeter.

The system was installed in the van and remained in place for the duration of the excursion. Measurements were performed from inside the van, utilizing a support structure that was constructed to be capable of damping vibration in transit and leveling out tilts at stationary points. The gravimeter was mounted on a modified version of the 80/20 aluminum cart built for the elevator test. This structure was shortened to carry only the gravimeter system, excluding the electronics rack, with wheels for mobility and extendable leveling feet for tilt adjustments. To protect the system against shock and vibration, the




Figure 5.4: Route taken for the White Mountains excursion. The gradient overlay is the gravity anomaly map reported by GGMplus [106]. Purple points are longer leveled measurements and lighter purple points are quicker unleveled measurements.

base of the support structure was comprised of inflated rubber inner tubes sandwiched between sheets of plywood (figure 5.5). The plywood base was wide enough to fill most of the area in the back of the van, providing a stable platform to which the 80/20 cart and electronics rack were fastened to side by side. The base was partially constructed on the side then inserted into place beneath the system frame in the van. The inner tubes were inflated such that the system was lifted off the floor of the van and there were no other points of contact with the van. The inflated tubes are an excellent damper of vibration, so shocks generated from driving along a road were significantly attenuated. At roughly 2X3X3 feet and 200 pounds, the system tended to sway dramatically while on the inner tube support structure, so loose straps to the floor and ceiling were added to prevent large shifts under significant acceleration.

To keep the system powered, two large battery backups were used. The system, when fully operational, draws 340 W of continuous power. This proved too large a power draw




Figure 5.5: Support structure for the system inside the van. Inflated inner tubes support a wooden base to which the gravimeter structure and electronics rack is mounted.

for the van’s internal battery, so additional power sources were required. A Yeti 3000X 3 kWh power station was tested to be able to supply 8 hours of continuous system operation, A Jackery Explorer 2000 Plus 2 kWh power station was tested to be able to supply 6 hours of continuous operation. To avoid transients in the sensor output due to various startup effects, it was decided to keep the system running uninterrupted throughout the entire excursion. The total excursion lasted about 16 hours, so the Yeti and Jackery alone were almost sufficient to provide continuous power from start to end. To supplement the power needs, the power stations were charged by the van while in transit. The van provided 8 A at 12 V while the engine was powered on, which allowed for roughly an additional 2 hours of run time. The power stations proved just barely sufficient for the entire excursion, lasting just long enough to switch back to wall power upon return to the greater Boston area.

The main set of data were stationary measurements performed at predetermined locations along the route. These measurements were taken at $T = 5 \text{ ms}$ ($a_\pi = 800 \text{ } \mu\text{g}$) and were set up to allow for best possible precision and accuracy. No IMU was present alongside the gravimeter, so calibrated inclinometers were used to define “local up,” and the hard metal legs of the cart were extended down to level out the system. Screw-adjust feet supported each leg to allow for fine adjustment of tilt angle. Engine noise proved

to be a substantial source of vibration with the legs deployed, so the van was turned off while static measurements were performed. The hard legs also prevent sway that may have occurred while floating on the inner tube support structure. Some sway still existed due to varying wind pressure on the van, but this proved to be small. Additionally, this small sway is zero-mean and could be averaged over by performing the measurement long enough. Each static measurement was performed for 30 minutes to allow for sufficient averaging over system and environmental noise.

To supplement the leveled measurements, an additional set of shorter static measurements were performed in which the system remained unleveled and floating on the inner tubes. These measurements then depended on the inclinometers to inform the tilt angle which can be used to calculate the full local gravity. Unleveled measurements were taken with 15 minute dwells to allow for decent precision while also keeping the total excursion time within the bounds of a reasonable day trip. Since the hard legs were not engaged, the inner tubes damped the engine vibration sufficiently to keep it running during unleveled measurements.

The gravimeter was run continuously throughout the trip, and while in transit the gravimeter was operated at a reduced accuracy of $T = 1 \text{ ms}$ ($a_\pi = 20 \text{ mg}$). The inclinometers have a bandwidth on the order of 1 second, so they are able to provide accurate measurements of static leveling but are not reliable real-time measurements of tilt while in motion. Given the high vibration threat of the engine and movement over uneven roads, the data produced while in motion did not provide any noteworthy measurement. The continuous operation did prove to be stable, however, and the gravimeter was able to inform a bias correction that never fringe hopped. This also proves that all components of the system were able to operate without issue while subject to the threat presented while in motion.

Precision of $< 1 \text{ }\mu\text{g}$ was achieved at each leveled measurement location and $< 2 \text{ }\mu\text{g}$ at each unleveled location. The leveled measurement at one point in Lincoln, NH was

performed three times distributed across the excursion to provide an evaluation of standard deviation for leveled measurements. The standard deviation of the three points was $1.9 \mu\text{g}$, which is well below predicted variations in local gravity.

The gravity variation along the trek was initially assumed to be linear as a function of elevation, as was the case in the local building-wide test. This proved quite inaccurate due to gravity anomalies of local terrain. Indeed, the gravity measured at the site in Lincoln, NH was found to be $29 \mu\text{g}$ larger than that measured in Boston, MA, despite a linear prediction of $-80 \mu\text{g}$ due to an increase in altitude of 800 ft. A more sophisticated gravity map was therefore used to determine expected values of gravity at the measurement locations. The global gravity model plus (GGMplus) [106] was developed by Curtin University and TU Munich, and provides gravity maps of the entire Earth's land masses with 200 meter resolution. GGMplus utilizes a few gravity models, notably GRACE [107] and EGM2008 [108], to calculate local gravity anomalies based on a set of sophisticated algorithms. From




Figure 5.6: Result of the measurements taken at each location plotted as a function of elevation. Comparison to the GGMplus prediction at each location demonstrates good agreement of leveled measurements up to a constant offset. Unleveled measurements are less accurate, with only 3 of 6 agreeing with predicted values.

GGMplus, the expected local gravity along the entire trek can be accurately determined. Measurements at stationary sites are then plotted against elevation at each site and

compared to the expected gravity.

Figure 5.6 shows the measurements as a function of altitude and the expected gravity at the corresponding location given by GGMplus. The leveled measurements were found to match quite well to the expected gravity anomaly, up to a constant bias of $7\ \mu\text{g}$. When a constant bias is removed, gravity measured at leveled sites align nearly perfectly with predicted values. The unleveled values met less success, some deviating from prediction by as much as $40\ \mu\text{g}$. It is uncertain why the unleveled measurements were so inaccurate, but one leading hypothesis is that the inclinometers had a substantial scale factor error. While the inclinometers' bias proved to be stable to $10\ \text{mdeg}$ ($\sim 15\ \text{ng}$), separate measurement of the scale factors demonstrated significant nonlinearity and temperature dependence. This means the inclinometers can be trusted for the purpose of leveling the gravimeter to point up, but are less reliable at providing an accurate measure of tilt angle away from zero. The unleveled measurements were therefore largely excluded, given the far better accuracy of the leveled results. Attempts to systematically account for tilt measurement errors had




Figure 5.7: Residuals of select measurements that are in close agreeance with the GGMplus prediction. Site 4, 8, and 11 are repeated measurements at the same site to demonstrate repeatability. Unleveled measurements that were $> 10\ \mu\text{g}$ from the prediction were excluded for visibility.

only a little success, and large errors could not be entirely accounted for when the system

was significantly tilted. Figure 5.7 demonstrates individual deviations of measurements from prediction. All included points lie within $3\text{ }\mu\text{g}$ of prediction with error bars of roughly $1\text{ }\mu\text{g}$. Gravity variations along the trek are on the order of 10s of μg , so the precision of the measurements are sufficient to distinguish gravity variations.

