APPARATUS FOR INERTIAL SENSING WITH COLD ATOMS

ADAM CURTIS CAREW

A DISSERTATION SUBMITTED TO THE FACULTY OF
GRADUATE STUDIES
IN PARTIAL FULFILMENT OF THE REQUIREMENTS
FOR THE DEGREE OF

DOCTORATE OF PHILOSOPHY

GRADUATE PROGRAM IN PHYSICS AND ASTRONOMY
YORK UNIVERSITY
TORONTO, ONTARIO
AUGUST 2018

Abstract

A variety of experimental techniques and equipment for the measurement of inertial effects are herein presented. The bulk of the work relates to improvements to an existing local gravitational acceleration "little-g" measurement apparatus. These improvements are predicted to push the statistical uncertainty in the measurement of g to less than 1 part-per-billion (ppb). To accomplish this goal, several other projects were undertaken. These include a finite-element model of the magnetic field coil setup used in the experimental apparatus, as well as the design and construction of a hermetically-sealed diode laser system with excellent long-term frequency stability. Additionally, a direct digital synthesis-based frequency generator was designed and built for a proposed frequency-domain atom interferometer experiment. Finally, a side-project involving the evaluation of the magnetic field uniformity/stability of a commercial optical isolator was performed, and its results are presented as an appendix.

Acknowledgements

Here, I would like to spend a few words to acknowledge the people without whom this thesis would never have happened.

I would like to start by thanking the support and technical staff of York University, particularly those working in the basement of the Petrie building. Thanks to the student machine shop supervisors that have come and gone in my time at York: Ator, John, Frank and Jerry. These guys helped with the fabrication of countless bits of experimental equipment, and modifications to the same. Big thanks too to Dirk from the glass shop for his help with our custom glasswork. It is vital that a physics or engineering department have support in these areas, and I always considered myself very lucky to have such great shops in the building.

In the realm of administration, I would like to point out the contributions of Marlene and Lauren, who made the department run smoothly. In particular, I would like to thank Marlene for all of her help with administrative hurdles throughout my long time at York, all the way from my last few years of undergrad, through my Masters' and to the nearly the end of my time as a PhD student. Thanks to Cristal for helping me across the finish line after Marlene's retirement.

My supervisory committee has been extremely supportive through my time at York, especially in the past few years as I worked to finish up my writing. Eric and

Roman, thank you for helping me realize the importance of finishing this work, and supporting me through a rough transition.

To the members of my extended thesis committee; Regina Lee, Tom Kirchner, and Amar Vutha, thank you all very much for your time and consideration.

My coworkers and lab-mates represent a cohort of fantastically talented and wonderful people. I have had many satisfying and beneficial professional relationships in my time at York, and what follows is a (probably incomplete) list of those people I have worked with. Thanks to: Matt Weel, Matt George, Dan Fitzackerly, Dan Comeau, Yaniv Gura, Francois Garofalo, Adam Bebko, Brynle Barrett, Robert Berthiaume, Iain Chan, Carson Mok, Hermina Beica, Kosuke Kato, and Alex Pouliot for the help, discussions, collaboration and friendship.

To Kumar, I would like to say thanks for all the support and encouragement along the way. I know that the trajectory of my degree was perhaps not the one that you had in mind when you invited me into your lab all those years ago, but I think that we can agree that things have turned-out alright in the end.

I would also like to thank my coworkers and colleagues at Fluidigm, in particular Dmitry Bandura and Sasha Loboda, for pushing me to finish this degree and supporting my efforts along the way. Fluidigm has been a fantastic environment for me, and I've truly enjoyed my time working there so far.

Thanks to my parents for creating me, nurturing my mind and providing a safe and stable home during my childhood.

To Olga, thanks for your unending support and understanding through this process.

Table of Contents

List of Figures

Abbreviations

Abbreviation	Long form
AI	Atom interferometer
AOM	Accousto-optic modulator
BEC	Bose-Einstein condensate
BS	Beam splitter
COM	Center of mass
DDS	Direct digital synthesis
DPAOM	Dual-pass accouso-optic modulator
ECDL	External cavity diode laser
EM	electromagnetic
FFW	Feed-forward
FWHM	Full width at half maximum
GECDL	Grating-stabilized external cavity diode laser
GSRF	Ground state Ramsey fringe
HWHM	Half width at half maximum
IF	Interference filter
IFECDL	Interference filter-stabilized external cavity diode laser
LIA	Lock-in amplifier
LIDAR	Light detection and ranging
MOT	Magneto-optical trap
PBS	Polarizing beam splitter
PID	Proportional-integral-derivative controller
PLL	Phase-locked loop
PMT	Photomultiplier tube
ppb	Parts per billion
ppm	Parts per million
RF	Radio frequency
RO	Read out
SNR	Signal-to-noise ratio
SW	Standing wave
TW	Traveling wave
VCO	Voltage-controlled oscillator

1 Outline

Inertial sensing comprises a variety of techniques in which properties of a physical system are used to monitor some aspect of the system's movement. Some classes of devices that fall under the umbrella of inertial sensing include seismometers, gyroscopes and accelerometers. A simple example of an accelerometer is a cantilever-style device, which uses deflection of a weighted spring to determine the magnitude of the acceleration being applied to the instrument [1]. The deflection of the spring can be measured by monitoring capacitance of a variable capacitor, one side of which is attached to the moving arm, or by monitoring the voltage produced by a piezo-electric element which is stressed by the movement of the arm. In both cases, deflection of the cantilever generates a signal that is proportional to the degree of deflection. The downside of this approach is that the spring must be carefully calibrated using a more precise accelerometer. Typical performance of a high-quality cantilever-based relative gravitometer (Scintrex CG-3M) yields an accuracy of 5 μGal over a range of 7000 mGal with a resolution of 1 μGal [2]. The unit of 1Gal = 10^{-2}m/s^2 is commonly-used in the gravity sensing community as a shorthand to discuss acceleration values.

In general, inertial sensors have a wide range of applications. Seismometers are used to monitor tectonic activity. Accelerometers are found in cell phones,

for orientation and position sensing, as well as in navigation chips for aerospace vehicles. Gyroscopes are often found in large cargo ships, where they are used for dead-reckoning navigation. Of particular interest to this work are local gravitational acceleration (g) sensors. These instruments frequently find use in prospecting for sub-surface deposits of natural resources [3].

While opto-mechanical inertial sensing is currently the industry standard, cold-atom-based techniques are rapidly approaching and in some cases exceeding the performance of these established methods in terms of achievable measurement precision. In these devices, a sample of atoms, usually alkali-earth elements, are trapped with a combination of specially-tuned light and magnetic fields. In this way, a cloud containing up to $\sim 10^9$ atoms can be collected and cooled to lower than 50 μK [4]. At these low temperatures, the quantum nature of the atoms can be exploited to extract information about the accelerations applied to them during an experiment.

In this work, the focus is on the construction of an experimental apparatus which uses cold atoms for measurements of g as well as measurements of Earth's rotation, Ω_E. Some preliminary data is presented, along with the corresponding theoretical background. Additionally, an experiment exploring the interaction of cold atoms with short light pulses is presented, as this behaviour is key to understanding the finer points of any inertial measurement that uses the techniques herein. Finally, a prototype laser development project is presented, in which a novel variation on an existing technology is implemented to achieve high frequency-stability of the laser's output, both on short and long timescales. This laser is intended to be used by our group in further research into cold atoms and atom interferometry, with possible applications outside our group in LIDAR, atomic spectroscopy and gravimetry.

1.1 Atom Interferometry

Atom interferometry (AI) has been a well-established field since the work done on atomic beams in the mid-20th century, involving excitation of atoms by microwave [5,6] and later optical [7] fields, or diffraction of atoms by micro-fabricated gratings [8, 9]. In these beam-based experiments, an oven would be used to produce a beam of thermal atoms, traveling at a few hundreds of metres per second. By carefully collimating the beam with apertures, the temperature of the atoms along the transverse axis (i.e., perpendicular to the direction of the beam's travel) could be lowered sufficiently to allow for these atoms to be probed successively by different fields. With the development of the Magneto-optical trap (MOT) in 1987 [10], the experimental timescale available to researchers increased greatly thanks to the very slow movement of the atoms confined by the MOT. Further works demonstrating atomic fountain techniques [11], cold atom interferometers [12] and cold atomic clocks [13] laid the groundwork for long-timescale atom-interferometry. This is important for gravity sensing experiments, because the effect on the phase of the interferometric signal due to an applied acceleration grows as the free-fall time squared. Therefore, the easiest way to increase the sensitivity of the system to accelerations is to increase the timescale of the experiment.

A cold atomic sample is not the only requirement for inertial sensing. In each of these experiments, it is necessary to split the atomic sample into two or more groups of atoms by application of light pulses, and to have these groups travel along separated trajectories before recombining them at a later time. Because the atom packets travel along different paths to their rendezvous, each state's wavefunction

will accumulate a different phase, which produces interference at the time of recombination. By measuring the variation in state amplitude at this point with respect to the pulse separation times, the acceleration can be extracted.

For our experiments, the phenomenon of Kapitza-Dirac diffraction [14] is used. Kapitza-Dirac diffraction was initially predicted for electrons, but was not observed until recently [15], due to the weak interaction of electrons with light. This effect is now most-commonly used to describe the diffraction of a sample of atoms by a near-resonant light beam, where the duration of the light-atom interaction is small enough that atomic motion during the interaction can be disregarded. This requires $\tau \ll 1/\omega_q$, where τ is the interaction time and $\omega_q = \hbar q^2/2M$ is the two-photon atomic recoil frequency, with $q = 2k$ being the wavenumber of the standing-wave (SW) excitation whose traveling wave components each have wavenumber k, and M is the atomic mass. Kapitza-Dirac diffraction was first demonstrated on atoms in 1986 [7], where it was used to split an atomic beam into several diffracted beams, each differing in momentum by an integer value of $\hbar q$. As shown in fig. (1.1), and as will be discussed in ch.2, the population of each momentum state after the diffraction is proportional to $|J_n(\Theta)|^2$, where Θ is the pulse area, which defines the strength of the atom-light interaction, and J_n is a Bessel function of the first kind of order n.

With the tools of atomic interferometry and magneto-optical trapping, it is now possible to trap and cool a sample of atoms before splitting and recombining them using pulses of light. By setting up the experimental geometry such that the atoms are separated along the same direction that the acceleration is applied, a measurement of this acceleration is possible [12,16,17]. If instead, the interferometer

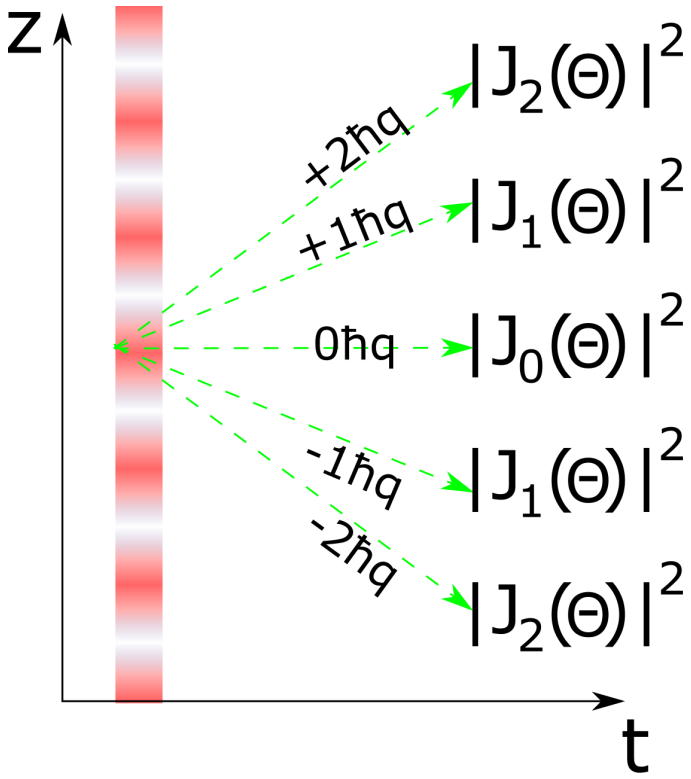

Figure 1.1: Recoil diagram illustrating the effect of Kapitza-Dirac diffraction on a cloud of cold atoms. The atom cloud, initially at rest, is diffracted into a set of momentum states by a standing wave pulse, each state having a momentum that is an integer multiple of $\hbar q$. Only the first two orders of momentum states are shown for clarity. Each momentum state's population as a fraction of the number of atoms is given by $|J_n(\Theta)|^2$, where Θ is the pulse area of the SW excitation.

is set up in such a way that the physical area enclosed by the interfering trajectories is normal to a rotation vector, the magnitude of that rotation can be measured [18]. In either case, extreme care must be taken to minimize mechanical vibrations of the apparatus which at best can cause loss of signal by washing-out fringe contrast, and at worst contribute a spurious component to the measured acceleration that cannot be separated from the desired signal.

1.2 Gravimetry

1.2.1 Opto-mechanical sensors

Measurements of local gravitational acceleration give information about the composition of Earth's crust at the point of measurement. A density map can be generated by repeating measurements of g on a grid pattern, allowing prospectors a method for detecting possible deposits of a desired resource [19]. In field applications of gravimetry, either a completely opto-mechanical, falling corner-cube device is used for absolute measurement of g [3], or a capacitative device [2] is used for a relative measurement. The capacitative device uses a calibrated spring which supports a mass connected to a capacitor plate. Changes in the magnitude of local gravity result in expansion/contraction of the spring, changing the capacitance of the device [20].

The absolute, falling corner-cube device is, in essence, a Mach-Zehnder-type interferometer, with one of its arms being composed of a light beam that is retro-reflected from a falling corner-cube retro-reflector [3]. As the corner-cube falls, this type of interferometer will exhibit periodic maxima in transmission at the detector

corresponding to instances of constructive interference. These maxima will occur at a spatial period of $\lambda/2$ in the motion of the corner cube, where λ is the wavelength of the light being used (~ 0.5 μm for visible light). Therefore, at the output of the interferometer, as the corner cube falls, successive maxima in transmitted intensity will be generated at frequency $2gt/\lambda$, where t is the elapsed time. Since the frequency of fringe generation grows linearly with t, the accumulated phase grows quadratically as the corner cube falls. State-of-the-art falling corner cube gravimeters claim precision of 0.15×10^{-6} m/s$^2/\sqrt{\text{Hz}}$, which means that they can achieve a 1 part per billion (ppb) measurement of g in about 3.75 minutes [21]. Longer measurement times produce higher precision, with strongly-diminishing returns (to get an improvement of a factor of 10 in precision requires a data set that is 100 times longer). It is important to consider the length of the experiment, as diurnal variations in the value of g due to tides and other effects must be accounted-for in long-term experiments. Tidal variations in g are at the level of 50 ppb in 1 hour.

The primary advantage of the absolute device over the capacitative device is that the absolute device produces a direct measurement of gravitational acceleration, without needing an external calibration. In fact, falling corner-cube gravimeters are used to calibrate the capacitative field instruments. Conversely, the capacitative device benefits from being "always on", producing immediate measurement of g.

1.2.2 AI gravimetry

In a cold-atom gravimeter, the role of the falling retro-reflector is played by a sample of cold atoms. In our case, these atoms are a cloud of ^{87}Rb atoms. Unlike the opto-mechanical gravimeter, the atomic gravimeter's falling retro-reflector (the

atom cloud) is not always reflective. Normally, when illuminated by a beam of monochromatic light, a sample of cold atoms will do one of two things: in the case of a far off-resonant beam, the atoms will not interact with the light, in essence being completely transparent. In the case where the beam is within the Doppler-broadened linewidth of the atomic transition, the light will be absorbed by the atoms, and later re-emitted as spontaneous or stimulated emission. Because spontaneous emission is randomly-directed, this light is not usefully reflected from the atomic sample, and can be considered lost from the experiment. There is also the notable phenomenon of electromagnetically-induced transparency (EIT), in which the resonant interaction is saturated by an intense light pulse, rendering the atomic sample transparent to light that is resonant with an atomic transition [22, 23].

In order to obtain an atomic sample with a non-negligible reflectivity, it is necessary to subject the atoms to an atom-interferometric technique while they are in free-fall. The technique we use is called the grating-echo atom interferometer [24], which will be discussed in detail in ch.2. In brief, this interferometer functions by diffracting the atomic sample into a multitude of momentum states. This is accomplished with a short, off-resonant SW pulse, applied at $t = 0$, as shown in fig. (1.2). Because this pulse is detuned from atomic resonance, after each interaction with the SW field, the atoms are returned to their ground state, having acquired a momentum kick of $\pm 2\hbar k$ along the axis defined by the SW beams (usually vertical), where k is the wavenumber of the SW light. Atoms can undergo multiple interactions with the SW field during the pulse, resulting in a manifold of momentum states with unequal amplitudes, each having momentum quantized in units of $2\hbar k$. Because of the Pendellösung effect [25], the atoms begin to move towards

the nearest node in the SW potential. The Pendellösung effect is a term borrowed from Bragg diffraction from crystals, and here refers to the 1D, harmonic motion of atoms in a SW potential. The atoms are pushed towards the nodes in the potential, and oscillate at a characteristic harmonic frequency. This causes the formation of a temporary modulation in the density of the atom cloud that is washed-out in ~ 1 μs by thermal motion of the atoms ($v_{\mathrm{RMS}} = 5$ cm/s for Rb atoms at 10 μK). From this point, the different momentum states will continue to separate spatially until time $t = T_{21}$, when a second SW pulse is applied to the atoms. This pulse re-diffracts the atoms, causing some momentum-space trajectories to intersect near time $t = 2T_{21}$, which is called the read-out time. At time $t = T_{21}$, the respective centers of mass (CoMs) of these two momentum states are still separated by a distance of $z_{21} = 2\hbar k T_{21}/M_{atom}$.

It is in the vicinity of time $t = 2T_{21}$, called the "echo time", that we observe a reformation of the periodic density modulation in the atomic sample, caused by interference between the matter waves. A traveling-wave beam, called the read-out beam, is applied to the atoms at this time, and some of that light is reflected into a sensitive photodetector. This light contains information about the depth of the density modulation, which is proportional to the reflectivity of the grating, as well as its phase relative to a static point of reference (the retro-reflecting mirror that forms the SW). It is this last piece of information, the phase of the grating, that allows accurate determination of how far the atoms have fallen during the experiment. Knowing this phase allows a measurement of g, which is accomplished by measuring the reflected intensity and phase while varying T_{21}.

As an example, we can consider one set of intersecting trajectories, shown in

fig. (1.2), which contributes to the signal at $t = 2T_{21}$: after the second SW pulse, some of the atoms originally in the $p = +2\hbar k$ state will be diffracted into the $p = 0$ state. Similarly, some of the initially stationary atoms will be diffracted into the $p = +2\hbar k$ state. At time $t = 2T_{21}$, these states will again be coincident. There are many such trajectory pairs that intersect at the echo time. What determines which pairs contribute to the signal is their difference in momentum. Only states differing by $\pm 2\hbar k$ will produce a spatial frequency which will support Bragg scattering of the readout beam. A more detailed treatment of this theory is provided in ch.2.