5.2 At-Sea Deployments

The ultimate goal of the gravimeter project was to aid in at-sea navigation. The initially funded project thus began with a tight development schedule that laid out a five year plan to design, build, and deploy the gravimeter in an ultimate shipboard navigation demonstration. The build and development phase began after roughly a year, and unfortunately coincided with the COVID-19 pandemic. Schedule delays in production and integration of the system left little time for system optimization and improvement, and thus many of the performance improvements demonstrated in chapter 4 had not been implemented by the time of the at-sea deployments. Nevertheless, the gravimeter was operational in time for the deployments, and function at-sea and improvement in INS solution were able to be effectively demonstrated. Two shipboard demonstrations were achieved – the first a simple deployment and operation at-sea, the second a more optimized demonstration of navigation aid.

5.2.1 First Deployment – Gulf of Mexico

The first at-sea deployment was a basic field operation test in February 2022. The system was transported to the US naval station in Norfolk, Virginia, where it was installed into a larger container. The container was then shipped to Gulfport, Mississippi, where it was deployed onto a small research ship. The vessel then carried out a five day underway in the Gulf of Mexico, journeying around the continental shelf.

The overall system was assembled with the rack-mounted electronics on top of the

outer support structure for a minimal footprint (see figure 5.8). A Marins M7 IMU was installed below the gravimeter baseplate and served as the basic INS solution to which the gravimeter would provide its correction. A simple insulation layer was constructed out of cardboard that encompassed the gravimeter baseplate and all components mounted therein. The cardboard layer provided some damping of ambient temperature fluctuations while still allowing some ventilation for heat generated by various components, notably the laser system.




Figure 5.8: The system inside the CONEX box. The electronics rack is mounted to a frame that surrounds the gravimeter structure with IMU mounted underneath.

The system was assembled as such in the lab and all basic functions demonstrated before

partial disassembly and shipping. Shipping by van was carried out by the science team personally to provide best-possible care to the system in transit. The larger gravimeter structure with all components in place was moved as a unit, as was the electronics rack, to minimize in-the-field assembly time and reduce risk of reintegration errors. The gravimeter structure and electronics rack were placed in a commercial van and mounted to a support structure similar to but simpler than the one described in section 5.1.2. The main features were the large wooden base floated on shock-absorbing inflated rubber inner tubes and safety straps affixed to the inner walls of the van. The support system proved sufficient to protect the system as it was driven from the development lab at Draper in Cambridge, MA down to Norfolk, VA, where it was installed into a larger CONEX box.

A 2X3 foot footprint was allocated for the gravimeter system inside a custom-built CONEX shipping container. Slightly smaller than a standard intermodal shipping container, the CONEX box's outer footprint is roughly 7X24 feet and 8 feet tall. Internally, it featured a roughly 3X20 foot solid floor area with regularly spaced threaded holes where various research systems could be secured. A similarly-sized front panel area was allocated for human access to systems. A full-faced door was located on one of the 7X8 foot short sides of the box to allow for installation of hardware, and a small human-sized door was located in the middle of the long side to allow access to equipment while underway. The walls were insulated and two air-conditioning units in extendable sections regulated the internal temperature of the box.

After installation into the CONEX box in Norfolk, initial system tests were performed to validate function post-transport. The CONEX box was then shipped via freight truck to Gulfport, MS, where it was craned onto the ship. The ship used for the first at sea excursion was the Point Sur, a small research vessel owned and operated by the University of Southern Mississippi [109]. At 135 feet long and 32 feet wide, the CONEX box was secured carefully at the stern of the ship (figure 5.9).

The initial vacuum pumpdown and MOT capture occurred only six months earlier in



Figure 5.9: The Point Sur docked in Gulfport, MS with the CONEX box installed at the stern of the ship.

August 2021, and atom interferometry was achieved in November 2021. The remaining time was spent integrating hardware and electronics into a portable setup for shipping to the launch location. With little to no time allowed for optimizations, performance of the gravimeter was quite poor. Methods of handling major contributors to stability had not yet been implemented such as optical intensity locks, AC Stark shift compensation, and state reversal. Most importantly, feedforward had not yet been implemented. Without active specific force/vibration phase suppression on the interferometer, the sensitivity had to be significantly lowered while underway.

With the system situated very near to the stern the ship, it was subject to less heave than other locations on the ship, but was also very close to the engine and propeller. Figure 5.16 demonstrates the inertial input along the gravimeter's input access (measured by the cosensor). A significant amount of vibration therefore coupled into the measurement despite the shock absorbers in the legs of the system. The full experiment cycle and therefore also the cosensor was sampled at 34 Hz, so the higher-frequency threat cannot be directly determined, but a significant amount of vibration can still be seen in the Fourier




Figure 5.10: Inertial environment at sea aboard the Point Sur as measured by the z-axis cosensor. (Left) Time domain during a particularly turbulent period sees as large as 0.6 g variation in z-axis specific force. (Right) Fourier transform shows a peak due to ship motion around 0.2 Hz and a shock mount resonance around 16 and 17 Hz that absorbs most higher frequency content.

transform of the time trace. The shock mounts absorbed much vibration, and a resonance near 16 Hz was observed. Wave-induced motion of the ship also produced a large, broad peak around 0.2 Hz, corresponding to a rocking motion of 5 seconds. The amount of motion varied from day to day, with surges as large as 0.6 g seen in the vertical axis on the most turbulent day. To handle such large specific force variations, the interferometer was operated at a very low $T = 0.1$ ms, which has a fringe size of $a_\pi = 2$ g (630 mg/rad). At this reduced sensitivity and a stability floor around 1 mrad, the gravimeter was limited to a precision of 400 mg (see figure 5.11).

The Point Sur excursion was executed solely for the science mission, so the track was free to be chosen for optimal gravity measurements. A number of maneuvers were chosen to demonstrate different gravity features, most notably a grid pattern centered over the continental shelf in the Gulf of Mexico (see figure 5.12). The variation in local mass beneath the ocean surface produced gravity variations as large as 100 μ g. Unfortunately, the precision of the gravimeter in its current state was an order of magnitude too large to discern these features.

The mission then ultimately served as a demonstration of deployment and ability to




Figure 5.11: Allan deviations of the gravimeter during the first at sea underway. With $T = 0.1$ ms, the sensitivity is severely limited. The stationary trace (blue) demonstrates performance of the gravimeter, while the at sea trace (red) represents real precision while in motion, which is larger than the stationary precision due to vibration and ship motion.

function aboard a surface ship. The gravimeter was designed with larger ships experiencing less severe dynamics in mind, so the Point Sur deployment served as a demonstration of function in a worst case environment. The fairly large waves during the middle of the journey produced large sway and heave, and the placement of the system very near to the engine and motor produced substantial vibration. Additionally, the CONEX box, while somewhat temperature-regulated, was subject to direct sunlight and saw around 5°C internal temperature swings over day/night cycles. The mechanical and thermal challenge to the system was therefore a thorough demonstration of its robustness in the field. The only source of failure was the LCM, which lost frequency lock three times over the course of the journey. The lock signal exhibited a temperature dependence that needed to be manually adjusted periodically. The lock loss events coincided with nightly ambient temperature drops combined with opening of the door for access to other experiments located in the CONEX box. The lock was quickly restored each time, and operation




Figure 5.12: At-sea track taken by the Point Sur as recorded by GPS. The EGM2008 geoid height prediction loosely correlates with gravity anomaly. Two grid search patterns were implemented near the continental shelf for maximum gravity variation.

of the gravimeter otherwise continued without issue. This temperature dependence has since been addressed, and the photonics box where the LCM is located is also kept more thermally insulated to prevent larger temperature fluctuations.

5.2.2 Second Deployment – Hawaii

The system returned to Draper immediately after the February deployment, where a number of critical upgrades were performed. It was returned in the same manner as it was shipped out, traveling via freight truck in the CONEX from Gulfport, MS back to the base in Norfolk, VA, then transported in the van from Norfolk back to Draper.

Over the next five months, significant improvements were made to the stability and performance of the gravimeter. Intensity locks were implemented to stabilize optical power fluctuations, the LCM signal was cleaned up to provide a more reliable frequency lock, and k-reversal was implemented (among other improvements). Most importantly, feedforward and feedback were implemented. Feedforward vibration suppression allowed for a much

more sensitive gravimeter scale factor. Integrated cosensor measurements disciplined by the gravimeter were sent to an INS algorithm at roughly 30 Hz, where the gravimeter measurement was coupled with the M7 IMU for a real-time navigation solution. Figure 5.13 demonstrates the significant improvement to sensitivity of the gravimeter achieved by the time of this second deployment.




Figure 5.13: Allan deviations at the time of the Gulf of Mexico excursion (blue) compared to that at the time of the Hawaii excursion (red). The five month period between the missions saw a significant increase in performance due to critical developments such as feedforward/feedback, k-reversal, and intensity locks.

The second excursion began in port at Pearl Harbor, Hawaii in August 2022. The system therefore had to be shipped by air, which occurred in two large sections – the integrated gravimeter structure and the electronics rack. The gravimeter structure was packed into a large plastic shipping crate. It was again floated on a small plywood base supported by inflated rubber tubes and affixed to the sides with ratchet straps. The ion pump was kept powered by a series of 1.5 V AA batteries, wired to provide 12 V continuous power throughout shipping. The electronics rack was boxed in a wooden crate and padded with multiple layers of bubble wrap. It was bolted at the base to prevent

movement within the box.

Both crates shipped via “overnight” air, which took three days from leaving the building in Cambridge, MA to receiving in Honolulu, HI. The batteries proved sufficient to keep the ion pump operational until wall power could be restored upon arrival in Hawaii. Passive single-use shock and temperature sensors were placed on the gravimeter structure and electronics rack to bound the environment they were subject to during shipping. Some of the smaller ~ 1 g shock sensors were tripped as well as the 10°C and 30°C temperature sensors. The more extreme sensors were not tripped, indicating a careful handling across shipment. No damage to the system was found, and operation was restored without issue upon installation into the CONEX box.



Figure 5.14: The HMNZS Aotearoa at sea. Photograph is taken from the bow of the ship looking back at the deck and control tower. The CONEX box is located on the deck at starboard, roughly mid-ship.

The setup was the same as the previous at-sea deployment, with the gravimeter installed in the CONEX box and insulated with cardboard. The CONEX box was craned onto

the ship and placed on the foredeck, roughly one-third of the length from the bow (figure 5.14). The ship chosen for the second excursion was the HMNZS Aotearoa [110, 111], a New Zealand Naval tanker ship. The ship was significantly larger than the Point Sur, at 568 feet long by 80 feet wide. The inertial threat was therefore significantly less – the motion of the ship due to waves was far smaller and the engine was very far away from the deck of the ship.

The at-sea journey lasted 24 days, with one half-day break at day eleven for a personnel change. Due to personnel limits on the ship, the system initially ran unattended for the first eleven days at sea. The ship then returned to dock, where the science crew was brought on board for a fourteen-day continuous journey. The ship was participating in a larger mission, RIMPAC 2022 [112], and the ship maneuvers were entirely dictated by the Aotearoa’s refuel-at-sea mission. The ship performed various exercises and routines, so the




Figure 5.15: At-sea track taken by the HMNZS Aotearoa as recorded by GPS. The EGM2008 geoid height prediction loosely correlates with gravity anomaly. Proximity to the Hawaiian islands provided measurable gravity anomaly variations throughout the excursion.

path could not be set by the science crew for optimal travel near significant local gravity features. The trek (figure 5.15) was therefore a realistic demonstration of gravimeter

operation and gravity-aided navigation on board a ship operating with other duties. The course did travel nearby a couple of the islands of Hawaii, which exhibit significant gravity variations due to the deep ocean base from which the volcanic islands arise.




Figure 5.16: Inertial environment at sea aboard the HMNZS Aotearoa as measured by the z-axis of the M7. (Left) Linear x-axis shows higher frequency content up to 100 Hz. Although much vibration is suppressed, some bleedthrough remains. (Right) Logarithmic x-axis shows lower frequency content, especially the ship motion around 0.1 Hz.

Given the significantly better environment and improvements to the gravimeter, the measure of gravity while underway was much more accurate. Figure 5.16 shows a maximum variation in specific force due to ship motion and vibration of about 1 mg (as opposed to 10 mg seen on the Point Sur in figure 5.10). Feedforward suppressed vibration at the scale of 100X, keeping residual input on the gravimeter phase to around 10 μ g. The cosensor-gravimeter input axis misalignment correction discussed in section 4.2 had not yet been implemented, so a tilt-dependent bias shift of roughly 2.7 mg per degree tilt was present (from equation 4.34, $\delta b/g/\theta \propto \xi \approx 2.7 \text{ mrad}$). This bias shift effectively limited the sensitivity at which the gravimeter could operate, since a sensitivity of $a_\pi < 3 \text{ mg}$ could result in fringe hopping. The sensitivity was thus chosen to be $T = 2.7 \text{ ms} \rightarrow a_\pi = 2.7 \text{ mg}$. Although tilts of a degree produced large phase shifts, the large ship was relatively stable, and such phase shifts were zero-mean given sufficient averaging time (greater than a few rocking periods of ~ 10 seconds).