This interferometer creates a spatially-modulated density in the atomic cloud, effectively a temporary diffraction grating. This modulation is very slight, having a Debye-Waller factor of ~ 0.1, resulting in low grating contrast. The Debye-Waller factor, $\beta = e^{-2k_L^2\omega_z^2}$, where k_L is the lattice vector of the atom cloud and ω_z is the width of each maximum in the density distribution, is a concept from solid-state physics that is used in this context to quantify the depth of the density modulation of the atoms [26]. By taking advantage of the Bragg effect, our technique is able to extract usable signals, even from such a weak reflector. The Bragg condition requires that successive maxima in the density of the atom cloud be separated by $\lambda/2$. When this occurs, the grating is an effective reflector for the incident light. However, one major shortcoming of this approach is that the density modulation in the atomic cloud is only present for a few microseconds, due to the continued motion of the atoms after the second SW pulse. Additionally, the atom cloud cannot be re-used in the experiment, as the non-zero temperature of the cloud results in its rarefaction below a useful density. For each repetition of the experiment, the atom cloud must be re-generated by trapping of background vapour. This limits

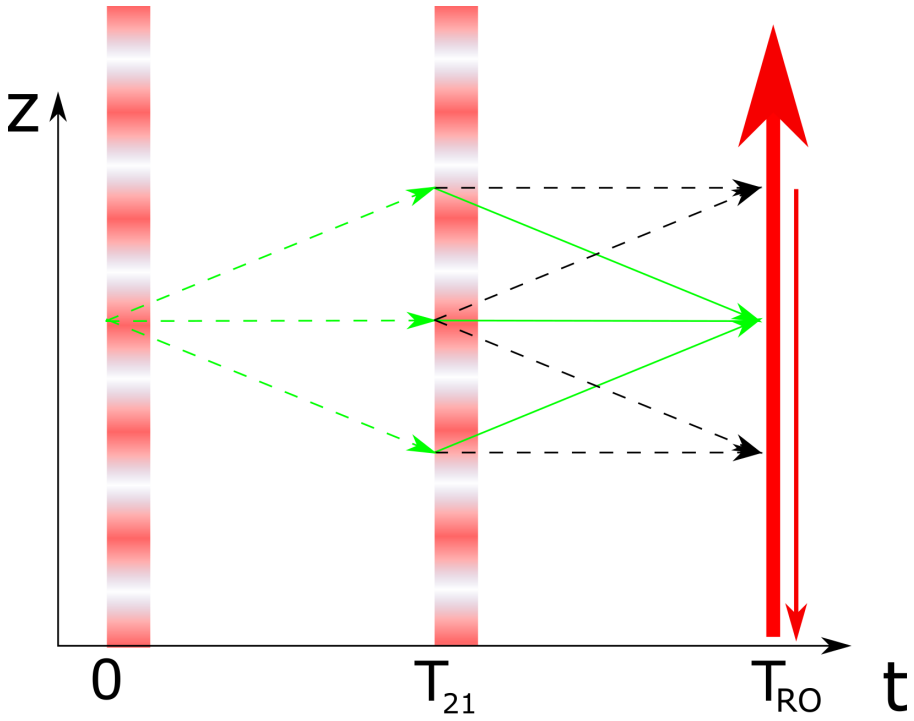

Figure 1.2: Simplified recoil diagram of the echo AI. SW pulses are applied to the atoms at times t=0 and t = T_{21}. Only a limited subset of the momentum states are shown, with four possible interferences. The paths marked in green form one pair of (coincident) interferences, and the dashed paths form two others. At time t = $T_{RO} \simeq 2T_{21}$, a traveling-wave read-out pulse is applied, which constructively back-scatters from the evanescent structure present in the atom cloud around this time. The time between the two SW pulses is varied to generate the experimental signal, which contains information about the acceleration of the atom cloud under gravity. Note that in this diagram, the effect of gravity is not shown. It would convert the straight-line atom trajectories into parabolic ones.

the maximum repetition rate of the experiment, as the background pressure of ^{87}Rb has to be kept low to limit the deleterious effect of collisions between trapped atoms and background gas atoms. In our work, we were unable to extract signals with high-enough signal-to-noise (SNR) for extraction of a measurement of g at repetition rates above 5 Hz. This limits the number of T_{21} values that can be used in a given measurement, as diurnal variations in the local environment coupled with tidal effects cause variations in the nominal value of g. For this reason, it is not possible to simply take more data to improve precision, as the quantity being measured is varying appreciably in the time it takes to conduct a measurement.

1.2.3 Studies of grating formation and optical lattice effects

If the first SW excitation is particularly long, which is a special case called the Bragg regime of atom interferometry [27], the atoms will be channeled to the nodes of the SW potential, and the resulting momentum states will be limited to $\pm 2\hbar k$. If this long SW excitation is applied at the same time as the MOT fields, the result is a stack of small traps, each separated by $\lambda/2$. This is usually referred to as an optical lattice [28–31]. Since the grating echo AI is so dependent on the formation of a modulated atomic density, studying the formation of atomic lattices and the transition between Kapitza-Dirac and Bragg regimes can throw some light on the underlying mechanisms, and reveal the steps necessary to improve contrast. In the Bragg regime, the SW pulses are long, resulting in a low spread of photon energy/momentum. This is due to the pulse-bandwidth effect, a particular form of the Heisenberg uncertainty principle. For such long pulses, only atomic momentum states having $\pm 2\hbar k$ will be populated. By contrast, the Kapitza-Dirac scattering

regime is one for which the pulses are short enough in duration to produce a wide spread in photon momentum. As a result, many atomic momentum states are excited, each having a momentum of an integer multiple of $2\hbar k$ [25].

Reference [30], which shows a 4-fold improvement in signal amplitude over the standard 2-pulse echo AI after pre-loading the atoms into an atomic lattice, is a result of particular interest related to grating formation dynamics.

To probe these effects, preliminary measurements of the formation of a density modulation in a cloud of trapped atoms were performed. These results are presented in ch.3. By applying a SW pulse followed immediately by a traveling-wave read-out pulse, the transient formation of the initial density modulation was observed. We referred to this as the "one-pulse" echo. Effects of SW detuning and pulse duration were explored, as well as effects relating to read-out pulse intensity. In the end, it was determined that the work required a phase-stable platform to generate SW potentials that would be coherent over several μs. Since this was not possible in the original experimental apparatus, which was not vibration stabilized, the work had to be postponed until the construction of the new apparatus for gravity measurements was completed.

1.2.4 Design and construction of a novel gravimeter

In our group's most recent attempt at a measurement of g, we obtained a statistical precision of 75 ppb [32, 33]. This experiment made use of a modified version of the grating echo AI to obtain this result. There were several shortcomings of this measurement, though, which needed to be addressed in the next iteration of the experiment. The following subsections outline the specific improvements

implemented in the context of this dissertation. More detail on these topics is presented in ch.4.

1.2.4.1 Non-magnetic construction

The work presented in Ref. [32] was performed in a stainless steel vacuum chamber. As our technique doesn't pre-select atoms into the field insensitive magnetic sublevel ($m_F = 0$) before the experiment, the atoms experience a force due to any stray magnetic fields and field gradients during the experiment. This causes decoherence of the atom cloud, and results in reduced signal lifetime. Some of this effect can be canceled by the application of external magnetic fields from an array of electromagnet coils, but the presence of ferromagnetic material in the vicinity of the atoms is problematic, due to the magnetizeability of these materials. Due to the pulsed magnetic field used in atom trapping, the chamber walls became slowly magnetized, resulting in a drift of the required field-cancellation settings. To overcome this limitation, the new gravimeter design made use of a borosilicate glass cell, in which the atoms were trapped and studied. This cell, coupled with careful experimental design maximized the distance between the atoms and any ferromagnetic materials.

1.2.4.2 Magnetic field modeling and coil design

The second improvement implemented in this experiment was in the design of the magnetic field canceling coils. In the old experimental cells, the magnetic field canceling coils were not designed according to any particular experimental criterion. Ideally, magnetic fields are canceled with Helmholtz coils. These coils consist of

a pair of identical circular cross section coils, separated by a distance equal to their radius. When the current flows through the two in the same direction, a zone of nearly-zero axial field gradient is produced at their mutual centre. If the currents run in opposite directions through the two coils, the result is called an anti-Helmholtz coil, which has a uniform field gradient at the centre. By combining a set of Helmholtz and anti-Helmholtz coils along each of three mutually-orthogonal directions, it is possible to cancel both magnetic field and magnetic field gradients at the geometric centre of the coil system.

In the old cells, only the coils acting along the vertical axis were positioned near the Helmholtz configuration (i.e.: coil pairs separated by their radius), with no consideration given to the fact that the cross-section of the coils was square rather than circular. It is much easier to build square cross-section coils at the scale required ($\sim$ 1 m diameter), but deviating from the ideal Helmholtz configuration causes deviation from ideal field character at the centre of the coils. To re-establish this condition, the coil separation had to be modified. The optimal coil spacing was determined by performing numerical simulation of the coils, which yielded results consistent with previous work in this field [34]. A set of three pairs of nested, mutually-orthogonal coils were then designed and built, to independently cancel stray fields in the vicinity of the trapping volume. The coils were built sufficiently-large to ensure that the zero-field/gradient volume was large enough to allow the atoms to experience the desired 100 ms drop time.

As a side project, a complete magnetic field model for a permanent magnet stack used in a commercial Faraday rotator was produced. A Faraday rotator is a key component in an optical isolator [35], which is a magneto-optical device that al-

lows light to propagate through it in one direction, while blocking the reverse path. These are very useful in applied optics, as diode lasers are particularly sensitive to optical feedback. By employing an optical isolator as an "optical diode", the influence of back-reflections is greatly reduced. By developing a finite-element model for the magnetic field of this isolator, it was possible to provide the manufacturer guidance in selecting new materials for the refinement of their design. A detailed breakdown of the approach employed in this project is presented in appendix (A).

1.2.4.3 Frame of reference motion post-correction

The third major improvement related to the frame of reference for the gravimeter. In all gravimeters, what is being measured is not gravity directly, but the acceleration of a falling mass with respect to a particular frame of reference, or the deflection of a mass supported by a spring of known strength. This allows for spurious accelerations to affect the measured output, as the device is incapable of discerning between the acceleration of the test mass under gravity and any movement of the apparatus as a whole during the measurement. For commercial capacitative gravimeters, it is usually sufficient to place the device on a firmly-rooted concrete piling driven into the ground. For an AI gravimeter, however, this approach is not possible. The reason for this is that the phase of the SW pulses is imparted onto the atomic wavefunction after each interaction. Since the phase of the SW is directly referenced to the spatial position of the retro-reflector (as there must be a node in the SW at the reflective surface due to electromagnetic boundary conditions), any vibrations that cause a movement of the retro-reflector will cause a corresponding movement of the SW. Therefore, it is necessary to perform

this measurement on a vibration-isolation platform to ensure phase stability. This technique is also used in ultimate-precision opto-mechanical gravimeters, such as the Fg5-x [21], which employs an active "superspring" vibration stabilizer. Mounting the experiment on such a platform means that we no longer have information about the position of the experiment with respect to Earth's gravity field. It is easy to show that small movements of the apparatus cause measurable changes in the value of g. Assume that the apparatus measures a value for the gravitational acceleration g_0 at a distance r_0 from Earth's centre. Now, since the apparatus is free to move, assume it takes its next measurement at $r_1 = r_0 + 1$ mm. Earth's mean radius is 6371 km, so if we assume that $r_0 = r_E$, this motion of the apparatus constitutes a 0.16 ppb change in the separation between our two masses (Earth and the atoms). Since the gravitational force between two bodies of masses M_1 and M_2 is given by $F_g = GM_1M_2/r^2$, we can calculate the fractional change in force caused by a shift in r. It will be given by $\delta r(\delta r + 2r_0)/(\delta r + r_0)^2$, which, for $r_0 = r_E$ and $\delta r = 1$ mm will be equal to a change of 0.31 ppb. As g is directly proportional to F_g, a shift of the experiment by 1 mm results in a 0.3 ppb change in the value of g recorded. This shift is comparable to the target precision of this measurement, namely 1 ppb, so it must be accounted for.

When dealing with effects due to gravity, it is usually necessary to examine the contribution of General Relativistic effects to see if they can be safely ignored. These effects include the differing "clock rates" seen by the atoms at different heights in the gravity field, as well as gravitational red-shift effects during the interferometer pulse sequence. Per reference [36], the GR effect in AI gravimeters is on the scale of $\sim 10^{-15}$g, which is much smaller than the expected statistical

precision of 1 ppb for this measurement, and thus can be safely ignored. Further advances in the precision of atomic interferometry will quickly result in these GR effects becoming measurable contributors to the total phase of the signal, but GR effects will be ignored for the rest of this work.

It is not only necessary to know the position of the apparatus at the start of an experiment, it is also necessary to know the acceleration of the reference frame during the experimental timeframe. In Ref. [16], the investigators constructed an active feedback loop that used signal from a seismometer placed on the retro-mirror to keep its position stable during the experiment. It has also been shown that it is possible to post-correct data to remove the effect of a moving reference frame [37–39], but only if there is good data tracking its position during the experiment. To that end, a commercial seismometer is used to track the motion of the reference frame. This is highly important, as a vibration with frequency 1 Hz and amplitude 0.1 mm will result in a peak acceleration of 0.0039 m/s^2, which is approximately 400 ppm of g. This is an overwhelmingly large acceleration when compared to the target precision of the device. It is for this reason that real-time monitoring of the frame-of-reference acceleration is so important. Statistical averaging will only help to smooth out the effects of completely random contributors to the signal. In this case, because we are dealing with the movement of a physical object on a spring, it is likely that some frequency component of the motion of the apparatus will be in phase with the repetition rate of the experiment. Therefore, there would be a persistent contribution to the acceleration that would not quickly average out.

With these improvements to the overall construction of the gravity experiment, further experimental work in this group is expected to produce high levels of sta-

tistical precision in measurements of g. With the planned increase in experimental timescale, the statistical precision should readily exceed 1 ppb, opening the door for an in-depth study of systematic effects.

1.3 The Sagnac effect and gyroscopes

Around the turn of the 20^{th} century, work was being done on a variety of interferometer arrangements, designed to detect the presence of the luminiferous ether. One such interferometer design was one published by Georges Sagnac in 1913 [40]. This interferometer uses a beam splitter to split incident light into two beams, which are directed oppositely along a common optical path. In modern applications, the common optical path is usually a fibre-optic cable of considerable (up to several km) length. In the case of a circular loop of fibre-optic cable rotating around its normal, the Sagnac effect can be conceptualized as a propagation delay experienced by the light that circulates in the same sense as the rotation. Similarly, for the light that propagates in the direction that runs against the rotation, the path will be slightly shortened when compared to the stationary case, as shown in fig. (1.3). These two effects add rather than cancel, resulting in a phase shift in the fringes developed at the output equal to $\delta\phi = \vec{\Omega} \cdot \vec{A}k/c$, where $\vec{\Omega}$ is the vector whose magnitude is the frequency of rotation, directed along the axis of rotation, $\vec{A}$ is the vector whose magnitude is the area enclosed by the loop, directed along the surface normal, k is the wavenumber of the light and c is the speed of light. While it is almost trivial to calculate the phase shift for a circular loop, it has been proven [41] that this relationship holds for any arbitrarily-shaped closed optical loop.

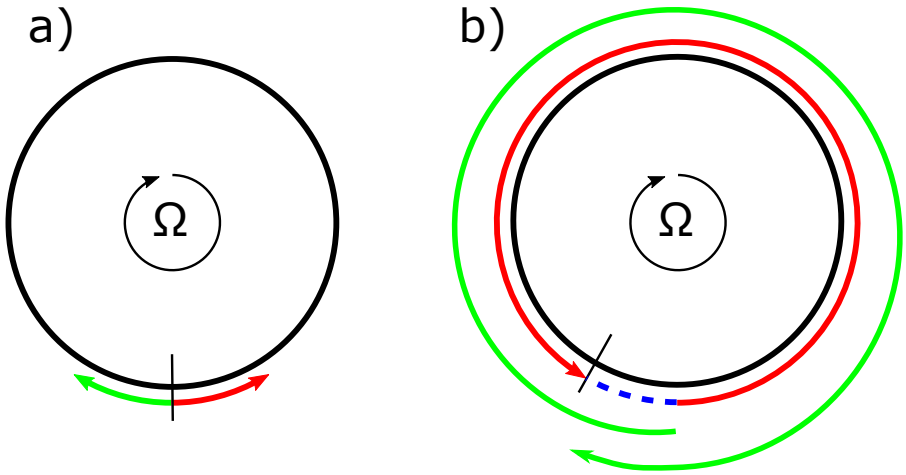

Figure 1.3: Illustration of the origin of the Sagnac phase shift. The closed-path of an optical interferometer is shown as the black circle, which rotates at the rate Ω, as shown. In part a), a pulse of light is split by the beam splitter (denoted by the short black line) with half of the light being directed in each direction through the interferometer. The beam that propagates along the direction of the rotation is denoted by the green arrow, and the other by the red arrow. In the time it takes the beams to transit the interferometer and rejoin one-another at the beam splitter, the interferometer has rotated a small amount, as denoted by the blue dashed line in part b). This path difference manifests itself as a phase shift in the output of the interferometer, and is proportional to the rotation rate.

Since, for fixed parameters, the phase shift at the interferometer output (called the Sagnac phase) is proportional to the rotation rate, measuring this phase accurately allows for determination of the rate of rotation. This is the basic function of a laser gyroscope. As with all interferometers, the same effects can be achieved with matter waves. Even the relevant quantities retain their functional forms, with the following substitutions: $c \to v$ and $k \to p/\hbar$, where v and p are the speed and momentum of the atoms, respectively. This results in a factor of $\sim 10^{10}$ increase in the size of the Sagnac phase for a given enclosed area. It is important to note that the area enclosed by a laser gyroscope is very large typically several m^2, constrained only by the desired package size. In contrast, typical cold atom gyroscopes have enclosed areas of only a few mm^2, resulting in a reduction in sensitivity by a factor of $\sim 10^6$. Even after this reduction, atom gyroscopes are still $\sim 10^4$ times more sensitive to the Sagnac shift than optical gyroscopes [42].

1.3.1 Experimental implementation of atom gyroscope

The planned experiment for measuring Earth's rotation is intended to proceed as follows. A sample of atoms will be trapped and cooled at the edge of a ~ 2.5 cm-wide SW excitation zone. The sample is prepared at the edge of the SW excitation zone in order to avoid needing multiple localized SW excitation beams through which the atoms would pass. In this way, the atoms travel across the spatial profile of the single SW beam, which is pulsed on and off to perform the experiment. This comes with the notable drawback of the atoms seeing differing interaction strengths (pulse areas) for the different SW interactions, due to the spatial profile of the beam. This can be compensated for by adjusting the power in the SW between the pulses

so that the local SW intensity in the position of the atoms is constant. Through the technique of moving molasses [11], the atoms will be launched horizontally across the excitation zone, where they will undergo the standard 2-pulse grating echo AI pulse sequence as described earlier in this document. By careful arrangement of the SW axis and the launch direction, the value of $\vec{\Omega} \cdot \vec{A}$ is maximized to the extent possible. Since Earth's axis of rotation does not pass through Toronto, the experiment will measure the component of Earth's rotation normal to the surface, whose value in Toronto is $\sim 5 \times 10^{-5}$ radians/s. With the available beam sizes and other experimental constraints, the maximum enclosed area of this interferometer will be ~ 0.75 mm^2, resulting in a phase shift in the output signal of 0.05 radians. This phase shift should be measurable in a single-shot measurement, using the read-out method from ref. [43]. By varying A, which is itself varied by T_{21} and the launch speed, it will be straightforward to extract Ω from the slope of the resulting $\delta\phi$ vs. A curve.

A preliminary implementation of this apparatus was constructed before it was realized that, for reasons of insufficient vibration stability, the experiment was not phase-stable enough to measure $\delta\phi$. It was determined that it wasn't feasible to correct this error, due to timing issues related to shared equipment, and the focus of research passed onto the gravity experiment. Although this measurement was not completed, its design and construction were very helpful in the subsequent work on the gravimeter described in ch.4, and the rotation measurement can potentially be completed in the new apparatus.