The gravimeter ran with few issues through the duration of the excursion. One problem

arose roughly eight days into the first portion of the excursion, during which the system was unmanned. An electrical/grounding fault produced a significant DC jump in the quad fluorescence detector, which corrupted the phase calculation in the feedback algorithm. Although the gravimeter was functioning without issue, the erroneous signal produced a bad bias correction in feedback that corrupted the total specific force measurement. The data during the remaining three unmanned days therefore did not produce a reliable measurement of gravity, but did demonstrate otherwise continued reliable operation of the gravimeter (optical stability, thermal robustness, etc.). Immediately upon arrival of the science crew eleven days in, the electrical jump was addressed by power cycling the detector. The issue did not arise again, and feedback bias correction was reinitialized. The gravimeter then continued to operate without issue and output a bias-corrected specific force measurement for the thirteen-day duration of the excursion.

CHAPTER 6

Conclusion

The work presented in this thesis represents a significant improvement to the operation of atom interferometric sensors in a real world platform. The construction of a more mature prototype capable of reliable long term measurements of gravity and its deployability outside of a laboratory environment is successfully demonstrated through multiple field excursions. The system has been packaged into a mobile configuration capable of being transported mostly-assembled and deployed with relative ease in the field. The practical challenges of performing a sensitive atom interferometer measurement while subject to a non-ideal environment were overcome and real applications of the gravimeter to aid in inertial navigation were realized.

Two deployments of the gravimeter at sea demonstrated the ability of the system to operate in a dynamic environment, where it was subjected to significant vibration and inertial variations. The first at-sea deployment occurred very early in the life of the sensor, and served primarily as a demonstration of robustness of the system. While performance of the gravimeter was rather poor at the time, the hardware proved to be impervious to large vibrations and surge/sway of the ship in fairly rough sea conditions. In particular, the optics in the AOD box and sense head were shown to have solid alignment that survived without issue through transport from Massachusetts to Virginia to Mississippi and back.

Better performance of the gravimeter was demonstrated at the second underway, at which point significant upgrades such as feedforward and k-reversal had been implemented. The resulting real-time measurement proved stable in real-world conditions over three weeks, and the output was successfully integrated into a demonstration of inertial navigation aid.

Further development of the gravimeter over the next few years aimed to improve sensor performance and reliability while remaining field-deployable. Improvements to the feedforward and feedback algorithms enabled more accurate integration of the hybrid cosensor-gravimeter operation both in the lab and in the field. A stability floor of 150 ng can reliably be achieved in about 1000 seconds and remains stable over weeks or longer. The absolute accuracy of the sensor has been bounded to less than 10 μ g and is a focus of current investigations of the interferometer implementation. A number of physical effects that generate phase shifts have already been studied and addressed, and others continue to be identified.

The gravimeter has proven to be reliable over five years with most of the system showing no signs of degradation throughout the lifetime of near-continuous operation. Minor issues in the photonics system were addressed, specifically the replacement of the PPLN frequency-doubler that proved to have a roughly six month lifetime with a more reliable part that has remained fully operational over nearly two years with no signs of decline of output power. The one major failure observed was the degradation of MOT optics inside the vacuum. The mechanism was studied and theorized to be due to chemical interactions of Rb with the HR coatings of the optics. Protective alumina coatings were applied to protect the optics and no sign of attenuation has been observed since the repair, though further observation over time is needed to determine if the solution is viable long term. Nearly all other mechanical, electrical, and optical components of the system have proven reliable long term and robust to field deployments.

6.1 Future Work

Development of the gravimeter is an ongoing effort, with multiple upgrades planned for the system. The efforts can be loosely grouped into mechanical and electrical engineering efforts and study of underlying physical effects in the atom interaction. The goal is to reduce the SWAP of the system while continuing to improve sensor performance characteristics such as accuracy, stability, and long-term reliability. Such efforts have short term targets at improving performance and longer term targets of dramatically reducing SWAP and demonstrating many-year longevity of the system.

Physics-based investigations into sources of error in the interferometer measurement are a current and future area of study. The absolute accuracy is currently offset by about $-7\text{ }\mu\text{g}$, the source of which is not presently understood. Given that one of the primary draws of the atom-based gravimeter is its ability to produce absolutely accurate measurements, the understanding and removal of effects that produce such offsets without need for external calibration is of high interest. Recent tests have identified contributions from parasitic DI interferometers due to cross coupling of the Raman beams. The characterization of these transitions as a manifestation of phase shifts in the main interferometer may be a pathway to significantly improved accuracy and stability in the gravimeter. Other sources of phase shifts being considered are nonuniformities in the Raman beam wavefronts and algorithmic errors in Kalman filter phase tracking that may allow common-mode effects to bleed through to the k-subtracted phase.

While such investigations into accuracy and continuing engineering work to improve the system are important to the development of the sensor, these changes make it difficult to evaluate long-term performance of the gravimeter. The construction of a second near-identical prototype would allow parallel investigations to be performed, such as multi-month stability and reliability demonstrations. These kinds of tests are important in demonstrating the gravimeter's promise of long-term accurate performance that make it an attractive candidate as a navigation aid. The ability to compare measurements of

the two systems would also demonstrate accuracy and repeatability of the gravimeter.

Electrical engineering efforts aim to reduce the SWAP of the system by integrating splayed-out electronics into integrated PCBs. Much of the electronics are groups of COTS evaluation boards assembled together for signal processing. In particular, the DACs and power delivery for sensors located on or near the sense head are bundled into messy groups that take up a lot of space on the gravimeter base plate. Moreover, the signal routing and grounding scheme of the sensors are likely introducing a lot of electrical noise onto the signals. These proximity electronics take up a roughly two cubic foot space due to the splayed-out nature of the design. A custom DAC board is currently being designed that will replace the entire two-cubic foot set of proximity electronics with a single board of roughly 6X6 inches. While significantly reducing overall SWAP, the board will also likely handle signal routing and grounding much more efficiently, leading to cleaner signals. In particular, the cosensors' analog signals are of high interest, and producing a more clean raw cosensor output will not only improve the short term noise of the hybrid feedback output, but may also reduce noise imprinted on the atoms by the cosensor through feedforward.

Mechanical engineering efforts aim to reduce the overall SWAP by decreasing the volume enclosed by the system in line with the shrunk integrated electronics and also decreasing the weight by using lighter support hardware. The original support structure was designed to also carry a large IMU that is no longer present, so the shock mounts are too stiff for the current system. A better selection of shock mounts for the current system weight will reduce mount resonance and therefore improve the suppression of mid-high frequency vibrations. The new system will have a weight reduction from roughly 140 kg to 70 kg (not including the electronics rack), and initial tests and calculations indicate a mount resonance decreased from 20 Hz to 14 Hz.

To further study the effects of motion and vibration on the gravimeter, a hexapod motion stage is being acquired capable of simulating fairly aggressive ship motion. The

motion stage will be able to produce translations and rotations that exceed expected sea states as well as vibrations from sub-Hz to hundreds of Hz, improving the ability to study effects of motion and vibration on the gravimeter. Such studies will lead to improvements in feedforward vibration suppression, which is a limiting factor of the system's ability to operate accurately in dynamic environments. A promising possible improvement involves factoring the bandwidth of the cosensor into the estimation of specific force provided to the gravimeter for feedforward. Additionally, an attitude, heading, and reference system (AHRS) is being installed on the system to provide the gravimeter with a measurement of dynamic motion. This enables transformation of the moving gravimeter measurement of specific force into a vertical measurement of gravity. The AHRS also serves as an IMU with which algorithms can be developed and improved upon to use the gravity measurement as an aid in the inertial measurement.