1.4 Laser Development

For much of the work done in any atomic physics lab, including the work presented in this thesis, it is necessary to use a laser with linewidth comparable to that of the atomic transition being probed. The laser linewidth represents the average deviation from the nominal frequency of the laser's output. The narrower the linewidth of the laser, the narrower the resonances that can be examined. State-of-the-art lasers for atomic physics research have linewidths in the sub-Hz range [44], which greatly overshoots the minimum linewidth needed to address typical atomic transitions ($\sim$ 1 MHz). These laser systems typically require very sophisticated external cavities to narrow the linewidth to such an extent, making them very expensive and difficult to implement. Most commercial laser linewidths are on the order of a few hundred kHz to 1 MHz, depending on the wavelength and technology in use, making them much more suitable for driving atomic transitions. In addition to linewidth concerns, the laser must also be stable on whichever timescale that is required in the application. Locking the laser to an atomic transition with a lock-in amplifier (LIA) [45, 46] stabilizes the output frequency to a certain degree, but the limitations of locking techniques make careful laser design a requirement. To this end, we designed a novel, hermetically sealed external-cavity diode laser (ECDL). Further motivating our need for low-linewidth, high-stability light sources was the failure of the titanium:sapphire laser that had been the primary light source in the lab since its formation, and had enabled all of the experiments performed pre-2015.

This particular ECDL, in contrast to the popular Littrow design [47], uses an interference filter [48–50] rather than a diffraction grating to select a cavity mode.

When we discuss the output of a laser, we often speak of the contrasting cases of multi-mode and single-mode operation. Single-mode operation is usually desired for atomic spectroscopy and trapping, and means that the laser is operating at a single frequency, with a given linewidth. Since laser diodes emit a large number of closely-spaced (in frequency) modes, it is necessary to select one of them to be preferentially amplified if single-mode operation is desired. As soon as the gain for one of the modes is made higher than that for the others, it becomes dominant in the output of the laser [51]. It is for this reason that ECDL's are used when narrow linewidths are required from diodes. Preliminary tests of this laser design have shown its linewidth to be in the vicinity of 1 MHz, which is a very competitive linewidth among the lasers at this price-point. A number of prototypes based on our design have realized signals relevant to gravimetry, magnetometry and precision measurements of atomic lifetimes.

Work in our research group [52] has shown that there is a direct correlation between atmospheric pressure and diode laser output frequency. This effect is due to some combination of index of refraction changes in the external cavity, and mechanical deformation of the laser diode package, which is a sealed chamber containing a gas. In the first case, the wavelength of the light in the external cavity is altered by $\lambda = \lambda_0/n$, where λ is the wavelength of the light in the cavity, λ_0 is the vacuum wavelength of the light, and n is the index of refraction, which is proportional to the pressure of the gas that fills the cavity (ie: the atmospheric pressure). In the case of mechanical deformation of the diode, any change in the length of the diode package will alter the mode spacing of the diode's output, thereby changing the modes present in the cavity. To remove these effects, we

needed to isolate the external cavity from atmospheric pressure fluctuations. This was accomplished by my design of a sealed enclosure for the laser, incorporating electrical and mechanical feedthroughs, which allow the cavity alignment to be adjusted without opening the case. Additionally, the laser case is equipped with a pump-out port, allowing the enclosure to be evacuated to a level of ~ 1 mTorr. This will allow for maximum thermal insulation of the laser internals from the environment, while making absolutely sure that the laser's frequency cannot be affected by pressure fluctuations. When combined with an auto-locking controller [53], this laser is expected to remain on atomic resonance for days at a time. The laser prototype currently operates at 780 nm and 633 nm, with other wavelengths being relatively simple to implement provided a laser diode and optical coatings at the required wavelength exist. In addition, the 780 nm laser can be used to seed a tapered amplifier, allowing it to be used as a high-stability light source for falling corner-cube gravimeters and cold atom experiments using rubidium MOTs. The 633 nm laser has performance comparable to He-Ne lasers, as shown in an experimental trial on a gravimeter [54]. A detailed description of the laser, including preliminary data is presented in ch.5.

1.5 RF Phase-locked loop frequency synthesizer development

To generate the optical frequencies required to address the transitions of interest in these experiments, it is necessary to use radio frequency (RF) devices. These devices operate in the 10-1000 MHz range, and they require input signals that are as

stable as possible. For example, when generating the light for the excitation pulses in a cold atom experiment, if the RF source is unstable over the experimental timescale, the signal phase will be similarly unstable. Therefore, a phase-locked loop-based frequency synthesizer was developed, which would take as input a stable 10 MHz signal from a commercial rubidium oscillator (Stanford Research Systems PRS10), and output a signal in the vicinity of 250 MHz, tunable in steps of 1 mHz. A PLL has the desirable property of a phase-stable output, regardless of the output frequency, unlike most other types of frequency synthesizer. The output of the PLL-based synthesizer was shown to have similar stability to the commercial clock, and allowed a great degree of flexibility in choosing the operating frequencies for experiments. Additionally, this equipment would allow for experiments to be performed in the frequency domain [55], rather than in the time domain. Here, a frequency domain experiment refers to one in which the time spacing between the SW pulses is fixed, but a frequency difference between the two TW components is added which can be varied to produce the signal. There are some types of experiment that can benefit from this approach, and they will be discussed in ch.6 along with the design of the synthesizer.

1.6 Thesis objectives

My purpose in pursuing this research has been to lay the groundwork for a precision measurement of gravitational acceleration using cold atoms and the echo AI technique. To that end, I have designed a purpose-built AI experimental cell, making full use of numerical simulation methods and CAD tools to ensure a minimum of

ferromagnetic materials in the vicinity of the atoms. Great care was taken to ensure that the apparatus was properly vibration-isolated, using a combination of laminar-flow pneumatic isolation and tuned-spring platforms. The magnetic field/gradient canceling coils were designed and modeled to ensure good uniformity of magnetic field in the experimental volume containing the atoms.

I conducted experiments to explore the mechanisms underlying density grating formation effects, and the effect of long SW pulses on the atomic density. Good understanding of these effects are necessary for future experiments involving optical lattice preloading of the atoms before the AI pulse sequence, which are expected to increase the signal amplitude by a large factor, which will, in turn, improve the SNR of gravimetric AI phase measurements.

Finally, I designed several pieces of experimental apparatus necessary for future experiments in the lab, including an interference-filter stabilized diode laser, Faraday isolators and a PLL-based frequency synthesizer. The laser exhibits a sub-MHz linewidth and good long-term stability, making it an ideal tool for atom trapping and AI experiments. The numerical model of the Faraday isolator improves our understanding of the home-built isolators in our lab, and shows ways to improve them. The work on the PLL-based frequency synthesizer sets the stage for future experiments that require high-stability frequency sources, including an AI gyroscope.

1.7 Long-term research objectives

Ultimately, we aim for a precision measurement of gravitational acceleration with 1 ppb statistical uncertainty. This would put the statistical precision of the cold atom-based gravimeter at a level comparable to that of the state-of-the-art optomechanical gravity sensors. At that point, it would be possible to run the two gravimeters side-by-side and compare performance between them in real time to characterize systematic effects in the cold-atom sensor.

Author contributions

Publications

1. A. Pouliot, H. Beica, **A. Carew**, A. Vorozcovs, G. Carlse, B. Barrett and A. Kumarakrishnan. Investigations of Optical Pumping for Magnetometry using an Auto-locking Laser System. *Laser Technology for Defense and Security XIV*, Proc. SPIE 106370A, 2018

2. A. Pouliot, H. Beica, **A. Carew**, A. Vorozcovs, G. Carlse and A. Kumarakrishnan. Auto-locking Waveguide Amplifier System for Lidar and Magnetometric Applications. *High-Power Diode Laser Technology XVI*, Proc. SPIE 105140S, 2018

3. H. Beica, **A. Carew**, A. Vorozcovs, P. Dowling, A. Pouliot, G. Singh and A. Kumarakrishnan. An Auto-locked Diode Laser System for Precision Metrology. *High-Power Diode Laser Technology XV*, Proc. SPIE 100860W, 2017

4. H. Beica, **A. Carew**, A. Vorozcovs, P. Dowling, A. Pouliot, B. Barron and A. Kumarakrishnan. An Auto-locked Laser System for Precision Metrology and Lidar Applications. *Laser radar technology and applications XXII*, Proc. SPIE 101910K-101910K-6, 2017.

5. B. Barrett, **A. Carew**, H. Beica, A. Vorozcovs, A. Pouliot and A. Kumarakrishnan, Prospects for Precise Measurements with Grating-Echo Atom Interferometry. *Atoms*, 4(19), 2016, MPDI, Special Issue on Atom Interferometry.

6. B. Barrett, **A. Carew**, S. Beattie and A. Kumarakrishnan. Measuring the Atomic Recoil Frequency using a Modified Grating-Echo Atom Interferometer. *Physical Review A*, 87:033626, 2013.

7. C. Mok, B. Barrett, **A. Carew**, R. Berthiaume, S. Beattie and A. Kumarakrishnan. Demonstration of Improved Sensitivity of Echo Interferometers to Gravitational Acceleration. *Physical Review A.*, 88(023614):1-15, 2013.

8. B. Barrett, I. Chan, C. Mok, **A. Carew**, I. Yavin, A. Kumarakrishnan, S. Cahn and T. Sleator. Chapter 3 - Time-Domain Interferometry with Laser-Cooled Atoms. In *Advances in Atomic, Molecular, and Optical Physics*, volume 60, pages 119-199. Elsevier Inc., 2011.

9. B. Barrett, S. Beattie, **A. Carew**, I. Chan, C. Mok, I. Yavin and A. Kumarakrishnan. Progress Towards a Precision Measurement of the Atomic Recoil Frequency using an Echo-type Atom Interferometer. Proc. SPIE 7993, 79930Y, 2010

Conference presentations

1. A. Pouliot, H. Beica, **A. Carew**, A. Vorozcovs, G. Carlse and A. Kumarakrishnan. Auto-locking Waveguide Amplifier System for Lidar and Magnetometric Applications, *SPIE Photonics West*, San Francisco, January 2018

2. A. Pouliot, H. Beica, **A. Carew**, A. Vorozcovs, G. Carlse, B. Barrett and A. Kumarakrishnan. Auto-locking Diode Laser System for Lidar and Magnetometric Applications, *SPIE Defense + Commercial Sensing*, Orlando, April 2018

3. G. Carlse, A. Pouliot, H. Beica, **A. Carew**, A. Vorozcovs, B. Baron and A. Kumarakrishnan. Development of a High Power Pulsed Laser System, *CUPC*, Ottawa, October 2017

4. H. Beica, **A. Carew**, A. Vorozcovs, P. Dowling, A. Pouliot, B. Barron, and A. Kumarakrishnan. Applications of Auto-locked Diode Lasers, *Canadian-American-Mexican Graduate Student Physics Conference*, Washington DC, August 2017

5. H. Beica, **A. Carew**, A. Vorozcovs, P. Dowling, A. Pouliot, B. Barron, and A. Kumarakrishnan. An Auto-locked Laser System for Precision Metrology and Lidar Applications, *SPIE DCS Defense + Security*, Anaheim, April 2017

6. H. Beica, **A. Carew**, A. Vorozcovs, P. Dowling, A. Pouliot, G. Singh, and A. Kumarakrishnan. Auto-locked Laser Systems for Precision Metrology, *SPIE Photonics West*, San Francisco, January 2017

7. **A. Carew**, B. Barrett and A. Kumarakrishnan. Rotation Measurement using a Grating Echo Interferometer, *APS DAMOP*, Anaheim, June 2012

8. B. Barrett, **A. Carew** and A. Kumarakrishnan. Measuring the Atomic Recoil Frequency with a Grating Echo Atom Interferometer, *APS DAMOP*, Anaheim, June 2012

9. C. Mok, **A. Carew**, B. Barrett, R. Berthiaume and A. Kumarakrishnan. Measurements of Gravitational Acceleration from an Echo Atom Interferometer, *APS DAMOP*, Anaheim, June 2012

10. I. Chan, B. Barrett, **A. Carew**, C. Mok, and A. Kumarakrishnan. Precise Determination of Atomic g-factor Ratios from a Dual Isotope Magneto-optical Trap, *APS DAMOP*, Anaheim, June 2012

11. B. Barrett, **A. Carew** and A. Kumarakrishnan. Measuring Atomic Recoil using a Grating Echo Atom Interferometer, *Canadian Institute for Advanced Research, Meeting on Cold Atoms*, Banff, February 2012

12. B. Barrett, **A. Carew**, S. Beattie, I. Chan, C. Mok, R. Berthiaume and A. Kumarakrishnan. Progress Towards a Precision Measurement of Atomic Recoil Frequency Using an Echo Interferometer, *APS DAMOP*, Houston, May 2010

13. B. Barrett, I. Chan, C. Mok, **A. Carew**, R. Berthiaume, A. Kumarakrishnan and I. Yavin. Simulations of a Multi-level Atom Interferometer, *APS DAMOP*, Houston, May 2010

14. I. Chan, R. Berthiaume, B. Barrett, C. Mok, **A. Carew**, and A. Kumarakrishnan. Precision Measurements of Atomic g factor Ratios using a Dual Species MOT, *APS DAMOP*, Houston, May 2010

15. **A. Carew**, R. Berthiaume, C. Mok, I. Chan, B. Barrett, M. Weel and A. Kumarakrishnan. Realization of an Inexpensive Multi-Channel Frequency Synthesizer, *APS DAMOP*, Houston, May 2010

2 The grating echo atom interferometer

With Louis de Broglie's 1924 postulation that electrons have a wavelength inversely proportional to their momentum, the early framework for matter wave interferometry was laid. Contemporary experimental work by Clinton Davisson and Lester Germer showed that it was true that electrons, heretofore understood to be "solid", indivisible particles, possessed wavelike properties. These ideas have since been expanded to include all matter. Even comparatively massive objects like atoms can be recast in a framework of wavelike properties. Particles having wavelengths opened-up the possibility for experiments and technologies that make use of the unique properties of waves. It is now commonplace to discuss the properties of a matter wave interferometer, and applications thereof.

2.1 Atom interferometry overview

To generate an atom interferometer, a sample of atoms has to be prepared in such a way that the quantum state of the atoms can be expected to last the duration of the experiment. In practice, this means that the atoms need to be cooled to the point that they remain in the region illuminated by the AI beams for long enough to be interrogated. This is the so-called transit-time limitation on experimental

timescales, and it is the principle effect limiting the length of experiments using cold atoms. Much work has been done to address this limitation, principally focused on the use of atomic fountains, which launch the atomic sample such that it transits the experimental volume over a longer time (typically about 2 times the transit time of non-fountain experiments) [11]. Additionally, energetic cold-cold or cold-hot collisions between atoms will cause electronic transitions, energy level shifts and state changes, resulting in loss of signal. These effects are why most atom interferometric experiments make use of magneto-optical cooling to reduce the temperature of the sample. In addition, the atoms need to be isolated from background gasses, which are often at room temperature. These hot particles will quickly decohere the sample if the partial pressure is too high. Finally, the atoms must be situated in a low-field environment, where the effects of electric (Stark effect) and magnetic (Zeeman effect) fields can be minimized.

It is worth noting, however, that some experiments are making use of room-temperature atomic vapours for interferometric studies [56]. The analogy here is similar to the comparison between traditional optical interferometry and low-coherence interferometry. By employing a variety of techniques, the experimenters are able to extract interferometric information from a low-coherence time sample. This techniques allows the experimenters to achieve very high repetition rates for their experiments (10 kHz, compared to $\sim$1 Hz for traditional AI schemes), and very high dynamic range as well. This comes at a cost of a severely reduced experimental timescale, however, reducing the ultimate precision of measurements that are possible with such a scheme.

Just as there are different types of optical interferometers (Michelson, Mach-

Zehnder, Fizeau, etc.), each with their advantages and disadvantages, so too are there different types of atom interferometers. The two main types of atom interferometer described here are the Raman AI and the grating echo AI.

2.2 The Raman AI

The Raman AI [12,57] uses a sequence of two-photon interactions to manipulate the internal state of the atoms between two metastable states during the experiment, as shown in fig. (2.1). At the end of the interferometer sequence, coherences between the ground and excited states are converted into populations in order to extract the signal. It should be noted that here, "excited state" refers to an upper ground state of the atom's hyperfine ground state manifold, which is coupled to a lower ground state by a two-photon transition. Each interaction is produced by a pair of counter-propagating beams, one with wavevector $\vec{k}_1$ and the other with $\vec{k}_2$. The magnitudes of the wavevectors are chosen to satisfy the two-photon resonance condition: $k_1 - k_2 = E_{2,1}/\hbar c$, where $E_{2,1}$ is the energy difference between the two states. For inertial measurements, the first pulse is tuned to have a pulse area of $\pi/2$, which places the atoms into an equal superposition of the excited and ground states. For the ground state atoms, there is no net change in their momentum after the pulse. Conversely, the excited-state atoms are left with a momentum of $\hbar\delta k$, where $\delta k = k_1 - k_2$. This means that the excited state and ground state will travel along separate spatial trajectories during the experiment. At a time T later, the atoms are subjected to a second pulse. This pulse is tuned to be a π-pulse, which inverts the excited and ground state populations and alters the

momenta of the two atom packets to cause them to re-combine at around time $2T$. At this time, a final $\pi/2$ pulse is applied to generate the interference between the two packets. The phase difference between the upper and lower trajectories, which is proportional to the acceleration experienced by the atoms, is encoded in the excited-state population at this point. This population is measured by applying a resonant pulse and measuring the atomic fluorescence.

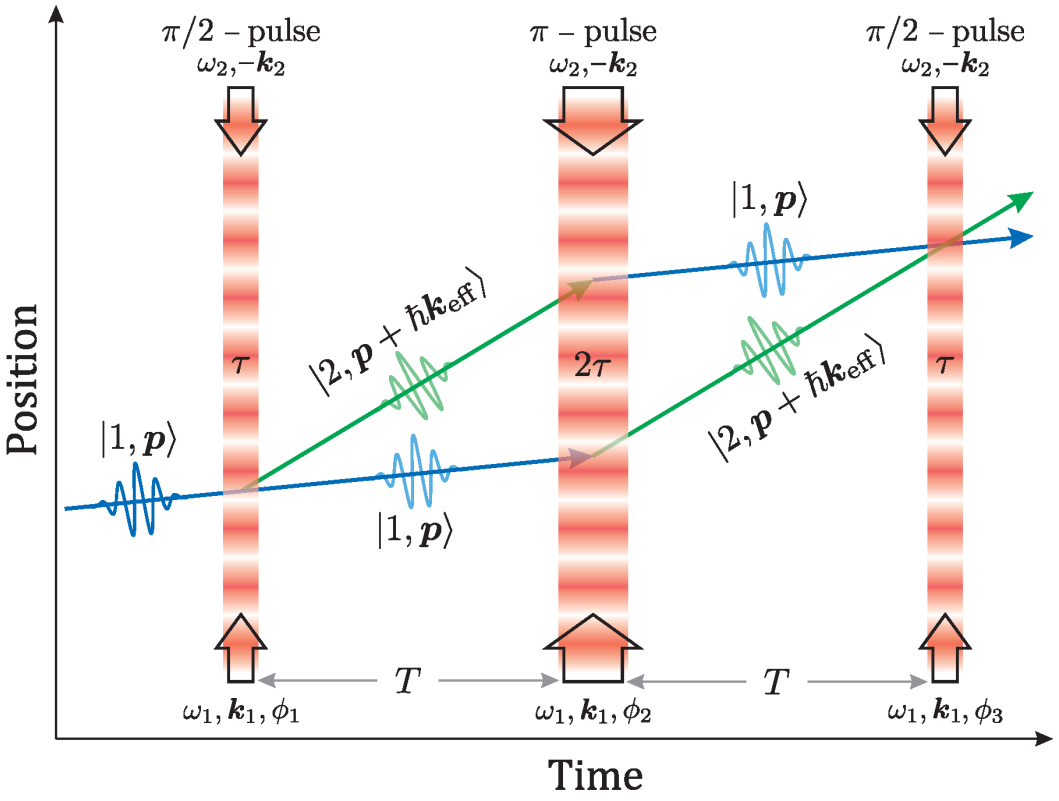

Figure 2.1: Mach-Zehnder Raman AI configured to measure accelerations. The excited state population at the read-out time oscillates sinusoidally as a function of the laser phase which is imprinted onto the atoms every time they undergo a momentum change due to interaction with the beams. This figure reproduced from ref. [57]

While the Raman interferometer is the most commonly-used AI for measuring inertial effects, it has some notable drawbacks. The first is that the atoms must be velocity selected and optically pumped into a single hyperfine magnetic sublevel before the experiment begins. This results in a small fraction ($\sim$ 1 in 25) of the trapped atoms that are available for the AI. This effect is somewhat mitigated by the resonant nature of the final readout allowing each atom to scatter thousands of photons. Another drawback to the Raman approach is the high level of experimental complexity. There are several different optical frequencies needed for even the most basic Raman experiment, all needing a high degree of stability with respect to a stable reference clock.