Longer term goals seek to greatly reduce the size of the system, which remains a fairly large cube with an accompanying electronics rack. The easiest candidate for size reduction is the remaining electronics and RF systems. These currently take up the bulk of the system's space and could likely be reduced to a very compact form factor with a moderate engineering effort. The thermal box that monitors and controls heaters could be reduced to a very small PCB. The RF box is more difficult since many of the microwave components such as amplifiers produce a significant amount of heat. While a slightly more difficult engineering task, the RF system could be reduced to something small enough to fit near the current sense head. The optical system also takes up a lot of space, with the photonics box in particular enclosing a lot of empty space. Tighter packing of fiber-coupled components could likely reduce the volume by a factor of 10X. The AOD box could also be reduced in volume by a factor of up to 8X. Further reductions in size of the optical system would likely involve photonics chip-scale integration, an effort which is increasingly being developed at Draper (and other places). Finally, the sense head would pose a greater challenge to significantly decrease in size. The limiting

factor is the choice to put optics inside the vacuum (to avoid wavefront errors). Size reduction of the vacuum chamber would require careful consideration of the optics to avoid impacts on performance of the gravimeter. All together, it is feasible that the entire system could fit readily within the 2X3 foot base plate with moderate engineering effort. Further reductions to size, weight, and power are possible and will be undertaken in future efforts. The system could ultimately fit within the volume of the sense head (10 inches diameter by 15 inches tall) with a more substantial redesign. The atom gravimeter could then be considered for larger-scale deployment as a mature sensor on ships as a use in navigation aid or as a field device for metrological studies.

References

- [1] Morris M Kuritsky, Murray S Goldstein, IA Greenwood, H Lerman, JE McCarthy, T Shanahan, M Silver, and JH Simpson. Inertial navigation. *Proceedings of the IEEE*, 71(10):1156–1176, 1983.
- [2] Neil M Barbour. Inertial navigation sensors. *NATO RTO Lecture Series, RTO-EN-SET-116, Low-Cost Navigation Sensors and Integration Technology*, pages 1–24, 2010.
- [3] Oliver J Woodman. An introduction to inertial navigation. Technical report, University of Cambridge, Computer Laboratory, 2007.
- [4] Guoquan Huang. Visual-inertial navigation: A concise review. In *2019 international conference on robotics and automation (ICRA)*, pages 9572–9582. IEEE, 2019.
- [5] Yuri Dachev and A Panov. 21st century celestial navigation systems. *18th Annual General Assembly*, 2017.
- [6] Qiong Wang, Chen Zheng, Peili Wu, and Xiaoyu Wang. Geomagnetic/inertial navigation integrated matching navigation method. *Heliyon*, 8(11), 2022.
- [7] Felix Goldenberg. Geomagnetic navigation beyond the magnetic compass. In *Proceedings of IEEE/ION PLANS 2006*, pages 684–694, 2006.
- [8] Justin A Richeson. Gravity gradiometer aided inertial navigation within non-gnss environments. In *Proceedings of the 20th International Technical Meeting of the Satellite Division of the Institute of Navigation (ION GNSS 2007)*, pages 1089–1100, 2007.
- [9] Hubiao Wang, Lin Wu, Hua Chai, Yaofei Xiao, Houtse Hsu, and Yong Wang. Characteristics of marine gravity anomaly reference maps and accuracy analysis of gravity matching-aided navigation. *Sensors*, 17(8):1851, 2017.

- [10] Iginio Marson. A short walk along the gravimeters path. *International Journal of Geophysics*, 2012(1):687813, 2012.
- [11] TM Niebauer, GS Sasagawa, James E Faller, R Hilt, and Fred Klopping. A new generation of absolute gravimeters. *Metrologia*, 32(3):159, 1995.
- [12] P Křen, V Pálinkáš, M Val’ko, and P Mašika. Improved measurement model for FG5/X gravimeters. *Measurement*, 171:108739, 2021.
- [13] L Vitushkin. Measurement standards in gravimetry. *Gyroscopy and navigation*, 2(3):184–191, 2011.
- [14] Sébastien Merlet, Quentin Bodart, Nicola Malossi, Arnaud Landragin, F Pereira Dos Santos, Olga Gitlein, and Ludger Timmen. Comparison between two mobile absolute gravimeters: optical versus atomic interferometers. *Metrologia*, 47(4):L9, 2010.
- [15] Pierre Gillot, Olivier Francis, Arnaud Landragin, F Pereira Dos Santos, and Sébastien Merlet. Stability comparison of two absolute gravimeters: optical versus atomic interferometers. *Metrologia*, 51(5):L15, 2014.
- [16] Shuqing Wu, Jinyang Feng, Chunjian Li, Duowu Su, Qiyu Wang, Ruo Hu, and Lishuang Mou. The results of 10th international comparison of absolute gravimeters (ICAG-2017). *Journal of Geodesy*, 95(6):63, 2021.
- [17] Chandrasekhara Venkata Raman. A new radiation. *Indian Journal of physics*, 2:387–398, 1928.
- [18] Isidor Isaac Rabi, Jerrold R Zacharias, Sidney Millman, and Polykarp Kusch. A new method of measuring nuclear magnetic moment. *Physical review*, 53(4):318, 1938.
- [19] Norman F Ramsey. A new molecular beam resonance method. *Physical Review*, 76(7):996, 1949.

- [20] Richard C Prince, Renee R Frontiera, and Eric O Potma. Stimulated Raman scattering: from bulk to nano. *Chemical reviews*, 117(7):5070–5094, 2017.
- [21] Olivier Carnal and Jürgen Mlynek. Young’s double-slit experiment with atoms: A simple atom interferometer. *Physical review letters*, 66(21):2689, 1991.
- [22] David W Keith, Christopher R Ekstrom, Quentin A Turchette, and David E Pritchard. An interferometer for atoms. *Physical review letters*, 66(21):2693, 1991.
- [23] F Riehle, Th Kisters, A Witte, J Helmcke, and Ch J Bordé. Optical Ramsey spectroscopy in a rotating frame: Sagnac effect in a matter-wave interferometer. *Physical review letters*, 67(2):177, 1991.
- [24] Mark Kasevich and Steven Chu. Atomic interferometry using stimulated Raman transitions. *Physical review letters*, 67(2):181, 1991.
- [25] Ch J Bordé. Atomic interferometry with internal state labelling. *Physics letters A*, 140(1-2):10–12, 1989.
- [26] Mark Kasevich and Steven Chu. Measurement of the gravitational acceleration of an atom with a light-pulse atom interferometer. *Applied Physics B*, 54:321–332, 1992.
- [27] Sébastien Bize, Philippe Laurent, Michel Abgrall, Harold Marion, Ivan Maksimovic, Luigi Cacciapuoti, Jan Grünert, Céline Vian, F Pereira Dos Santos, Peter Rosenbusch, et al. Cold atom clocks and applications. *Journal of Physics B: Atomic, molecular and optical physics*, 38(9):S449, 2005.
- [28] Vladislav Gerginov, Gregory Hoth, Thomas Heavner, Thomas Parker, Kurt Gibble, and Jeffrey Sherman. Accuracy evaluation of primary frequency standard NIST-F4. *Metrologia*, 2025.
- [29] Nathan Hinkley, Jeff A Sherman, Nathaniel B Phillips, Macro Schioppo, Nathan D

- Lemke, Kyle Beloy, Marco Pizzocaro, Christopher W Oates, and Andrew D Ludlow. An atomic clock with 10^{-18} instability. *Science*, 341(6151):1215–1218, 2013.
- [30] Jeffrey B Fixler, GT Foster, JM McGuirk, and MA Kasevich. Atom interferometer measurement of the Newtonian constant of gravity. *Science*, 315(5808):74–77, 2007.
- [31] G Rosi, L Cacciapuoti, F Sorrentino, M Menchetti, Marco Prevedelli, and GM Tino. Measurement of the gravity-field curvature by atom interferometry. *Physical Review Letters*, 114(1):013001, 2015.
- [32] Chao Xue, Jian-Ping Liu, Qing Li, Jun-Fei Wu, Shan-Qing Yang, Qi Liu, Cheng-Gang Shao, Liang-Cheng Tu, Zhong-Kun Hu, and Jun Luo. Precision measurement of the Newtonian gravitational constant. *National Science Review*, 7(12):1803–1817, 2020.
- [33] Dennis Schlippert, Jonas Hartwig, Henning Albers, Logan L Richardson, Christian Schubert, Albert Roura, Wolfgang P Schleich, Wolfgang Ertmer, and Ernst M Rasel. Quantum test of the universality of free fall. *Physical Review Letters*, 112(20):203002, 2014.
- [34] Peter Asenbaum, Chris Overstreet, Minjeong Kim, Joseph Curti, and Mark A Kasevich. Atom-interferometric test of the equivalence principle at the 10^{-12} level. *Physical Review Letters*, 125(19):191101, 2020.
- [35] Marco Giacinto Tarallo, Tommaso Mazzoni, Nicola Poli, DV Sutyurin, X Zhang, and GM Tino. Test of Einstein equivalence principle for 0-spin and half-integer-spin atoms: search for spin-gravity coupling effects. *Physical review letters*, 113(2):023005, 2014.
- [36] Savas Dimopoulos, Peter W Graham, Jason M Hogan, Mark A Kasevich, and Surjeet Rajendran. Atomic gravitational wave interferometric sensor. *Physical Review D—Particles, Fields, Gravitation, and Cosmology*, 78(12):122002, 2008.

- [37] Oliver Buchmueller, John Ellis, and Ulrich Schneider. Large-scale atom interferometry for fundamental physics. *Contemporary Physics*, 64(2):93–110, 2023.
- [38] Benjamin P Abbott, Richard Abbott, Thomas D Abbott, Matthew R Abernathy, Fausto Acernese, Kendall Ackley, Carl Adams, Thomas Adams, Paolo Addesso, Rana X Adhikari, et al. Observation of gravitational waves from a binary black hole merger. *Physical review letters*, 116(6):061102, 2016.
- [39] Pau Amaro-Seoane, Heather Audley, Stanislav Babak, John Baker, Enrico Barausse, Peter Bender, Emanuele Berti, Pierre Binetruy, Michael Born, Daniele Bortoluzzi, et al. Laser interferometer space antenna. *arXiv preprint arXiv:1702.00786*, 2017.
- [40] Benjamin Elliott Garber. *Strontium Clock Atom Interferometry for Fundamental Physics*. PhD thesis, Stanford University, 2024.
- [41] Peter W Graham, Jason M Hogan, Mark A Kasevich, Surjeet Rajendran, and Roger W Romani. Mid-band gravitational wave detection with precision atomic sensors. *arXiv preprint arXiv:1711.02225*, 2017.
- [42] Mahiro Abe, Philip Adamson, Marcel Borcean, Daniela Bortoletto, Kieran Bridges, Samuel P Carman, Swapan Chattopadhyay, Jonathon Coleman, Noah M Curfman, Kenneth DeRose, et al. Matter-wave atomic gradiometer interferometric sensor (MAGIS-100). *Quantum Science and Technology*, 6(4):044003, 2021.
- [43] Jason M Hogan, David MS Johnson, Susannah Dickerson, Tim Kovachy, Alex Sugarbaker, Sheng-wei Chiow, Peter W Graham, Mark A Kasevich, Babak Saif, Surjeet Rajendran, et al. An atomic gravitational wave interferometric sensor in low earth orbit (AGIS-LEO). *General Relativity and Gravitation*, 43:1953–2009, 2011.
- [44] Brynle Barrett, I Chan, and A Kumarakrishnan. Atom-interferometric techniques for measuring uniform magnetic field gradients and gravitational acceleration. *Physical Review A—Atomic, Molecular, and Optical Physics*, 84(6):063623, 2011.

- [45] Yuval Cohen, Krishna Jadeja, Sindi Sula, Michela Venturelli, Cameron Deans, Luca Marmugi, and Ferruccio Renzoni. A cold atom radio-frequency magnetometer. *Applied Physics Letters*, 114(7), 2019.
- [46] TL Gustavson, Arnaud Landragin, and MA Kasevich. Rotation sensing with a dual atom-interferometer Sagnac gyroscope. *Classical and Quantum Gravity*, 17(12):2385, 2000.
- [47] Alexandre Gauguier, Benjamin Canuel, Thomas Lévêque, Walid Chaibi, and Arnaud Landragin. Characterization and limits of a cold-atom Sagnac interferometer. *Physical Review A—Atomic, Molecular, and Optical Physics*, 80(6):063604, 2009.
- [48] MM Beydler, ER Moan, Z Luo, Z Chu, and CA Sackett. Guided-wave Sagnac atom interferometer with large area and multiple orbits. *AVS Quantum Science*, 6(1), 2024.
- [49] MJ Snadden, JM McGuirk, P Bouyer, KG Haritos, and MA Kasevich. Measurement of the Earth’s gravity gradient with an atom interferometer-based gravity gradiometer. *Physical review letters*, 81(5):971, 1998.
- [50] Nan Yu, JM Kohel, JR Kellogg, and L Maleki. Development of an atom-interferometer gravity gradiometer for gravity measurement from space. *Applied Physics B*, 84:647–652, 2006.
- [51] Wei Lyu, Jia-Qi Zhong, Xiao-Wei Zhang, Wu Liu, Lei Zhu, Wei-Hao Xu, Xi Chen, Biao Tang, Jin Wang, and Ming-Sheng Zhan. Compact high-resolution absolute-gravity gradiometer based on atom interferometers. *Physical Review Applied*, 18(5):054091, 2022.
- [52] Achim Peters, Keng Yeow Chung, and Steven Chu. High-precision gravity measurements using atom interferometry. *Metrologia*, 38(1):25, 2001.
- [53] Guillaume Stern, Baptiste Battelier, Rémi Geiger, Gaël Varoquaux, André Villing,