2.3 The grating echo AI

The grating echo AI is a single state AI that uses far off-resonant SW excitations to produce interfering momentum states, as shown in fig. (2.2). Unlike the Raman AI, the grating echo AI uses a single colour and a (comparatively) simple optical layout. The atoms are returned to their ground state after each interaction, reducing certain systematic effects like the DC Stark shift. Additionally, the echo AI does not require optical pumping (although it is required for eliminating the effect of magnetic field gradients), and it requires no velocity selection. To generate the echo AI signal in its simplest form, first the atoms are separated by a short (Kapitza-Dirac), pulsed SW beam into a manifold of momentum states, each with momentum $\pm 2n\hbar k$, where n is an integer and k is the wavenumber of the light. Immediately after application of the pulse, a $\lambda/2$-periodic density modulation appears in the atomic sample. This

structure washes-out quickly due to the thermal motion of the atoms, and must be re-formed at a later time through the echo technique. The atomic momentum states are allowed to propagate through space until a time T_{21} after the first pulse, where a second pulse is applied. This second pulse re-diffracts the atoms, setting up a rendezvous for certain momentum states whose momenta differ by $2\hbar k$ at a time $2T_{21}$. At this time, called the read-out time, the density of the atom cloud again becomes modulated at the wavelength of the SW. This modulation creates the conditions necessary for Bragg scattering to occur, which allows a readout of the density grating phase and amplitude using a TW (traveling wave) pulse. By examining the amplitude and phase of the reflected light from the grating, the acceleration of the atoms along the experimental axis can be determined. This leads to a measurement of the gravitational acceleration acting on the atoms.

Due to rubidium's nature as an alkali element, it can be treated as a one-electron atom without too much loss of information about the atomic state. We will adopt this assumption in order to derive the signal of the so-called two-pulse grating echo.

2.4 The one-pulse echo signal

What follows is a calculation of the expected signal from the echo interferometer, based substantially on previous developments of similar expected signals [24, 58–65]. This derivation preserves a fair bit of generality, and can be applied to atom interferometers at either "cold" ($\sim 10\ \mu$K, as seen in typical MOTs) or "ultra-cold" (~ 10 nK, as seen in BECs) temperatures. The effect of spontaneous emission is

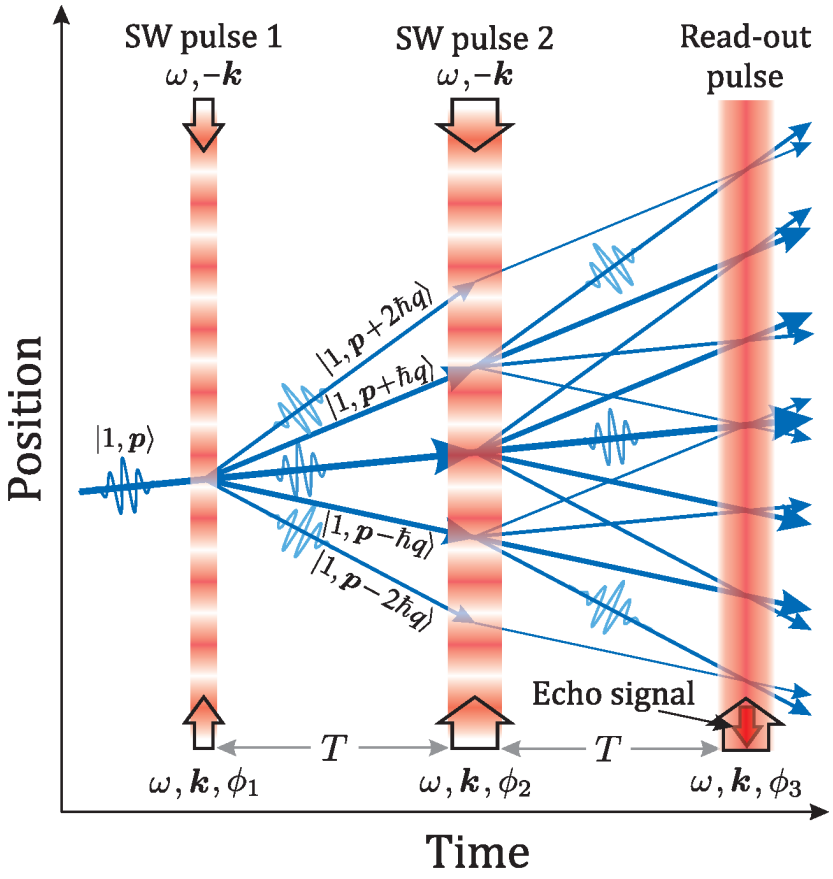

Figure 2.2: Two-pulse echo AI configured to measure accelerations. The signal (proportional to the acceleration of the system) manifests itself in the phase of the density modulation produced at time 2T, which is detected in the phase of the back-scattered read-out light at this time. This figure reproduced from ref. [57]

included throughout, but effects due to the magnetic sublevels of the atom and due to the spatial profile of the AI beams are omitted herein.

We model our system as a two level atom, governed by the time-dependent Schrödinger equation:

$$\hat{H}\begin{pmatrix} a_e \\ a_g \end{pmatrix} = i\hbar \begin{pmatrix} \dot{a}_e \\ \dot{a}_g \end{pmatrix}. \tag{2.1}$$

Adopting the dipole and rotating-wave approximations [66,67], the Hamiltonian in eq. (2.1) is given by:

$$\hat{H} = \hbar \begin{pmatrix} -\Delta - i\gamma & \Omega(\mathbf{r}) \\ \Omega(\mathbf{r}) & 0 \end{pmatrix}, \tag{2.2}$$

where $\Delta = \omega - \omega_0$ is the detuning of the TW components of the SW light from the atomic resonance, $\gamma = -\Gamma/2$ is the coherence decay rate of the atom and $\Omega(\vec{r}) = \Omega_0 \cos(\mathbf{k} \cdot \mathbf{r})$ is the position-dependent Rabi frequency of the SW potential. $\Omega_0 = \mu_{eg} E_0/\hbar$ is the Rabi frequency of the interaction, with μ_{eg} being the dipole matrix element linking the excited and ground states and E_0 being the electric field magnitude of each traveling-wave component of the SW.

If we substitute eq. (2.2) into eq. (2.1), we obtain a pair of coupled differential equations for the ground and excited state amplitudes. By assuming that $|\dot{a}_e| \ll |\Delta|$, which is true for far-off-resonant SW beams, we obtain:

$$a_e = \frac{\Omega(\mathbf{r})a_g}{(\Delta + i\gamma)}, \tag{2.3a}$$

$$\dot{a}_g = \frac{-i\Omega^2(\mathbf{r})a_g}{(\Delta + i\gamma)}. \tag{2.3b}$$

It is helpful to make the substitution $\theta = \arctan(-\Gamma/2\Delta)$ into eq. (2.3). This yields:

$$a_e = \frac{\Omega(\mathbf{r})e^{i\theta}a_g}{\sqrt{\Delta^2 + \gamma^2}}, \tag{2.4a}$$

$$\dot{a}_g = \frac{-i\Omega^2(\mathbf{r})e^{i\theta}a_g}{\sqrt{\Delta^2 + \gamma^2}}. \tag{2.4b}$$

In eq. (2.4), θ acts as a phase that characterizes the effect of spontaneous emission during the pulse. To obtain an expression for the ground state amplitude, we integrate eq. (2.4b):

$$a_g^{(1)}(\mathbf{r}) = a_g^{(0)}(\mathbf{r})e^{-i\Theta_1 \cos \mathbf{q}\cdot\mathbf{r}}, \tag{2.5}$$

where the superscript on the state amplitude denotes the number of Kapitza-Dirac SW pulses it has been subjected to, $\mathbf{q} = \mathbf{k}_1 - \mathbf{k}_2 = 2\mathbf{k}$ (for perfectly-counterpropagating TW beams) and $\Theta_1 = u_1 e^{i\theta}$ is the complex pulse area of the first applied pulse, whose magnitude is given by:

$$u_1 = \frac{\Omega_0^2 \tau_1}{2|\Delta|}\left[1 + \left(\frac{\Gamma}{2\Delta}\right)^2\right]^{-1/2}, \tag{2.6}$$

where τ_1 is the length of the first SW pulse.

It is important to point out that the Hamiltonian, eq. (2.1), used here is non-Hermitian. As such, it does not necessarily preserve normalization of the wavefunction at all times. To deal with this, a constant term in the exponent of eq. (2.5) has been omitted, which removes an unphysical decay of the ground state amplitude during the pulse. In order to treat this interferometer signal "properly" with Hermitian operators, we would need to use the density matrix formalism, which is beyond the scope of this work, or ignore the effect of spontaneous emission [59].

To simplify eq. (2.5), we apply the Jacobi-Anger expansion:

$$e^{-iz\cos\alpha} = \sum_n (-i)^n J_n(z) e^{in\alpha}, \tag{2.7}$$

which yields:

$$a_g^{(1)}(\mathbf{r}) = a_g^{(0)}(\mathbf{r}) \sum_n (-i)^n J_n(\Theta_1) e^{in\mathbf{q}\cdot\mathbf{r}}, \tag{2.8}$$

where $J_n(z)$ is a Bessel function of the first kind.

The ground state amplitude is then made-up of a sum of momentum states, each having a momentum value that is an even multiple of $\hbar k$. These states correspond to an integer number of interactions with the SW field, wherein the atom absorbs a photon from one traveling-wave field component direction and emits into the other.

The excited state amplitude can be written by substituting eq. (2.8) into eq. (2.4a):

$$a_e^{(1)}(\mathbf{r}) = \frac{\Omega_0 e^{i\theta}}{2(\Delta^2+\gamma^2)^{1/2}} a_g^{(0)}(\mathbf{r}) \sum_n (-i)^n J_n(\Theta_1) \left[e^{i(2n-1)\mathbf{k}\cdot\mathbf{r}} + e^{i(2n+1)\mathbf{k}\cdot\mathbf{r}} \right]. \quad (2.9)$$

Since our SW excitation light is far-detuned from atomic resonance ($\Omega_0 \ll |\Delta|$), we can ignore the excited state amplitude. Additionally, we assume that the length of the SW pulse is short enough that we can neglect the motion of the atoms during the pulse. This is the Raman-Nath approximation.

After the first pulse ends, the atomic state evolves freely. To model this evolution most simply, we transform the ground state amplitude into momentum space:

$$a_g^{(1)}(\mathbf{p}) = \frac{1}{(2\pi\hbar)^{3/2}} \int a_g^{(1)}(\mathbf{r}) e^{-i\mathbf{p}\cdot\mathbf{r}/\hbar} d^3r. \quad (2.10)$$

We assume that the initial wavefunction is a plane wave:

$$a_g^{(0)}(\mathbf{r}) = \frac{1}{V^{1/2}} e^{i\mathbf{p}_0\cdot\mathbf{r}/\hbar}, \quad (2.11)$$

where V is the interaction volume and $\mathbf{p}_0$ is the initial momentum of the atom. Therefore, the momentum-space ground-state amplitude after the pulse is then:

$$a_g^{(1)}(\mathbf{p}) = \frac{(2\pi\hbar)^{3/2}}{V^{1/2}} \sum_n (-i)^n J_n(\Theta_1) \delta^3(\mathbf{p} - \mathbf{p}_0 - n\hbar\mathbf{q}). \quad (2.12)$$

The Dirac delta function in eq. (2.12) ensures that the momentum of the atom ground state wavefunction is quantized in units of $\hbar q$.

The Hamiltonian for the case of free-evolution has only a kinetic energy term, so the momentum-space solution of the Schrödinger equation is:

$$a_g^{(0)}(\mathbf{p},t) = a_g^{(1)}(\mathbf{p},0)e^{\left[\frac{-ip^2}{\hbar 2M}t\right]}. \tag{2.13}$$

Now, we apply eq. (2.13) to eq. (2.12) and transform back into position-space:

$$a_g^{(1)}(\mathbf{r},t) = \frac{1}{(2\pi\hbar)^{3/2}}\int a_g^{(1)}(\mathbf{p},t)e^{i\mathbf{p}\cdot\mathbf{r}/\hbar}d^3p \tag{2.14a}$$

$$= \frac{1}{V^{1/2}}\sum_n (-i)^n J_n(\Theta_1)e^{-i(\mathbf{p}_0+n\hbar\mathbf{q})^2 t/2M\hbar}e^{i(\mathbf{p}_0+n\hbar\mathbf{q})\cdot\mathbf{r}/\hbar}. \tag{2.14b}$$

Making the substitution:

$$\frac{(\mathbf{p}_0+n\hbar\mathbf{q})^2}{2M} = \epsilon_0 + n\hbar\mathbf{q}\cdot\mathbf{v}_0 + n^2\hbar\omega_q, \tag{2.15}$$

where ϵ_0 is the kinetic energy and $\omega_q = \hbar q^2/2M$ is the recoil frequency of the atom, we see that the ground state amplitude after the first pulse is given by:

$$a_g^{(1)}(\mathbf{r},t) = \frac{e^{i(\mathbf{p}_0\cdot\mathbf{r}-\epsilon_0 t)/\hbar}}{V^{1/2}}\sum_n (-i)^n J_n(\Theta_1)e^{in\mathbf{q}\cdot\mathbf{r}}e^{-in\mathbf{q}\cdot\mathbf{v}_0 t}e^{-in^2\omega_q t}. \tag{2.16}$$

At this point, we pause in the development of the full echo AI signal derivation (which requires at least 2 SW pulses in its simplest form) to examine the case of only a single applied SW pulse. Since we extract our information about the atoms by probing the sample with a traveling-wave readout pulse, the quantity of most interest is the atomic density distribution, $\rho = a_g^* a_g$:

$$\rho_g^{(1)}(\mathbf{r},t) = \frac{1}{V}\sum_{n,n'}(i)^{n'-n}J_n(\Theta_1)J_{n'}(\Theta_1^*)e^{-i(n'-n)\mathbf{q}\cdot\mathbf{r}}e^{i(n'-n)\mathbf{q}\cdot\mathbf{v}_0 t}e^{i(n'^2-n^2)\omega_q t} \tag{2.17a}$$

$$= \frac{1}{V}\sum_{\eta,n}(i)^{\eta}J_n(\Theta_1)J_{n+\eta}(\Theta_1^*)e^{-i\eta\mathbf{q}\cdot\mathbf{r}}e^{i\eta\mathbf{q}\cdot\mathbf{v}_0 t}e^{i\eta(2n+\eta)\omega_q t}, \tag{2.17b}$$

where $\eta = n' - n$.

The value of η is called the order of the interference, and it denotes the difference in momentum between interfering states at the time of the read out. The order of the interference also determines the wavelength of the spatially-periodic density modulation in the atomic cloud at the time of the read out pulse. This is important, because our detection scheme relies on Bragg scattering, and we use the same wavelength of light for the SW and RO beams. The Bragg condition (at normal incidence): $2d = m\lambda$, coupled with spacing of the density maxima: $d = 2\pi/\eta|\mathbf{q}|$ yields:

$$m\lambda = 2\left(\frac{2\pi}{\eta|\mathbf{q}|}\right) \implies m\eta = 1. \tag{2.18}$$

Since both m and η are constrained to be positive integers, the only values that satisfy the result of eq. (2.18) is $m = \eta = 1$. This simplifies the sum in eq. (2.17):

$$\rho_g^{(1)}(z,t) = \frac{1}{V}\sum_{n} iJ_n(\Theta_1)J_{n+1}(\Theta_1^*)e^{-iqz}e^{iqv_0 t}e^{i(2n+1)\omega_q t}, \tag{2.19}$$

where we have taken $\mathbf{q}\cdot\mathbf{r} = qz$, which is the same as considering the case where the TW components of the SW potential are aligned with the z-axis, and perfectly counter-propagating each-other.

We see that eq. (2.19) contains a term, e^{iqv_0t}. This term contains the information about Doppler dephasing of the interference due to the different atomic velocity classes present at the beginning of the experiment. Each of these velocity classes will contribute a different frequency to the total wavefunction of the atoms, resulting in dephasing of the density modulation. To determine the timescale of this dephasing, we must integrate the density over the speed distribution of the atoms. This is simplified by realizing that the only term in eq. (2.19) that depends on v_0 is the afore-mentioned exponential term. The velocity distribution of the atoms is:

$$f(v) = \frac{1}{(\pi\sigma_v^2)^{3/2}} e^{-v^2/\sigma_v^2}, \tag{2.20}$$

where $\sigma_v = \sqrt{2k_BT/M}$ is the width of the velocity distribution. The integral over the initial-velocity-dependent term in the atomic density is:

$$\int f(v_0) e^{iqv_0t} d^3v_0 = e^{-(q\sigma_v t/2)^2}. \tag{2.21}$$

Therefore, after replacing the e^{iqv_0t} term in eq. (2.19) with the L.H.S. of eq. (2.21), the atomic density becomes:

$$\rho_g^{(1)}(z,t) = \frac{1}{V} \sum_n i J_n(\Theta_1) J_{n+1}(\Theta_1^*) e^{-iqz} e^{i(2n+1)\omega_q t} e^{-(q\sigma_v t/2)^2}. \tag{2.22}$$

As shown, the amplitude of the density modulation will decay on a timescale of $\tau_{coh} = 2/q\sigma_v$, which for Rb atoms at 1 μK is ~ 9 μs.

The signal from the atoms after a single pulse is generated by applying a traveling-wave read out pulse to the atoms. The back-reflected electric field amplitude from the density modulation of the atoms will be proportional to the amplitude

of the Fourier-component, e^{iqz} in eq. (2.22):

$$E^{(1)}(t) \propto E_{RO} e^{-q^2\sigma_v^2 t^2/4} \sum_n i J_n(\Theta_1) J_{n+1}(\Theta_1^*) e^{i(2n+1)\omega_q t}, \tag{2.23}$$

where E_{RO} is the amplitude of the applied read out pulse.

This expression can be simplified by applying a Bessel function summation identity:

$$\sum_n J_n(\Theta) J_{n+\eta}(\Theta^*) e^{i2n\phi} = i^\eta e^{-i\eta\phi} J_\eta(\omega) \left(\frac{\sin(\phi-\theta)}{\sin(\phi+\theta)} \right)^{\eta/2}, \tag{2.24}$$

where $\Theta = ue^{i\theta}$, $\omega = 2u(\sin(\phi+\theta)\sin(\phi-\theta))^{1/2}$ and $\phi = \eta\omega_q t$. The read out signal becomes:

$$E^{(1)}(t) \propto -E_{RO} e^{-q^2\sigma_v^2 t^2/4} J_1(\omega_1) \left[\frac{\sin(\omega_q t-\theta)}{\sin(\omega_q t+\theta)} \right]^{1/2}, \tag{2.25}$$

where $\omega_1 = 2u_1\sqrt{\sin[\omega_q t+\theta]\sin[\omega_q t-\theta]}$.