- Frédéric Moron, Olivier Carraz, Nassim Zahzam, Yannick Bidel, W Chaibi, et al. Light-pulse atom interferometry in microgravity. *The European Physical Journal D*, 53:353–357, 2009.
- [54] Yannick Bidel, Olivier Carraz, Renée Charrière, Malo Cadoret, Nassim Zahzam, and Alexandre Bresson. Compact cold atom gravimeter for field applications. *Applied Physics Letters*, 102(14), 2013.
- [55] Yin Zhou, Can Zhang, Peijun Chen, Bing Cheng, Dong Zhu, Kainan Wang, Xiaolong Wang, Bin Wu, Zhongkun Qiao, Qiang Lin, et al. A testing method for shipborne atomic gravimeter based on the modulated coriolis effect. *Sensors*, 23(2):881, 2023.
- [56] Yannick Bidel, N Zahzam, C Blanchard, A Bonnin, Malo Cadoret, Alexandre Bresson, D Rouxel, and MF Lequentrec-Lalancette. Absolute marine gravimetry with matter-wave interferometry. *Nature communications*, 9(1):627, 2018.
- [57] Bin Wu, Can Zhang, Kainan Wang, Bing Cheng, Dong Zhu, Rui Li, Xiaolong Wang, Qiang Lin, Zhongkun Qiao, and Yin Zhou. Marine absolute gravity field surveys based on cold atomic gravimeter. *IEEE Sensors Journal*, 23(20):24292–24299, 2023.
- [58] Yannick Bidel, Nassim Zahzam, Alexandre Bresson, Cédric Blanchard, Malo Cadoret, Arne V Olesen, and René Forsberg. Absolute airborne gravimetry with a cold atom sensor. *Journal of Geodesy*, 94:1–9, 2020.
- [59] Alexis Bonnin, Yannick Bidel, Jeanne Bernard, Cédric Blanchard, Malo Cadoret, Nassim Zahzam, Sylvain Schwartz, and Alexandre Bresson. Marine and airborne gravimetry with an absolute cold atom sensor. In *2022 IEEE International Symposium on Inertial Sensors and Systems (INERTIAL)*, pages 1–4. IEEE, 2022.
- [60] Yannick Bidel, Nassim Zahzam, Alexandre Bresson, Cédric Blanchard, Alexis Bonnin, Jeanne Bernard, Malo Cadoret, Tim E Jensen, René Forsberg, Corinne Salaun, et al. Airborne absolute gravimetry with a quantum sensor, comparison with classical

- technologies. *Journal of Geophysical Research: Solid Earth*, 128(4):e2022JB025921, 2023.
- [61] Richard Stoner, David Butts, Joseph Kinast, and Brian Timmons. Analytical framework for dynamic light pulse atom interferometry at short interrogation times. *Journal of the Optical Society of America B*, 28(10):2418–2429, 2011.
- [62] Krish Kotru. *Timekeeping and accelerometry with robust light pulse atom interferometers*. PhD thesis, Massachusetts Institute of Technology, 2015.
- [63] Kathryn Moler, David S Weiss, Mark Kasevich, and Steven Chu. Theoretical analysis of velocity-selective Raman transitions. *Physical Review A*, 45(1):342, 1992.
- [64] Paul R Berman. *Atom interferometry*. Academic press, 1997.
- [65] Kai Bongs, R Launay, and Mark A Kasevich. High-order inertial phase shifts for time-domain atom interferometers. *Applied Physics B*, 84:599–602, 2006.
- [66] Guglielmo M Tino and Mark A Kasevich. *Atom interferometry*, volume 188. IOS Press, 2014.
- [67] Charles Antoine. Matter wave beam splitters in gravito-inertial and trapping potentials: generalized ttt scheme for atom interferometry. *Applied Physics B*, 84:585–597, 2006.
- [68] Alain Miffre, Marion Jacquey, Matthias Büchner, Gérard Tréneç, and Jacques Vigué. Atom interferometry. *Physica Scripta*, 74(2):C15, 2006.
- [69] Sven Abend, Matthias Gersemann, Christian Schubert, Dennis Schlipfert, Ernst M Rasel, Matthias Zimmermann, Maxim A Efremov, Albert Roura, Frank A Narducci, and Wolfgang P Schleich. Atom interferometry and its applications. In *Foundations of quantum theory*, pages 345–392. IOS Press, 2019.
- [70] Lambert Paul Parazzoli, Aaron M Hankin, and Grant W Biedermann. Observa-

- tion of free-space single-atom matter wave interference. *Physical review letters*, 109(23):230401, 2012.
- [71] Patrick Cheinet, Benjamin Canuel, Franck Pereira Dos Santos, Alexandre Gauguier, Florence Yver-Leduc, and Arnaud Landragin. Measurement of the sensitivity function in a time-domain atomic interferometer. *IEEE Transactions on instrumentation and measurement*, 57(6):1141–1148, 2008.
- [72] David L Butts, Joseph M Kinast, Krish Kotru, Antonije M Radojevic, Brian P Timmons, and Richard E Stoner. Coherent population trapping in Raman-pulse atom interferometry. *Physical Review A—Atomic, Molecular, and Optical Physics*, 84(4):043613, 2011.
- [73] Erwin L Hahn. Spin echoes. *Physical review*, 80(4):580, 1950.
- [74] Eric L Raab, Mara Prentiss, Alex Cable, Steven Chu, and David E Pritchard. Trapping of neutral sodium atoms with radiation pressure. *Physical review letters*, 59(23):2631, 1987.
- [75] Harold J Metcalf and Peter Van der Straten. *Laser cooling and trapping*. Springer Science & Business Media, 1999.
- [76] Christopher D Wallace, Timothy P Dinneen, Kit-Yan N Tan, A Kumarakrishnan, Phillip L Gould, and Juha Javanainen. Measurements of temperature and spring constant in a magneto-optical trap. *Journal of the Optical Society of America B*, 11(5):703–711, 1994.
- [77] AM Steane, M Chowdhury, and CJ Foot. Radiation force in the magneto-optical trap. *Journal of the Optical Society of America B*, 9(12):2142–2158, 1992.
- [78] M Drewsen, Ph Laurent, A Nadir, G Santarelli, A Clairon, Y Castin, D Grison, and C Salomon. Investigation of sub-Doppler cooling effects in a cesium magneto-optical trap. *Applied Physics B*, 59:283–298, 1994.