The one-pulse signal is then the squared field amplitude:

$$S^{(1)}(t) \propto E_{RO}^2 e^{-q^2\sigma_v^2 t^2/2} [J_1(\omega_1)]^2 \frac{\sin(\omega_q t-\theta)}{\sin(\omega_q t+\theta)}. \tag{2.26}$$

A plot of eq. (2.26) with temperature set to zero ($\sigma_v = 0$) is shown in fig. (2.3a). Note the zeroes in the signal in the vicinity of the recoil period, τ_q, which would allow measurement of the recoil frequency if $T = 0$ were achievable in reality. If we allow the temperature to increase to a realistic value, we get the signal shown in fig. (2.3b). Note the different horizontal scales. In this case, it is clear that the signal amplitude decreases on a timescale shorter than the recoil period, making

recoil measurements or long time-scale experiments impossible. As the area of the pulse increases, the one-pulse signal begins to exhibit high-frequency oscillatory behaviour.

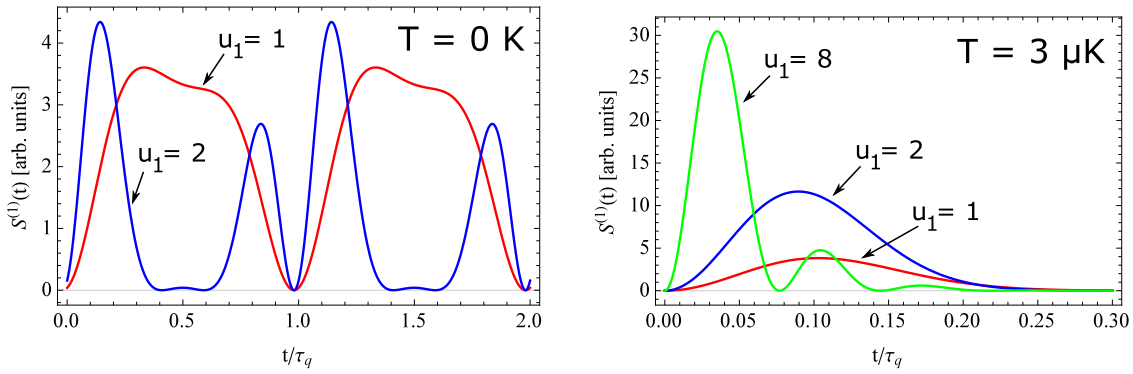

Figure 2.3: Plots of eq. (2.26) for different temperatures and pulse areas. For all curves, $\theta = -0.062$ Radians, a typical value for our experiments. Note the onset of oscillatory behaviour for non-zero temperature as the pulse area is increased. This allows for some limited experimental probing of the density modulation formation on short timescales.

2.5 Two-pulse signal

Due to the rapid wash-out of the one-pulse signal due to Doppler-dephasing of the atomic density grating, it is unsuitable for high-precision measurements of accelerations or atomic recoil. In the case of recoil measurements, this is because the recoil period, ~ 30 μs is much longer than the coherence time of the signal $\tau_{coh} \sim 9$ μs. Since precise measurement of a frequency requires that multiple periods are observed, such a short experimental window is unsuitable for recoil. In the case of gravity, the phase shift of the fringes accumulates as $\propto gt^2$, so longer timescales

mean larger phase shifts, which are more easily and accurately measurable.

For these reasons, we use an echo technique, conceptually similar to a photon echo, to recover the density modulation of the atom cloud at a later time. This technique allows re-formation of the periodic density structure in the atomic cloud at any later time, theoretically only limited by the transit time of the atoms (i.e.: the time for the atoms to leave the interaction volume, based on their initial velocities). Since the interaction volume for our experiments is on the scale of 1 cm^3, the transit time for 1 μK atoms is on the scale of 1 s. In practice, it is difficult to approach this experimental timescale due to the influence of magnetic gradients, Doppler shifts due to gravitational acceleration and collisional effects, so we find that our ultimate timescale is limited to less than $\sim$ 100 ms. Nonetheless, this represents a significant improvement over the $\sim 0.1\tau_q \simeq 3$ μs maximum timescale of the one-pulse experiments, as shown in fig. (2.3).

To derive the two-pulse signal, we allow the atoms to evolve for a time T_{21} after the first pulse. The ground state amplitude at the time of the second pulse is then given by eq. (2.16) evaluated at $t = T_{21}$:

$$a_g^{(1)}(\mathbf{r}, T_{21}) = \frac{e^{i(\mathbf{p}_0\cdot\mathbf{r}-\epsilon_0 T_{21})/\hbar}}{V^{1/2}} \sum_n (-i)^n J_n(\Theta_1) e^{in\mathbf{q}\cdot\mathbf{r}} e^{-in\mathbf{q}\cdot\mathbf{v}_0 T_{21}} e^{-in^2\omega_q T_{21}}. \qquad (2.27)$$

As shown in eq. (2.5), the ground state amplitude after a SW pulse is simply equal to the amplitude at the time of the pulse onset, multiplied by the factor $exp(-i\Theta_n \cos\mathbf{q}\cdot\mathbf{r})$, where n is the number of the applied pulse. We are free to take this approach because there is no explicit time-dependence in the state amplitude

of eq. (2.27). If we apply the second pulse in this way, we obtain:

$$a_g^{(2)}(\mathbf{r}, T_{21}) = \frac{e^{i(\mathbf{p}_0\cdot\mathbf{r}-\epsilon_0 T_{21})/\hbar}}{V^{1/2}} \sum_n (-i)^n J_n(\Theta_1) e^{in\mathbf{q}\cdot\mathbf{r}} e^{-in\mathbf{q}\cdot\mathbf{v}_0 T_{21}} e^{-in^2\omega_q T_{21}} e^{-i\Theta_2\cos(\mathbf{q}\cdot\mathbf{r})}. \tag{2.28}$$

To evolve the amplitude in the time following the second pulse, we transform to momentum space:

$$\begin{aligned} a_g^{(2)}(\mathbf{p}, T_{21}) = \frac{(2\pi\hbar)^{3/2}}{V^{1/2}} e^{-i\epsilon_0 T_{21}/\hbar} \sum_{n,m} (-i)^{n+m} J_n(\Theta_1) J_m(\Theta_2) \\ \times e^{-in\mathbf{q}\cdot\mathbf{v}_0 T_{21}} e^{-in^2\omega_q T_{21}} \delta^3(\mathbf{p} - \mathbf{p}_0 - (n+m)\hbar\mathbf{q}). \end{aligned} \tag{2.29}$$

This momentum-space wavefunction contains a Dirac delta function, $\delta^3(\mathbf{p} - \mathbf{p}_0 - (n+m)\hbar\mathbf{q})$, which ensures that the momenta of the wavepackets that make up the function are discretized in units of $\hbar\mathbf{q}$. Remaining in momentum space, we allow the wavefunction to evolve until time t, before transforming back to position space:

$$\begin{aligned} a_g^{(2)}(\mathbf{r}, T_{21}+t) = \frac{e^{i(\mathbf{p}_0\cdot\mathbf{r}-\epsilon_0(\tau_1+T_{21}+\tau_2+t))/\hbar}}{V^{1/2}} \sum_{n,m} (-i)^{n+m} J_n(\Theta_1) J_m(\Theta_2) \\ \times e^{i(n+m)\mathbf{q}\cdot\mathbf{r}} e^{-i\mathbf{q}\cdot\mathbf{v}_0[nT_{21}+(n+m)t]} e^{-i\omega_q[n^2T_{21}+(n+m)^2t]}. \end{aligned} \tag{2.30}$$

The density is found by calculating $a_g^{(2)*} a_g^{(2)}$:

$$\rho_g^{(2)}(\mathbf{r}, T_{21}+t) = \frac{1}{V}\sum_{n,m,n',m'}(-i)^{n+m-n'-m'}J_n(\Theta_1)J_m(\Theta_2)J_{n'}(\Theta_1^*)J_{m'}(\Theta_2^*)$$
$$\times e^{i(n+m-n'-m')\mathbf{q}\cdot\mathbf{r}}e^{-i\mathbf{q}\cdot\mathbf{v}_0[(n-n')T_{21}+(n+m-n'-m')t]}$$
$$\times e^{i\omega_q\{(n^2-n'^2)T_{21}+[(n+m)^2-(n'+m')^2]t\}}. \tag{2.31}$$

As for the one-pulse case, there are some simplifications to be made at this point. In particular, we make the substitutions: $\bar{\eta} = n' + m' - n - m$ and $\bar{N} = (n - n')/\eta$. The physical meaning of $\bar{\eta}$ is the difference in momentum between the interfering states at time $\tau_1 + T_{21} + \tau_2 + t$, in units of $\hbar\mathbf{q}$. As for $\bar{N}$, it represents the ratio of the difference between the interfering states in the first pulse to the second pulse. With this, the density becomes:

$$\rho_g^{(2)}(\mathbf{r}, T_{21}+t) = \frac{1}{V}\sum_{n,m,\bar{\eta},\bar{N}} i^{\bar{\eta}}J_n(\Theta_1)J_{n-\bar{\eta}\bar{N}}(\Theta_1^*)J_m(\Theta_2)J_{m+\bar{\eta}(\bar{N}+1)}(\Theta_2^*)$$
$$\times e^{-i\bar{\eta}\mathbf{q}\cdot\mathbf{r}}e^{i\mathbf{q}\cdot\mathbf{v}_0(\bar{\eta}t-\bar{\eta}\bar{N}T_{21})}$$
$$\times e^{i\omega_q[-\bar{\eta}\bar{N}(2n-\bar{\eta}\bar{N})T_{21}+\bar{\eta}[2(n+m)+\bar{\eta}]t]}, \tag{2.32}$$

which allows us to separate the sums over $\bar{\eta}$, $\bar{N}$, n and m:

$$\rho_g^{(2)}(\mathbf{r}, T_{21}+t) = \frac{1}{V}\sum_{\bar{\eta},\bar{N}} i^{\bar{\eta}} e^{-i\bar{\eta}\mathbf{q}\cdot\mathbf{r}} e^{i\mathbf{q}\cdot\mathbf{v}_0(\bar{\eta}t-\bar{\eta}\bar{N}T_{21})} e^{i\omega_q[(\bar{\eta}\bar{N})^2 T_{21}+\bar{\eta}^2 t]}$$
$$\times \sum_n J_n(\Theta_1) J_{n-\bar{\eta}}(\Theta_1^*) e^{i2n\omega_q[\bar{\eta}t-\bar{\eta}\bar{N}T_{21}]}$$
$$\times \sum_m J_m(\Theta_2) J_{m+\bar{\eta}(\bar{N}+1)}(\Theta_2^*) e^{i2m\omega_q(\bar{\eta}t)}. \tag{2.33}$$

We are now free to apply the Bessel function identity, eq. (2.24), to collapse some of the sums in eq. (2.33):

$$\rho_g^{(2)}(\mathbf{r}, T_{21}+t) = \frac{1}{V}\sum_{\bar{\eta},\bar{N}} i^{\bar{\eta}} e^{-i\bar{\eta}\mathbf{q}\cdot\mathbf{r}} e^{i\phi_D(\mathbf{v}_0)} e^{i\omega_q[(\bar{\eta}\bar{N})^2 T_{21}+\bar{\eta}^2 t]}$$
$$\times i^{-\bar{\eta}\bar{N}} e^{i\bar{\eta}\bar{N}\phi_1} J_{-\bar{\eta}\bar{N}}(\omega_1) \left(\frac{\sin(\phi_1-\theta)}{\sin(\phi_1+\theta)}\right)^{\bar{\eta}\bar{N}}$$
$$\times i^{\bar{\eta}(\bar{N}+1)} e^{-i\bar{\eta}(\bar{N}+1)\phi_2} J_{\bar{\eta}(\bar{N}+1)}(\omega_2) \left(\frac{\sin(\phi_2-\theta)}{\sin(\phi_2+\theta)}\right)^{\bar{\eta}(\bar{N}+1)/2}, \tag{2.34}$$

where $\phi_D(\mathbf{v}_0)$ is the initial velocity-dependent Doppler phase, ϕ_1 and ϕ_2 are the recoil phases induced by the first and second pulses, respectively, and ω_1 and ω_2 are the amplitudes of the interferences caused by the first and second pulses, re-

spectively. Their values are:

$$\phi_D(\mathbf{v}_0) = \bar{\eta}\mathbf{q}\cdot\mathbf{v}_0(t-\bar{N}T_{21}) \tag{2.35a}$$

$$\phi_1 = \bar{\eta}\omega_q(t-\bar{N}T_{21}) \tag{2.35b}$$

$$\phi_2 = \bar{\eta}\omega_q t \tag{2.35c}$$

$$\omega_j = 2u_j\sqrt{\sin(\phi_j+\theta)\sin(\phi_j-\theta)}, j\in[1,2]. \tag{2.35d}$$

Many of the exponential terms in eq. (2.34) cancel, leaving:

$$\rho_g^{(2)}(\mathbf{r}, T_{21}+t) = \frac{1}{V}\sum_{\bar{\eta},\bar{N}}(-1)^{\bar{\eta}}e^{-i\bar{\eta}\mathbf{q}\cdot\mathbf{r}}e^{i\phi_D(\mathbf{v}_0)}J_{-\bar{\eta}\bar{N}}(\omega_1)J_{\bar{\eta}(\bar{N}+1)}(\omega_2) \times\left(\frac{\sin(\phi_1-\theta)}{\sin(\phi_1+\theta)}\right)^{-\bar{\eta}\bar{N}/2}\left(\frac{\sin(\phi_2-\theta)}{\sin(\phi_2+\theta)}\right)^{\bar{\eta}(\bar{N}+1)/2}. \tag{2.36}$$

Note that this expression for the ground state density carries its recoil modulation in the terms containing $\phi_{1,2}$ and $\omega_{1,2}$. We now integrate over the initial velocity distribution of the atoms, similar to the case for the one-pulse signal. The atomic density becomes:

$$\rho_{g,d}^{(2)}(\mathbf{r}, T_{21}+t) = \frac{1}{V}\sum_{\bar{\eta},\bar{N}}(-1)^{\bar{\eta}}e^{-i\bar{\eta}\mathbf{q}\cdot\mathbf{r}}e^{-[\bar{\eta}(t-\bar{N}T_{21})/\tau_{coh}]^2}J_{-\bar{\eta}\bar{N}}(\omega_1)J_{\bar{\eta}(\bar{N}+1)}(\omega_2) \times\left(\frac{\sin(\phi_1-\theta)}{\sin(\phi_1+\theta)}\right)^{-\bar{\eta}\bar{N}/2}\left(\frac{\sin(\phi_2-\theta)}{\sin(\phi_2+\theta)}\right)^{\bar{\eta}(\bar{N}+1)/2}. \tag{2.37}$$

The amplitude of the density modulation of the atoms will be negligible at all times except near certain "echo" times, for which $t_{echo} = (\bar{N}+1)T_{21}$. In the vicinity of t_{echo}, the amplitude of the density modulation will be detectable for a time $\sim \tau_{coh}/\bar{\eta}$ before velocity dephasing washes-out the contrast of the grating. The coherence time, $\tau_{coh} = 2/q\sigma_v$ is the length of time needed for a representative atom in the thermal velocity distribution of the MOT to travel $\lambda/2$. Its value is around 3 μs for typical MOT temperatures.

To extract the signal from the two-pulse AI, we apply the same traveling-wave RO pulse. The backscattered signal is proportional to the amplitude of the $e^{-i\mathbf{q}\cdot\mathbf{r}}$ term in the atomic density, due to the need to satisfy the Bragg condition. This means that only the terms for which $\bar{\eta} = \pm 1$ will contribute to the observed signal. Since the density function is even with respect to $\bar{\eta}$, we drop every term except for $\bar{\eta} = 1$. The backscattered amplitude for the echo of order $\bar{N}$ is:

$$\begin{aligned} E_g^{(2)}(\Delta t) \propto\ & E_{RO}(-1)^{\bar{N}+1} e^{-(\Delta t/\tau_{coh})^2} \\ & \times \left(\frac{\sin(\omega_q \Delta t + \theta)}{\sin(\omega_q \Delta t - \theta)}\right)^{\bar{N}/2} \left(\frac{\sin[\omega_q(\Delta t + \bar{N}T_{21}) - \theta]}{\sin[\omega_q(\Delta t + \bar{N}T_{21}) + \theta]}\right)^{(\bar{N}+1)/2} \\ & \times J_{\bar{N}}\left(2u_1\sqrt{\sin(\omega_q \Delta t + \theta)\sin(\omega_q \Delta t - \theta)}\right) \\ & \times J_{\bar{N}+1}\left(2u_2\sqrt{\sin[\omega_q(\Delta t + \bar{N}T_{21}) + \theta]\sin[\omega_q(\Delta t + \bar{N}T_{21}) - \theta]}\right), \end{aligned} \quad (2.38)$$

where $\Delta t = t - t_{echo}$ is the time relative to the $\bar{N}^{\text{th}}$ echo time.

Figure (2.4) shows the effect of varying the parameters in eq. (2.38). In part a), we see that by increasing the pulse area of either the first or second SW pulses,

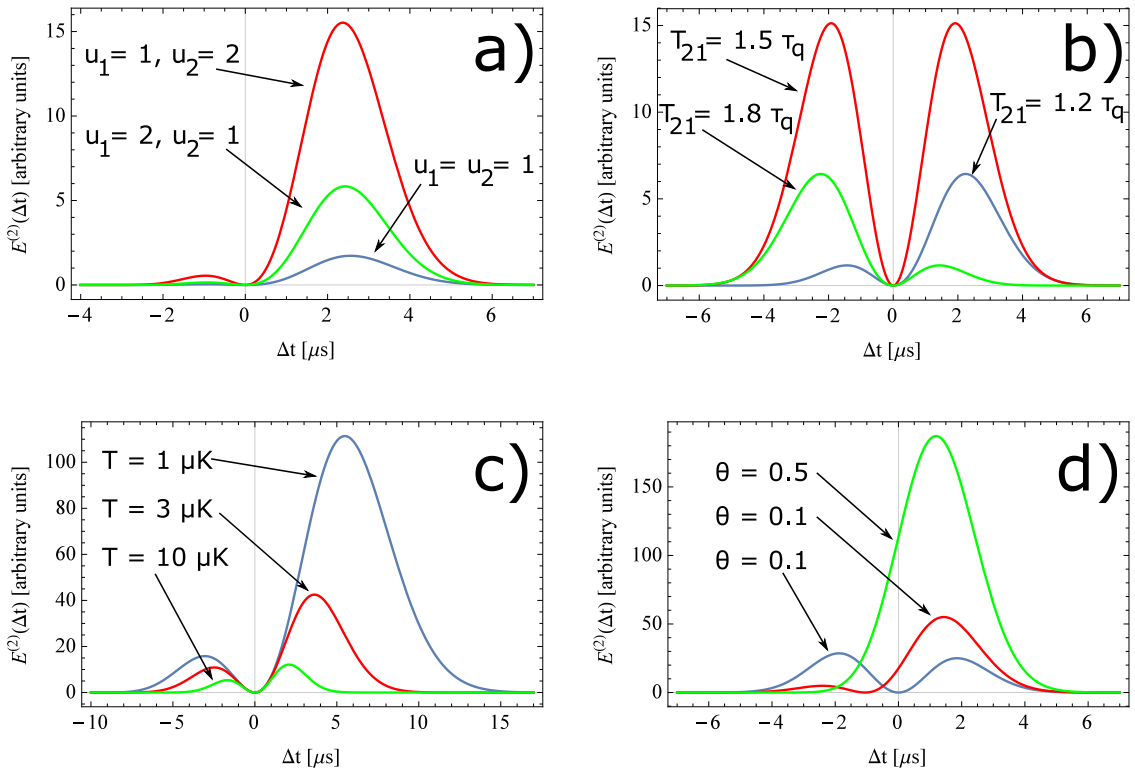

Figure 2.4: Plots of the effect of varying the parameters in eq. (2.38) on the two-pulse signal shape. Part a) shows the effect of changing the relative intensities of the two pulses. Here the values of $T = 10\ \mu K$, $T_{21} = 1.1\tau_q$ and $\theta = 0$ are fixed for the three curves. In part b) the effect of varying T_{21} is shown. The fixed parameters are $u_1 = u_2 = 1$, $T = 10\ \mu K$ and $\theta = 0$ for all curves. For part c), the effect of varying trap temperature is demonstrated. All curves have $u_1 = u_2 = 1$, $T_{21} = 1.3\tau_q$ and $\theta = 0$. Finally, in part d), the influence of spontaneous emission is shown.

the resulting signal amplitude increases. The effect of increasing the second pulse produces a much larger effect on the signal than the same increase applied only to the first pulse. Mathematically, this can be explained by the realization that $J_2(x) > J_1(x)$ for $\sim 2.6 < x <\sim 6.1$, which is exactly the area of parameter space that we find ourselves in. Part b) shows the effect of changing the the intra-pulse time, T_{21}. The observed behaviour is periodic in ω_q, with the two lobes of the back-scattered signal being equal in amplitude at the point of maximum grating contrast. For part c), the effect shown is that of varying temperature for fixed AI parameters. Smaller temperatures produce larger, longer-lived signals, which is a direct consequence of the effect shown in fig. (2.3). Finally, the effect of varying the parameter θ that controls spontaneous emission is shown in d). Changing θ changes the distribution of reflected intensity about the zero point, and also moves the zero point, which is normally found at $\Delta t = 0$ for $\theta = 0$, towards lower time. This effect has been studied in detail elsewhere [62].

To explain why we use the first echo ($\bar{N} = 1$), which is modulated at ω_q, rather than a higher-order echo, modulated at $\bar{N}\omega_q$, we have to consider the ratio of two sequential Bessel function of the first kind:

$$\frac{J_{n-1}(z)}{J_n(z)} = \frac{2n}{z}. \tag{2.39}$$

For a small argument z, the ratio in eq. (2.39) becomes very large, which means that in the back-scattered signal amplitude expression from eq. (2.38), the $J_{\bar{N}}$ term will dominate the amplitude of the signal. With this in mind, if we take the ratio of the amplitudes of the echo signals from the $\bar{N}^{\text{th}}$ and $(\bar{N}-1)^{\text{th}}$ echoes, we can

see that:

$$\frac{A_{\bar{N}-1}}{A_{\bar{N}}} \sim \frac{J_{\bar{N}-1}(2u_1)}{J_{\bar{N}}(2u_1)} = \frac{\bar{N}}{u_1}. \tag{2.40}$$

Therefore, the amplitude of the $\bar{N}^{\text{th}}$ echo is decreased by a factor of $\sim \bar{N}/u_1$ compared to the $(\bar{N}-1)^{\text{th}}$ echo. This provides strong incentive to use only the first-order echo for experiments, due to the unavoidable loss of signal-to-noise ratio (SNR) for higher-order echoes. For typical experimental parameters of $\tau_1 = 1$ μs, $\Delta = 1$ GHz and beam intensity $I = 1000$ Wm^{-2}, $u_1 \simeq 50$. Therefore, the 2^{nd} echo will be ~ 25 times smaller than the 1^{st}, while the 3^{rd} echo will be ~ 400 times smaller than the 1^{st}. The work in ref. [24] shows the first- and second-order echoes for a particular implementation of a similar experiment. However, the displayed data were produced using different polarizations and pulse intensities, meaning that the 25-fold reduction in SNR is not obvious. Nonetheless, for fixed experimental parameters, the SNR of the echo signal will very quickly decay for higher-order echoes. For this reason, this work focuses entirely on the first-order echo signal. So although the higher-order echoes are modulated at integer multiples of ω_q, which can be useful in principle, SNR decreases so quickly as $\bar{N}$ increases that measurements of recoil modulation for $\bar{N} \neq 1$ are very difficult in practice. The work in ref. [64] explores a variety of different echo types which will not be considered in any detail here.

2.6 The effect of gravity

When a constant force acts on the atoms during the AI pulse sequence, the resulting read-out signal is modified as follows [59]:

$$E_{grav}^{(2)} = E_g^{(2)} e^{i\phi_g^{(2)}}, \tag{2.41}$$

where the two-pulse gravitational phase is given by:

$$\phi_g^{(2)} = -\frac{qg}{2}[\bar{N}(\bar{N}+1)T_{21}^2 + 2(\bar{N}+1)T_{21}\Delta t + \Delta t^2]. \tag{2.42}$$

In the case of a first-order echo, the gravitational phase simplifies to:

$$\phi_g^{(2)} = -\frac{qg}{2}[2T_{21}^2 + 4T_{21}\Delta t + \Delta t^2]. \tag{2.43}$$

Since the gravitational acceleration, g, only appears in this term in the final signal, there are two ways that one can extract the value of g from measurements of echo signals. The first is done by measuring the phase of the signal within the echo envelope for a fixed T_{21}. This approach is substantially hindered by the difficulty of extracting an accurate phase measurement from an oscillatory signal whose amplitude is changing rapidly on the scale of the oscillations. Example echo shapes are shown in fig. (2.5) parts a) and b).

The second way to measure g from an echo signal is to vary T_{21}, which varies the phase of the signal due to the changing gravitational phase, but also modulates the amplitude at the recoil frequency. As a result, the SNR of a measurement done with this technique will vary, depending on how close T_{21} is to a multiple of τ_q.

Some example signals for this measurement method are shown in fig. (2.5) parts c) and d).

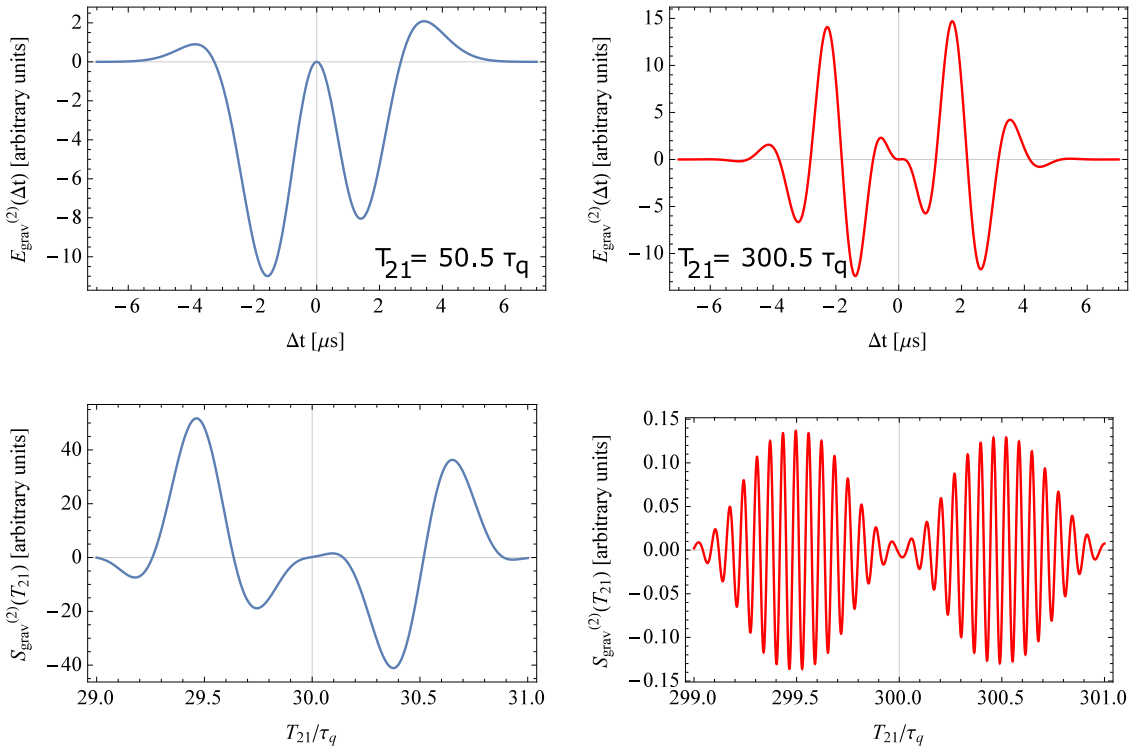

Figure 2.5: Plots of the expected behaviour of the two-pulse echo signal with gravity for different experimental timescales. For all graphs, $u_1 = u_2 = 1$, $T = 10\ \mu K$ and g is fixed at 9.81 m/s^2. The top two graphs show actual echo envelopes for both low and high values of T_{21}. Note that the frequency of the signal oscillation within the echo envelope increases with increasing T_{21} due to the acceleration of the atoms due to gravity. For the bottom two graphs, the echo has been integrated over Δt and the result plotted as a function of T_{21}. The signal oscillates both at the recoil frequency ω_q, and at the gravitational frequency ω_g, which increases linearly with T_{21}.

Regardless of the specific method used to measure g, it is necessary to measure the phase of the read-out beam. This requires a stable phase reference, so that the phase of the reflected light has a specific meaning in the context of the measurement.

Usually, this reference is taken to be the node in the SW potential that is formed at the surface of the retro-reflecting mirror that is used to generate the SW.

3 Atom lattice formation effects and studies of Bragg scattering

3.1 Introduction

Bragg scattering of light is a well-established and widely-used technique for studying the properties of ordered structures [68]. In early work [69], these structures consisted of crystal planes, from which it was possible to scatter coherent X-rays. This technique later became a standard tool for examining ordered structures in chemistry and biophysics [70]. Recent advances in technology have allowed for the manufacture of micro-fabricated structures such as distributed Bragg reflectors, which use the properties of Bragg scattering to act as narrow-band wavelength selectors for laser systems [71]. It is also possible to use this technique to investigate ephemeral ordered structures in ensembles of cold atoms, such as optical lattices [28]. We investigated the coherent transient scattering of light from a 1-dimensional cold-atom density grating produced via the grating echo technique. We have also examined the parameters that affect the efficiency of Bragg scattering from this grating, namely density, light detuning and sample temperature.

In refs. [33, 58], statistical precision has been demonstrated at the level of 37

parts per billion (ppb) and 75 ppb, for the recoil and gravity measurements, respectively. While these results approach the statistical precision of the leading methods [16, 72], substantial characterization of systematic effects is still necessary. The signal-to-noise ratio (SNR) in previous echo AI measurements, including the above-mentioned recoil and gravity measurements, was limited by a grating reflectivity of $\sim 0.2\%$, which implies a scattering ratio of ~ 0.0004 photons/atom, which is far below the theoretical limit of 1 photon/atom (see section 3.2). The mechanisms responsible for the limitation in reflectivity are multiple scattering effects associated with sample density, and grating contrast. We therefore consider the properties of Bragg scattering as it applies to the grating echo AI with the main motivation being to gain a deeper understanding of these effects.

An effective method for improving the grating contrast of an echo interferometer was demonstrated in ref. [30], where atoms were channeled into an optical lattice by application of an extended first AI pulse. The more recent work of Schilke et. al. [31] demonstrated scattering efficiencies approaching unity by transferring atoms from a MOT into an off-resonant dipole trap and carefully tuning the spatial frequency of the resulting lattice to maximize reflection of a probe beam.

The scattering of light from periodic structures in cold atoms has been the subject of much study, both in the case of atoms in optical lattices [28, 29, 31], and in the case of structures generated by interferometry [64]. The work presented in Ref. [28] demonstrates the typical Bragg reflection spectrum for a spatially-modulated atomic density distribution. The authors of this work also demonstrate the dependence of the shape of the spectrum on atomic density, and record a maximum reflectivity of $\sim 0.3\%$. The shape of these reflection spectra was explained

in ref. [29] using a transfer matrix formalism. This work also demonstrated a maximum reflectivity of $\sim$ 30%. The highest recorded reflectivity from an ordered sample of cold atoms is $\sim$ 80% [31], suggesting the possibility of a significant improvement in the Bragg scattering efficiency of the echo AI. Our technique is a marked departure from many previous experiments [28, 73] in which light is scattered from an equilibrated optical lattice, in that ours is a highly dynamic, transient structure.

One particularly interesting property of the echo AI is that even a very small "bunching" of the atoms towards the nodes of the SW excitation pulse is enough to produce a measurable back-scatter of the read out light. One of our goals was to quantify the grating contrast and scattering efficiency of the AI. First presented are results of a single-pulse experiment as in ref. [60]. Here, a single SW excitation of duration 1 μs is applied to the atoms at $t = 0$. Immediately thereafter, the read out pulse is applied to detect any $\lambda/2$-periodic spatial modulation in the sample. In a sample with no initial velocity distribution, we would expect that the maximum in the back-scattered signal would occur when the atoms had been channeled to the nodes in the SW potential. For our conditions, the dynamic evolution of density modulation in the sample is rapidly washed-out by the velocity distribution of the sample along the SW axis. For this reason, we find a peak in the back-scattered signal within a few μs of the end of the SW excitation. We use numerical simulation techniques [74] to quantitatively understand the magnitude of the density modulation for the first time. By increasing the pulse length, we explore the channeling of the atoms in the SW potential [75–78], comparing the results to simulations.

To understand how the echo AI signal size is affected by multiple scattering, we make use of the two-pulse AI technique, outlined earlier. Here, we vary the read out detuning for a range of sample densities, recording the reflection spectrum for each case. As expected [28], the two-peaked spectrum of a high-density sample evolves into a single-peaked spectrum at low density. Our simulations suggest that in the low density regime we approach the theoretical limit for the scattering efficiency of one photon/atom. The reduced density also decreases the influence of the index of refraction of the sample [79], an important systematic effect in measurements of recoil [58] and gravity [33].

The remainder of this section is organized as follows: in section 3.2 we describe the theoretical framework for our analysis, in section 3.3 we outline the experimental procedures, and in section 3.4 we discuss the results of our measurements and compare them to simulation.

3.2 Theory and Simulations

When a near-resonant beam is incident normal to a set of atomic planes separated by $\lambda/2$, there is a strong attenuation of the transmitted field along the forward direction due to interference between the forward and backward traveling beams [67, 73, 80, 81]. This gives rise to a strong enhancement in the reflected field compared to a disordered medium, which would show no particular directionality of the scattered/re-emitted light. For sufficiently dense media or for media with a large number of atomic layers, multiple reflections between adjacent atomic layers become important and a high reflectivity can be produced over a wide range of

frequencies (i.e. a photonic band gap, or PBG) where light has a low probability of transmitting through the medium.

Throughout this work, we restrict ourselves to the case of 1-D Bragg scattering. This is because the read out light is applied along the axis defined by the SW. Furthermore, the entire experiment is performed under conditions in which the frequency of the excitations and read out differ by only about 1 part per million (ppm). Our studies represent a system in which a fixed-period grating is probed by small changes in the frequency of the read out light. This means that the variation in back-scattered intensity can be entirely attributed to changes in the grating contrast, which is controlled by careful tuning of the pulse parameters and the sample density. Atomic grating-based experiments are, to first-order, unaffected by the internal structure of the atoms. As shown in fig. (3.1), at the echo time, $2T_{21}$, the atomic momentum states that differ by $2\hbar k$ are superimposed at periodic sites along the SW axis. At particular values of T_{21}, determined by the recoil frequency, these positions are separated by $\lambda/2$, and result in maximum scattering efficiency. At such times, application of a near-resonant read out pulse causes the atomic wavefunction, which is in a superposition of momentum states, to collapse into a single state with no $\lambda/2$-periodic component. Therefore, multiple coherent scattering events from a single atom are not possible. This is the origin of the one photon/atom limit pertaining to scattering efficiency. Since this is an elastic scattering process, momentum conservation requires that any coherently-scattered light be back-scattered. It is worthwhile to note that although this is a single-atom process, our signal arises from the contributions of many randomly-distributed atoms whose positions and velocities are given by uncorrelated Gaussian distributions.

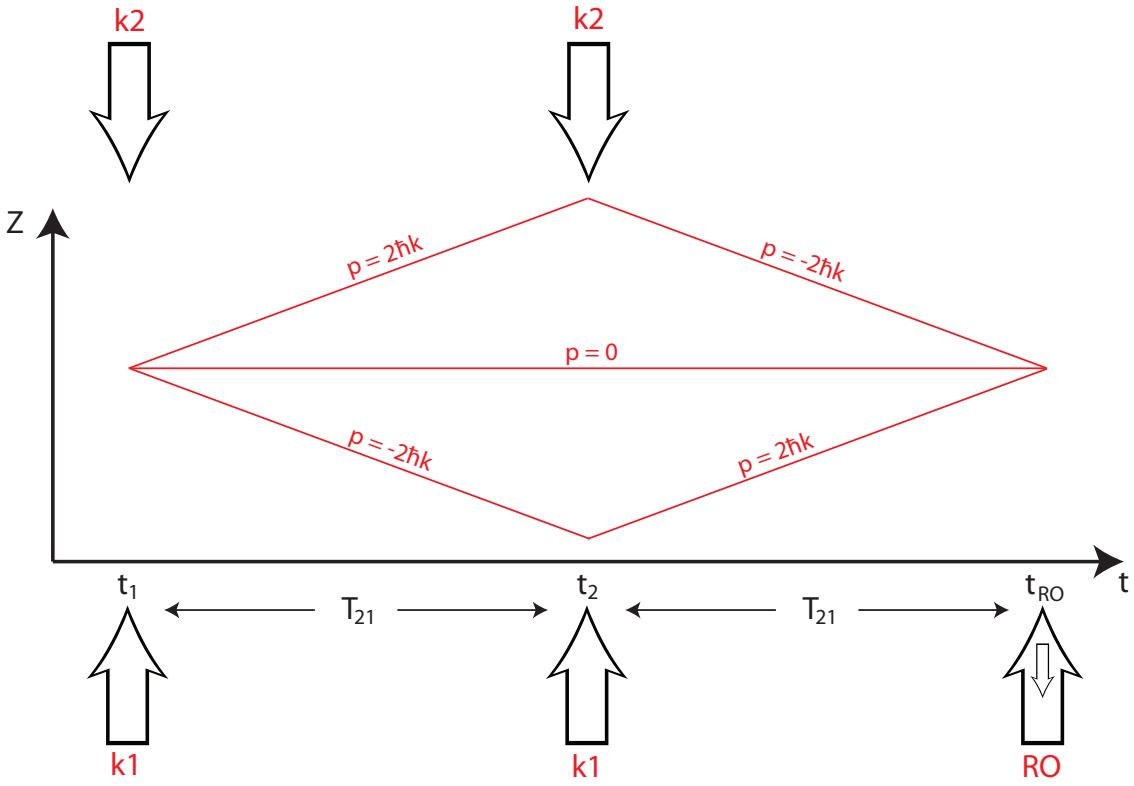

Figure 3.1: Simplified recoil diagram for Bragg scattering and lattice formation measurements.

In the semi-classical treatment, the atoms in our sample can be treated as a collection of dipole radiators. For such a collection, the peak intensity of the radiated light in the far-field is expected to scale as N^2, where N is the number of radiators. In the presence of multiple scattering effects, we expect significant deviation from this ideal behavior.

The reflectivity of the grating can be characterized by the Debye-Waller factor [82], which is given as: $\beta = exp(- < u^2 > K^2)$, where $< u^2 >$ is the mean-squared width of the reflection sites, and $K = 2k$ is the momentum transfer due to the Bragg scattering [28]. For a sample devoid of density modulation, $< u^2 > \simeq \infty$, meaning that $\beta \simeq 0$. In the case of a high-contrast grating, the localization can be as narrow as $\sim \lambda/10$ [83], resulting in $\beta = 0.45$. For our conditions, we have $\beta \simeq 0.1$, resulting in a reduced reflectivity.

To establish the reflectivity of the atomic sample, we first state that:

$$E_R(t) = r(t)E_I, \tag{3.1}$$

where E_R is the reflected electric field, $r(t)$ is the time-varying reflectivity of the sample, and E_I is the incident field. Since the electric field is a complex quantity, the reflectivity is similarly complex-valued. The Fresnel reflectivity from a dielectric boundary is given by:

$$r(t) = \frac{n_1(t) - n_2(t)}{n_1(t) + n_2(t)}, \tag{3.2}$$

where $n_{1,2}$ is the refractive index of the high-density and low-density portions of the atomic distribution. Here, we are simplifying matters by assuming that the atomic

density varies stepwise rather than continuously.