- [79] Paul D Lett, Richard N Watts, Christoph I Westbrook, William D Phillips, Phillip L Gould, and Harold J Metcalf. Observation of atoms laser cooled below the Doppler limit. *Physical review letters*, 61(2):169, 1988.
- [80] Harold J Metcalf and Peter Van der Straten. *Laser cooling and trapping*, chapter 8. Springer Science & Business Media, 1999.
- [81] K Dieckmann, RJC Spreeuw, M Weidemüller, and JTM Walraven. Two-dimensional magneto-optical trap as a source of slow atoms. *Physical Review A*, 58(5):3891, 1998.
- [82] KI Lee, JA Kim, HR Noh, and W Jhe. Single-beam atom trap in a pyramidal and conical hollow mirror. *Optics letters*, 21(15):1177–1179, 1996.
- [83] Matthieu Vangeleyn, Paul F Griffin, Erling Riis, and Aidan S Arnold. Laser cooling with a single laser beam and a planar diffractor. *Optics letters*, 35(20):3453–3455, 2010.
- [84] Jérôme Estève. Trapped by nanostructures. *Nature nanotechnology*, 8(5):317–318, 2013.
- [85] CC Nshii, Matthieu Vangeleyn, Joseph P Cotter, Paul F Griffin, EA Hinds, Charles N Ironside, Patrick See, AG Sinclair, Erling Riis, and Aidan S Arnold. A surface-patterned chip as a strong source of ultracold atoms for quantum technologies. *Nature nanotechnology*, 8(5):321–324, 2013.
- [86] James P McGilligan, Kaitlin R Moore, Argyrios Dellis, GD Martinez, Emeric de Clercq, PF Griffin, AS Arnold, E Riis, Rodolphe Boudot, and John Kitching. Laser cooling in a chip-scale platform. *Applied Physics Letters*, 117(5), 2020.
- [87] Eric Imhof, Benjamin K Stuhl, Brian Kasch, Bethany Kroese, Spencer E Olson, and Matthew B Squires. Two-dimensional grating magneto-optical trap. *Physical Review A*, 96(3):033636, 2017.

- [88] Daniel A Steck. *Quantum and Atom Optics*, chapter 5. Apr 2025. available online at <http://steck.us/teaching> (revision 0.16.3).
- [89] D.W. Allan. Statistics of atomic frequency standards. *Proceedings of the IEEE*, 54(2):221–230, 1966.
- [90] Naser El-Sheimy, Haiying Hou, and Xiaoji Niu. Analysis and modeling of inertial sensors using Allan variance. *IEEE Transactions on instrumentation and measurement*, 57(1):140–149, 2007.
- [91] Mavis Driscoll. Private communication, Aug 2023.
- [92] J Le Gouët, TE Mehlstäubler, Jaewan Kim, Sébastien Merlet, Andre Clairon, Arnaud Landragin, and F Pereira Dos Santos. Limits to the sensitivity of a low noise compact atomic gravimeter. *Applied Physics B*, 92:133–144, 2008.
- [93] Bess Fang, Nicolas Mielec, Denis Savoie, Matteo Altorio, Arnaud Landragin, and Remi Geiger. Improving the phase response of an atom interferometer by means of temporal pulse shaping. *New Journal of Physics*, 20(2):023020, 2018.
- [94] S Templier, J Hauden, P Cheiney, F Napolitano, H Porte, P Bouyer, B Barrett, and B Battelier. Carrier-suppressed multiple-single-sideband laser source for atom cooling and interferometry. *Physical Review Applied*, 16(4):044018, 2021.
- [95] F Bloch and A Siegert. Magnetic resonance for nonrotating fields. *Physical Review*, 57(6):522, 1940.
- [96] Lyndon Emsley and Geoffrey Bodenhausen. Phase shifts induced by transient Bloch-Siegert effects in NMR. *Chemical physics letters*, 168(3-4):297–303, 1990.
- [97] Richard Stoner. Two-photon light shift and the AC Stark shift of NMR. Unpublished manuscript, 2024.
- [98] A. Gauguier, T. E. Mehlstäubler, T. Lévêque, J. Le Gouët, W. Chaibi, B. Canuel,

- A. Clairon, F. Pereira Dos Santos, and A. Landragin. Off-resonant Raman transition impact in an atom interferometer. *Physical Review A*, 78(4), Oct 2008.
- [99] Erik Østreng, Henrik H Sønsteby, Timo Sajavaara, Ola Nilsen, and Helmer Fjellvåg. Atomic layer deposition of ferroelectric LiNbO₃. *Journal of Materials Chemistry C*, 1(27):4283–4290, 2013.
- [100] Steven M George. Atomic layer deposition: an overview. *Chemical reviews*, 110(1):111–131, 2010.
- [101] Jian Liu, Mohammad N Banis, Xifei Li, Andrew Lushington, Mei Cai, Ruying Li, Tsun-Kong Sham, and Xueliang Sun. Atomic layer deposition of lithium tantalate solid-state electrolytes. *The Journal of Physical Chemistry C*, 117(39):20260–20267, 2013.
- [102] Hannu Korhonen, Aki Syväluoto, Jari TT Leskinen, and Reijo Lappalainen. Optically transparent and durable Al₂O₃ coatings for harsh environments by ultra short pulsed laser deposition. *Optics & Laser Technology*, 98:373–384, 2018.
- [103] A Carmel Mary Esther, N Sridhara, Shaji V Sebastian, Parthasarathi Bera, Chinnasamy Anandan, Singanahalli T Aruna, Dinesh Rangappa, Anand Kumar Sharma, and Arjun Dey. Optical and RF transparent protective alumina thin films. *Journal of Materials Science: Materials in Electronics*, 26:9707–9716, 2015.
- [104] CE Kennedy, RV Smilgys, DA Kirkpatrick, and JS Ross. Optical performance and durability of solar reflectors protected by an alumina coating. *Thin Solid Films*, 304(1-2):303–309, 1997.
- [105] James A Oke and Tien-Chien Jen. Atomic layer deposition and other thin film deposition techniques: from principles to film properties. *Journal of Materials Research and Technology*, 21:2481–2514, 2022.
- [106] Christian Hirt, Sten Claessens, Thomas Fecher, Michael Kuhn, Roland Pail, and

- Moritz Rexer. New ultrahigh-resolution picture of Earth's gravity field. *Geophysical research letters*, 40(16):4279–4283, 2013.
- [107] Byron D Tapley, Srinivas Bettadpur, Michael Watkins, and Christoph Reigber. The gravity recovery and climate experiment: Mission overview and early results. *Geophysical research letters*, 31(9), 2004.
- [108] Nikolaos K Pavlis, Simon A Holmes, Steve C Kenyon, and John K Factor. The development and evaluation of the Earth gravitational model 2008 (EGM2008). *Journal of geophysical research: solid earth*, 117(B4), 2012.
- [109] The University of Southern Mississippi. R/V Point Sur, Mar 2025. Available online at <https://www.usm.edu/vessels/rv-point-sur.php>.
- [110] Geoff Slocombe. RNZN's largest ever vessel joins the fleet. *Asia-Pacific Defence Reporter (2002)*, 46(8):34–36, 2020.
- [111] New Zealand Defense Force. HMNZS Aotearoa, Mar 2025. Available online at <https://www.nzdf.mil.nz/nzdf/our-equipment/ships-and-watercraft/hmnzs-aotearoa/>.
- [112] Peter Dombrowski and Simon Reich. Multilateral maritime exercises, grand strategy, and strategic change: The American case and beyond. *Journal of Global Security Studies*, 9(3):ogae017, 2024.