The refractive index of a dilute gas is given by [84]:

$$n \simeq 1 - \rho \frac{\mu_{eg}^2}{2\epsilon_0 \hbar \Gamma} \frac{\Delta/\Gamma}{1 + (\Delta/\Gamma)^2}, \tag{3.3}$$

where ρ is the density, μ_{eg} is the dipole moment between the excited and ground states, Δ is the detuning of the light from resonance, and Γ is the atomic linewidth.

The Bragg-scattered optical power from atoms confined to a periodic potential is given by [26, 85]:

$$P_R = 2 \left(\frac{\pi \beta}{\lambda D} \right)^2 N^2 \left| \frac{\alpha(\Delta_B)}{\epsilon_0} \right|^2 I_I, \tag{3.4}$$

where D is the diameter of the atomic sample that is illuminated by the interrogation beam, N is the number of atoms in the sample, λ is the wavelength of the interrogation light, Δ_B is the detuning of that light from the Bragg condition, β is the Debye-Waller factor and $\alpha(\Delta)$ is the detuning-dependent atomic polarizability of the sample, given by:

$$\alpha(\Delta) = \frac{3\epsilon\lambda^3}{4\pi^2} \frac{S_{F,F'}^2}{i + 2\Delta/\Gamma}, \tag{3.5}$$

where $S_{F,F'}$ is the oscillator strength corresponding the the transition $(F \rightarrow F')$.

The reflection coefficient of the periodic atom structure is then given by:

$$R(\Delta_B) = \frac{18}{\pi^3} \left(\frac{\lambda_B}{D} \right)^4 \beta^2 N^2 \frac{S_{F,F'}^2}{1 + 2\Delta_B/\Gamma^2}, \tag{3.6}$$

where λ_B is the spatial wavelength of the lattice, D is the diameter of the atoms as

seen by the read out light, N is the total number of atoms, Δ_B is the detuning of the read out light from the Bragg condition, and Γ is the natural linewidth of the atoms.

A numerical simulation of the experimental effects studied in this work was performed by Brynle Barrett [86], and his description of the results is presented in Appendix B. In section 3.4, whenever appropriate, the simulation results are presented side-by-side with the experimental results, for qualitative comparison.

3.3 Experiment

Our trapping and excitation beams were derived from the same Ti:Sapphire laser, operating at $\sim$ 1.2 W output power. The laser is stabilized with respect to a hyper-fine absorption peak from a saturated absorption spectrum recorded from a sample of room-temperature gaseous rubidium. The various detunings used in the experiment are achieved with accousto-optic modulators (AOMs).

Figure (3.2) shows the "core" of our optical setup where the Ti:sapph laser is locked to the appropriate peak in the saturated absorption spectrum of ^{87}Rb. The overall detuning of the entire experiment is controlled by an dual-pass AOM (DPAOM) called the "Locking" AOM, which operates at -236 MHz, producing a total shift of -473 MHz. The output of this AOM is aligned through a saturated-absorption setup for frequency stabilization. The laser is locked to the lowest-frequency cross-over peak in the spectrum with a lock-in amplifier. Due to the shift of the Locking AOM, the Ti:sapph output will be at $\omega_0 + 128$ MHz. We generate the appropriate optical frequency for the MOT by using another DPAOM called the

"Trapping" AOM, which shifts the light a total of -148 MHz to generate an optical frequency of $\omega_0 - 20$ MHz, necessary for trapping. After the turn-off of the trapping light, the undiffracted beam from this AOM is used for atom interferometry. There are two different AOM chains used in the one-pulse and two-pulse experiments that are shown in fig. (3.3) and fig. (3.5), respectively.

The additional optics for the one-pulse experiment are shown in fig. (3.3). Here, the light from point "A" is split by a pair of $\lambda/2$-PBS combinations to feed both the AI and RO AOMs. The AI AOM upshifts the light away from resonance to $\omega_0 + 378$ MHz, which is the frequency of the SW AI beams. The light for the RO beam is first sent through a dual-pass "gate" AOM, which increases its detuning by 160 MHz. This light is then sent through a single-pass, "read-out" AOM, which reduces the detuning to 35 MHz above resonance. This detuning corresponds to the highest scattering efficiency for our parameters, as shown in fig. (3.9). The gating AOM increases the isolation ratio of the RO switching, which is necessary due to the near-resonant frequency of the RO beam. Such near-resonant light applied to the atoms during the AI pulse would result in signal loss. Figure (3.3), in conjunction with fig. (3.2) describes the complete optics setup for the generation of the AI and read out light for the one-pulse experiment.

The light derived by the optics in fig. (3.2) is sent to the experiment through an optical fiber for MOT generation. The traveling-wave components of the SW excitations are sent to the atom trapping apparatus through two additional independent fibers. The atom trapping apparatus, described in ref. [87] is placed on a simple wooden table with no vibration isolation. The outputs of the AI fibers are directed along the vertical, with the up-going beam called k1 and the down-going

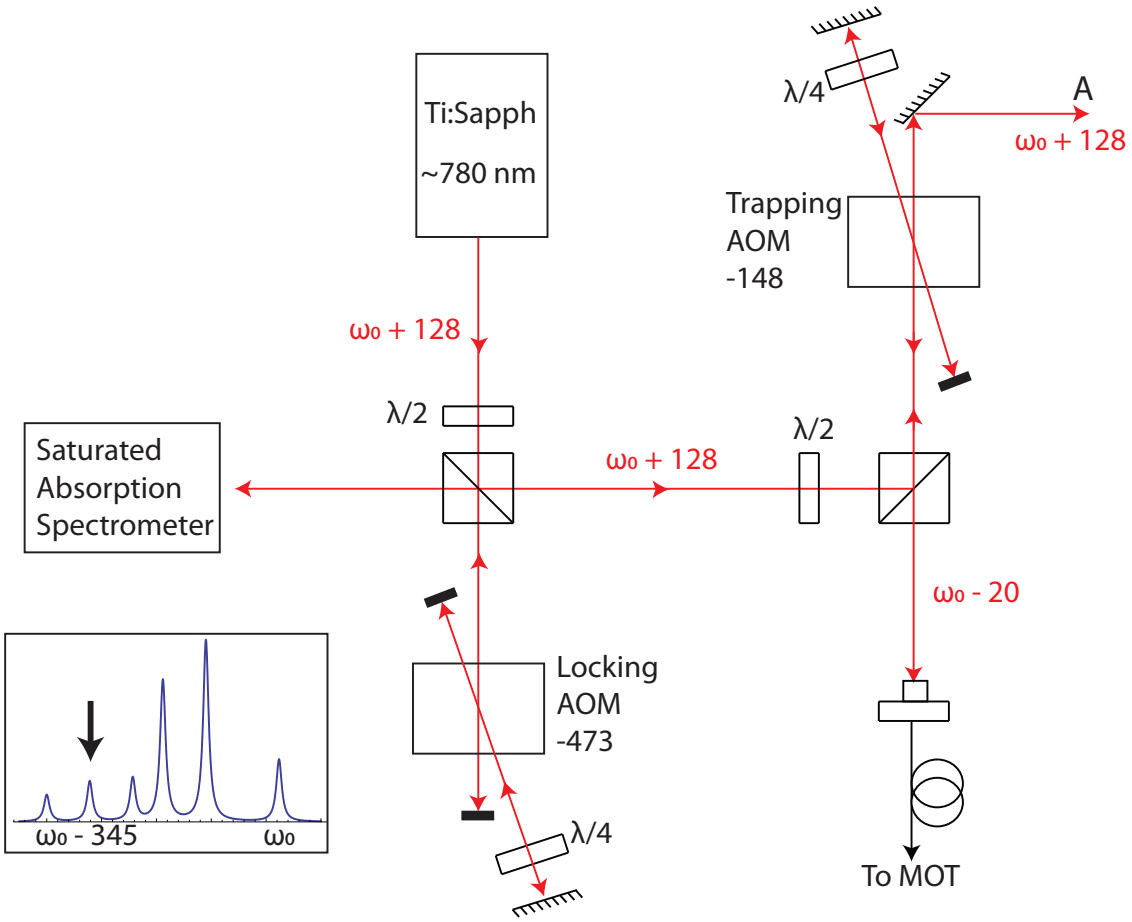

Figure 3.2: Schematic of the optics common to both one-pulse and two-pulse experiments with inset of ^{87}Rb "trapping" transition spectrum. Beam detunings are marked with respect to ω_0, which is the $|F = 1\rangle \rightarrow |F' = 2\rangle$ transition frequency in ^{87}Rb. AOM frequency shifts are marked with respect to the input light. All shifts and detunings are in units of MHz unless otherwise noted. The undiffracted beam from the Trapping AOM is used to generate the AI and read out beams for either the one-pulse or two-pulse experiment. The optics diagrams for these experiments are shown in fig. (3.3) and fig. (3.5), respectively, and they connect to this diagram at the point labeled "A".

beam called k2. The read out light is coupled through the same fiber as the k1 beam, and the back-scattered signal is detected along k2 using a gated photomultiplier tube (PMT).

The gating circuit (built into the PMT) disables the amplification stages of the PMT by providing a large positive voltage between the photocathode and the first dynode in the PMT during the SW excitation. However, there is still a recovery time after the intense k2 excitation beam strikes the PMT's photocathode. For this reason, the PMT can be temporarily "blinded" by the depletion of ejectable electrons from the photocathode surface. Because the one-pulse grating forms within 1 μs of the end of the excitation, the PMT does not have enough time to recover from the effect of the k2 beam before the arrival of the signal light. Therefore, in this setup, it is impossible to record the "turn-on" behavior of the signal without blocking the k2 light from hitting the PMT. Mechanical shutters, with closing times on the order of milliseconds are much too slow for this application. Therefore an AOM, with a turn-on time of around 100 ns is used to deflect the signal light into the PMT, as shown in fig. (3.4).

The two-pulse experiment requires a more complicated optical setup, as shown in fig. (3.5), and connects to optics shown in fig. (3.2) at point "A". Due to the limited tuning bandwidths of the available AOMs, we used a chain of three DPAOMs to obtain a detuning range of $\omega_0 \pm 60$ MHz. The first AOM, called the "Resonant" AOM, is operated at -128 MHz (-64 MHz dual pass) so that its output is on resonance with the $|F = 1\rangle \rightarrow |F' = 2\rangle$ transition in ^{87}Rb. This light then passes through two tunable DPAOMs in series. The first is called the "Upshift" AOM, and can provide a frequency shift between +130 MHz and +190 MHz. The

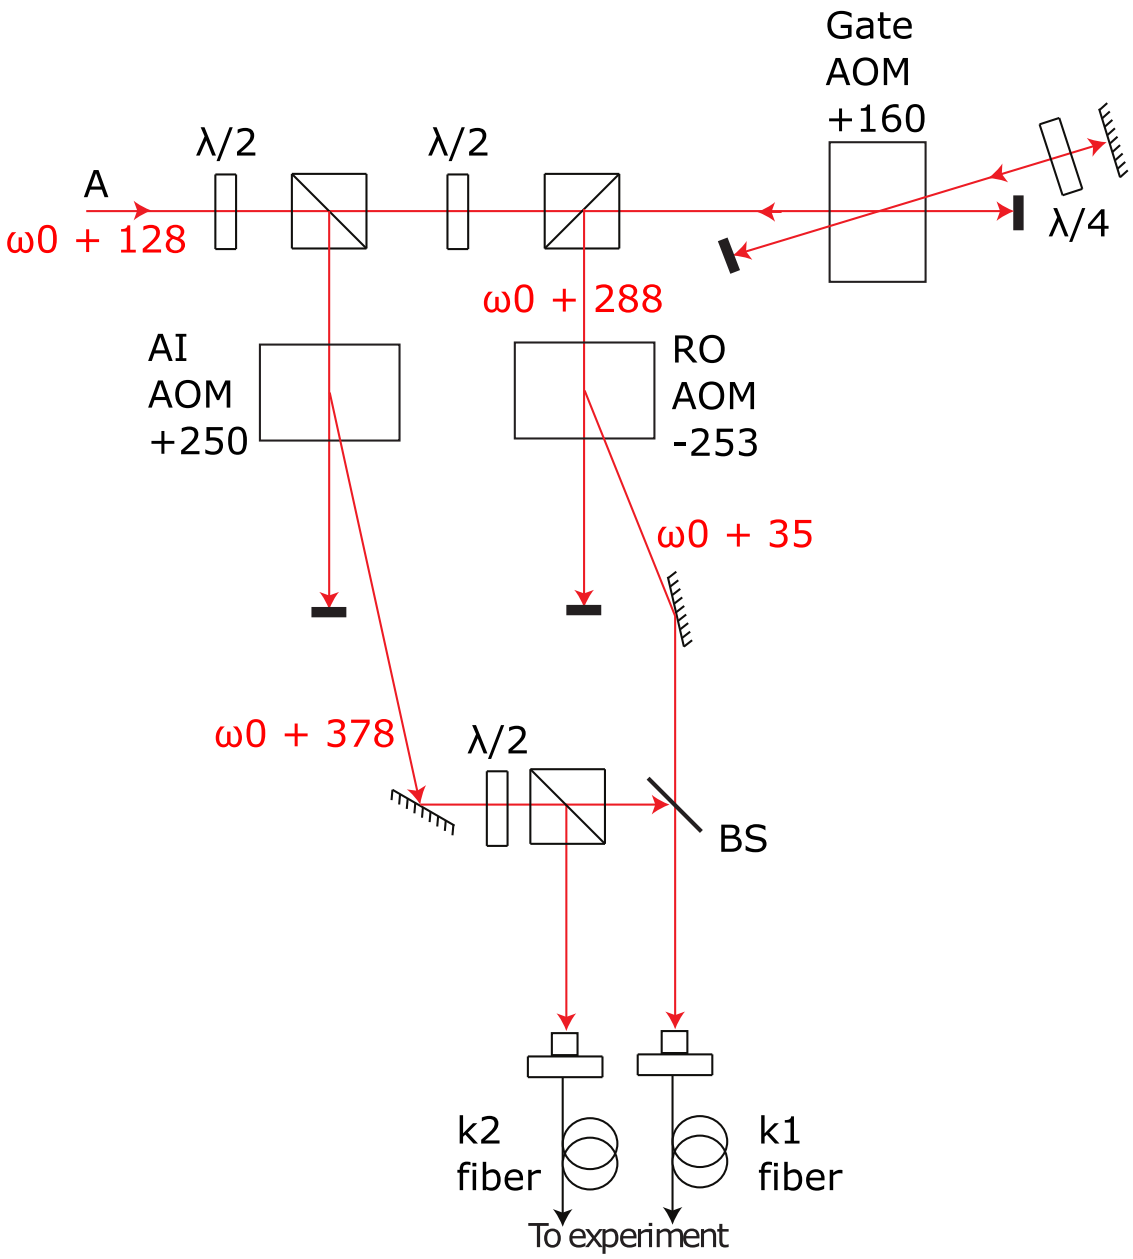

Figure 3.3: Schematic of the additional optics for the one-pulse experiment. The light derived by the optics in fig. (3.2) enters at the point marked "A". Beam detunings are marked with respect to ω_0, which is the $|F = 1\rangle \rightarrow |F' = 2\rangle$ transition frequency in ^{87}Rb. AOM frequency shifts are marked with respect to the input light. All shifts and detunings are in units of MHz unless otherwise noted.

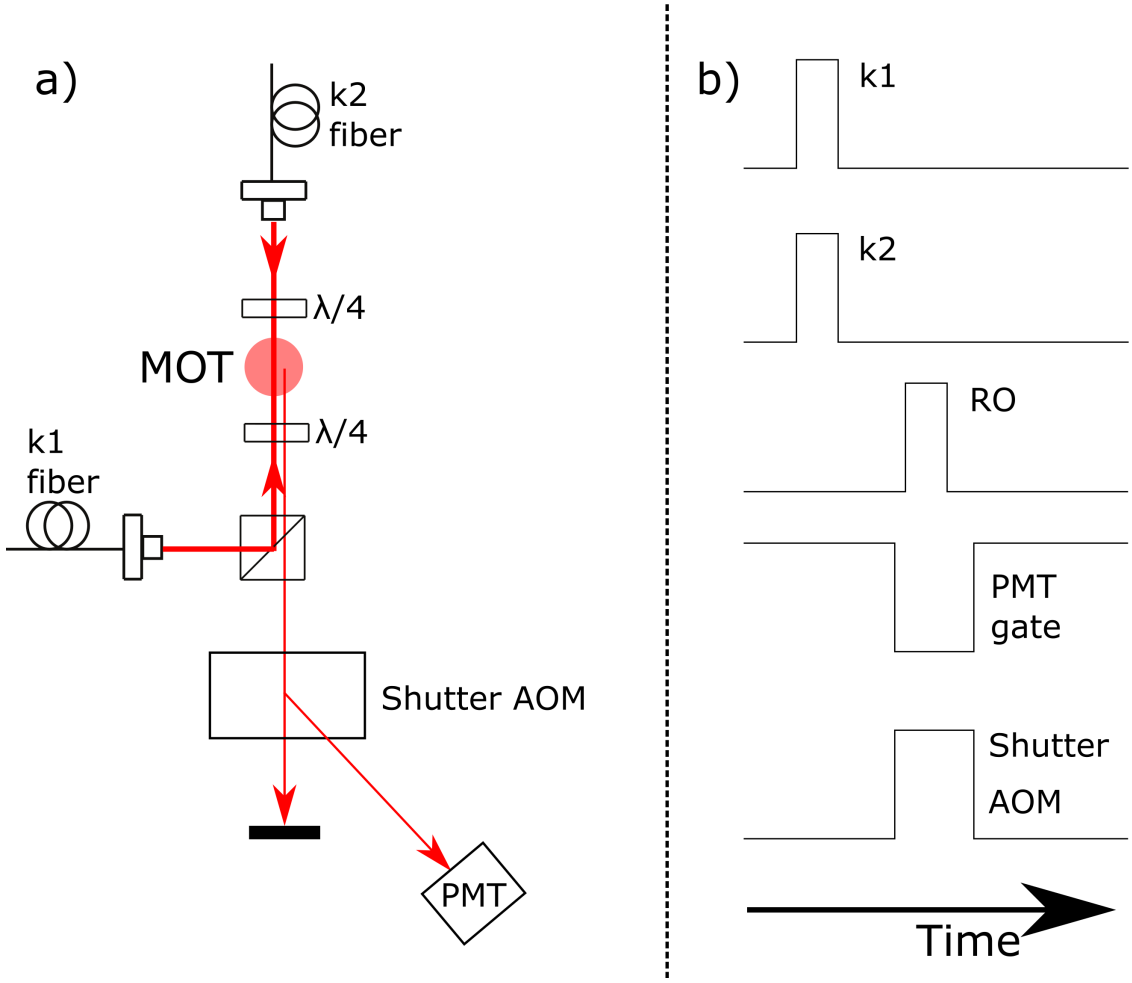

Figure 3.4: Detection optics and timing diagram for the one-pulse experiment. Part a) shows the optical layout of the detection scheme, including the outputs of the fibers from fig. (3.3). The k1 and k2 light is applied along the vertical direction, followed closely by the RO light, which comes in along the k1 direction. The RO light is reflected from the periodic structure that forms in the MOT due to the effect of the SW excitation. The reflected RO travels back through the shutter AOM, and strikes the gated PMT. In part b), the timing diagram for an experiment is shown. The k1 and k2 beams are turned-on simultaneously, followed shortly by the RO. At the time of the RO, the shutter AOM is activated and the PMT gate is deactivated to allow detection of the weak signal. Both of these events are timed to occur 100 ns before the arrival of the RO pulse.

next is called the "Downshift" AOM, and provides the same tuning range, but as a downshift. Since these two AOMs are independently-controlled, a final tuning range of $\omega_0 - 60$ MHz to $\omega_0 + 60$ MHz is achieved. To compensate for the power losses incurred through these three AOMs, which amounts to a factor of 50%, the light is then amplified by passage through a tapered-amplifier (TA).

3.4 Results

Here, we present the results of varying the available experimental parameters, and compare, if possible to the same results from the numerical simulations [86]. In each case, the comparison shows good qualitative match between experiment and simulation.

Figure (3.6) shows the evolution of the single-shot signal amplitude. To record this plot, we apply a traveling-wave pulse to the atomic sample immediately after a SW excitation, and integrate the intensity of the coherently back-scattered light over Δt. By keeping the intensity of this pulse well below the saturation intensity, we ensure that the grating dephases without the influence of spontaneous emission effects. It is important to note just how slight the movement of the atoms is due to these SW pulses, and how quickly we can detect the contrast of the grating. Classically, an atom that receives a momentum kick of $2\hbar$k will move at $2\hbar \mathrm{k}/\mathrm{M_{Rb}} \simeq$ 1.2 cm/s. This atom will have moved $\sim$ 12 nm in the ~ 1 μs that it takes for the signal to become discernible, as shown in fig. (3.6).

To investigate the motion of the atoms in a SW potential, we apply a pulse whose duration is sufficiently long that the Raman-Nath condition is violated. However the

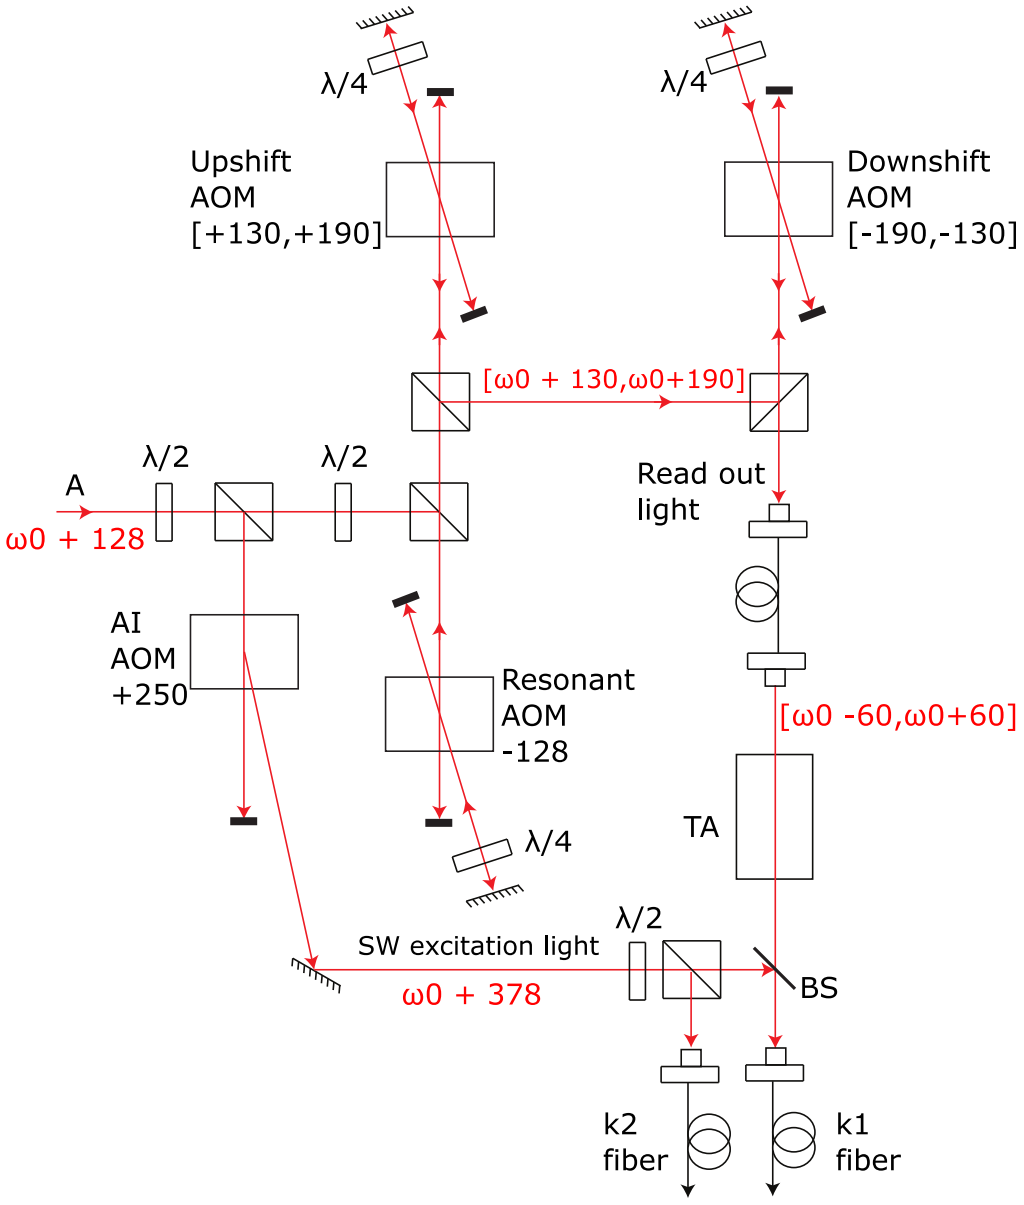

Figure 3.5: Schematic of the additional optics for the two-pulse experiment. The light derived by the optics in fig. (3.2) enters at the point marked "A". Beam detunings are marked with respect to ω_0, which is the $|F = 1\rangle \rightarrow |F' = 2\rangle$ transition frequency in ^{87}Rb. AOM frequency shifts are marked with respect to the input light. All shifts and detunings are in units of MHz unless otherwise noted.

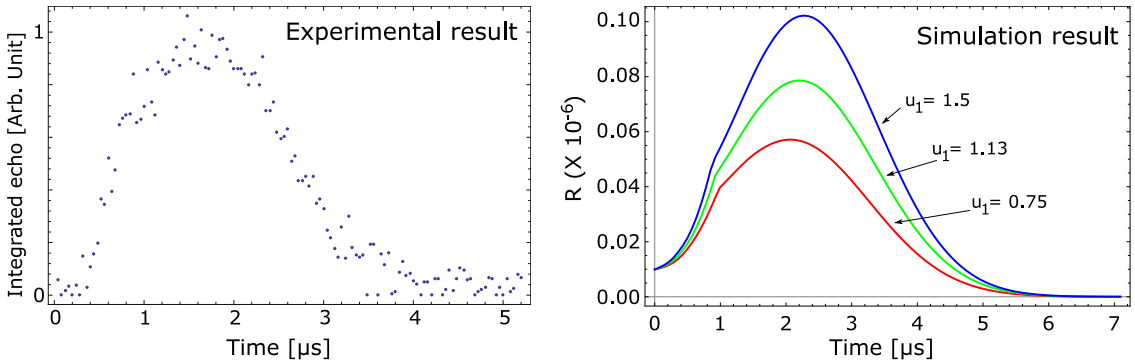

Figure 3.6: Rapid evolution of the density grating contrast following a single AI pulse. After ~ 1.5 μs the contrast begins to decrease. This is a result of the non-zero width of the atomic velocity distribution. By ~ 4 μs after the SW pulse ends, the grating contrast has completely disappeared. Each data point is the result of integrating the complete echo over Δt.

length of this pulse must be shorter than the ~ 20 μs limit imposed by the motion of the mirrors due to mechanical vibrations. Immediately after the application of this pulse, a read out pulse is applied, and the back-scattered light is again integrated over Δt. Our results are consistent with work done by groups working with long, intense optical lattice pulses [83], where the atoms are "channeled" towards the nodes in the potential.

As shown in fig. (3.7), as the length of the excitation pulse increases, there is a maximum in the reflectivity of the sample, reached after $\tau_{osc} \simeq 4$ μs, which is consistent with the oscillation period of atoms in a harmonic potential, $T_h = \lambda_L(M_a/4V_0)^{1/2}$, where λ_L is the wavelength of the laser light, M_a is the mass of the atom and $V_0 = \hbar\Omega^2/8\Delta_L$ is the SW potential. This period corresponds to twice the time it takes the atoms to move $\lambda/4$. This displacement generates a $\lambda/2$-periodic structure from which light can be Bragg-scattered. A further peak in reflectivity is seen near time $3\tau_{osc}$, where the atoms have reformed the $\lambda/2$-periodic structure.

We do not observe additional revivals of the reflectivity, which we attribute to long-term instability of the SW phase.

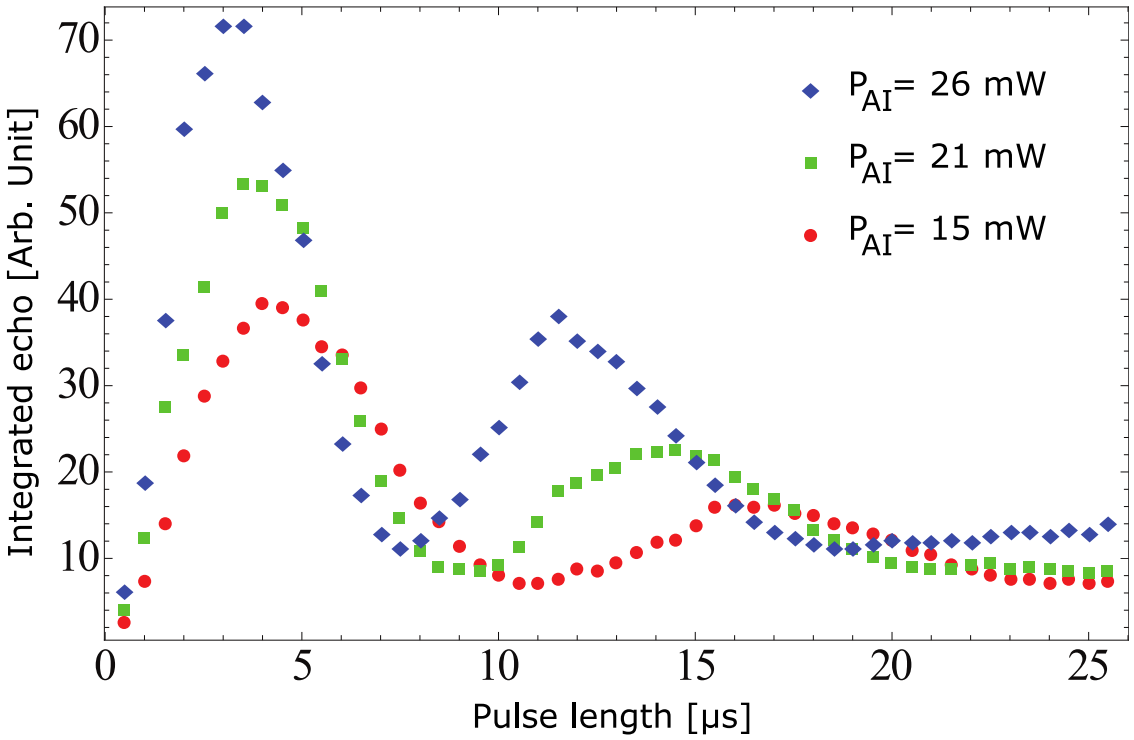

Figure 3.7: Plot of the variation in the period of contrast revivals as a function of AI light power. As the power increases, the peaks in the back-scattered signal occur closer together. Each of these data sets were taken by varying the length of the SW pulse and then applying a single traveling-wave RO pulse 100 ns after the turn-off of the AI pulse. Each data point is the result of integrating the complete echo over Δt.

We now discuss the properties of the back-scattered echo signal from the two-pulse AI. Figure (3.8) shows the echo signal envelope for various applied read out intensities. Data is shown for both high and low sample densities, where the low-density regime is achieved by allowing the MOT to expand before performing the experiment. This data was recorded with a AI beam detuning of +378 MHz and a read out detuning of +40 MHz. The shape of this signal has been studied as a

function of excitation pulse area and detuning [62], and modeled with numerical simulations [60, 74].

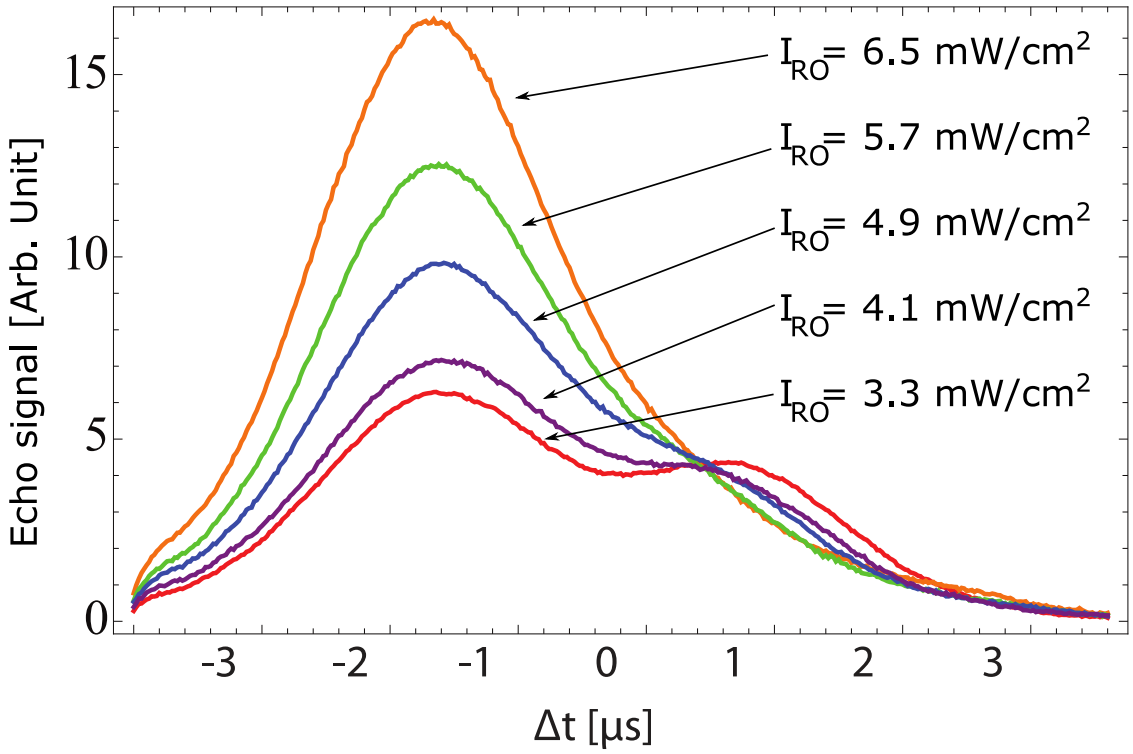

Figure 3.8: Plot showing the effect of increasing the read-out intensity in the two-pulse experiment. As the intensity is increased, a gradual transition from a two-lobed structure to a single peak is observed. This behavior is due to the tendency of an intense read-out pulse to decohere the sample during its application, resulting in a loss of signal after the echo time, $\Delta t = 0$. Note that each of these curves represents a single echo: no integration has been performed.

Finally, figure (3.9) shows the effect of varying the RO detuning from atomic resonance for a two-pulse experiment. We see two peaks in the resulting curve, at ± 30 MHz. This behaviour can be explained by realizing that for resonant RO light, there is no preferred direction for the scattered light due to the isotropic nature of spontaneous emission. As a result, we see a minimum in the reflected light at

the resonant frequency. The simulation results show some "jagged" features that aren't present in the experimental data, which we attribute to a lack of averaging over the echo envelope in the simulations. Otherwise, there is a good qualitative match between the simulation results and the experimental data.

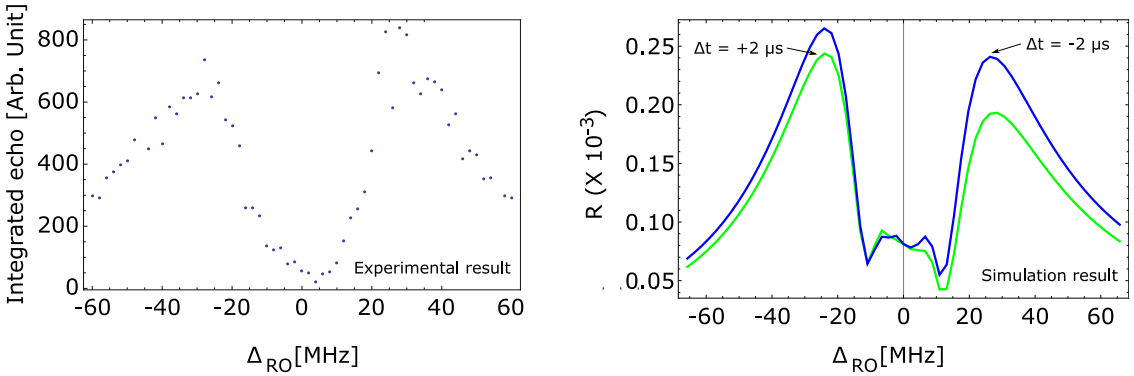

Figure 3.9: Plot of the optical energy back-scattered from the grating echo as a function of read-out detuning with comparison to simulation results (see appendix B for details on simulations). The dip at zero detuning is caused by spontaneous emission effects and the high density of our atomic sample. Each data point is the result of integrating the complete echo over Δt. The slight asymmetry of the spectrum is caused by a small misalignment of our read-out beam with respect to the density grating's principle axis.

4 AI gravimeter

Cold atoms are unique test samples for acceleration measurements. Due to the high-vacuum needed for the production of cold atoms, the atoms are free to move in a drag-free environment [4, 13]. More importantly, the atoms have a modifiable internal structure, allowing resonant interactions with light fields and the encoding of inertial information in the atomic wavefunction. The quantum nature of atoms allows for unique experiments wherein the interaction of the test mass (the atoms) with its environment can be carefully controlled. The cooling of the atoms allows for extended experimental timescales, not otherwise achievable with room-temperature atoms.

At the most fundamental level, atom-interferometric measurements of inertial effects (linear accelerations and rotations) are accomplished by exciting a sample of cold atoms with light pulses, and allowing the packets of atoms to travel along separated paths through space. In a frame in which the atoms experience an acceleration, the different paths will result in a differing rate of accumulation of phase in the atomic wavefunction, which manifests itself as interference when the atom packets are recombined at the end of the experiment. The theoretical development of the AI signal is presented in ch.2.

4.1 Atom trapping and experiment overview

The rest of this chapter will be presented in two subsections: the first will outline the experimental steps needed to produce a MOT and to perform an echo-AI experiment with it. Following this, a comparison will be made between the previous gravimeter apparatus used in refs. [32, 33]. The new apparatus has several notable improvements which will be described in detail.

4.1.1 Procedure for AI experiment

The experiment is performed in a series of steps, outlined below:

1. First, the MOT is loaded from the background vapour by allowing a set of 3 retro-reflected trapping beams to illuminate a volume inside a vacuum chamber containing rubidium vapour at a pressure of $\sim 5 \times 10^{-9}$ torr. This pressure, which is much lower than the room-temperature vapour pressure of rubidium, 3×10^{-7} torr, is maintained by the combination of a high-throughput ion pump, and a heated ampoule containing a small sample of atomic rubidium. The cell is heated gently with a resistive heater tape to increase the Rb pressure, while the ion pump works constantly to remove the Rb vapour as well as any undesirable gas species (from leaks, outgassing, etc.) from the vacuum chamber. Equilibrium pressure is attained when the rate of outgassing of the rubidium cell plus the rate of addition of non-rubidium species to the chamber matches the rate of removal by the ion pump. This condition is reached iteratively, by slowly increasing the heat applied to the rubidium ampoule until the pressure of rubidium in the cell is optimal for trap

formation. To ensure high pumping conductance to the ion pump, the vacuum chamber is constructed from large diameter (6" ConFlat) components having a constant inner diameter of 95.25 mm between the glass vacuum cell and the ion pump. To perform the initial pump-out and bake of the system, a small turbo-molecular pump with a mechanical backing pump is attached to the system through an all-metal valve. To prevent contamination of the chamber by oil from the backing pump, the flexible bellows connecting the backing and turbo pumps was immersed in liquid nitrogen during pump-down.

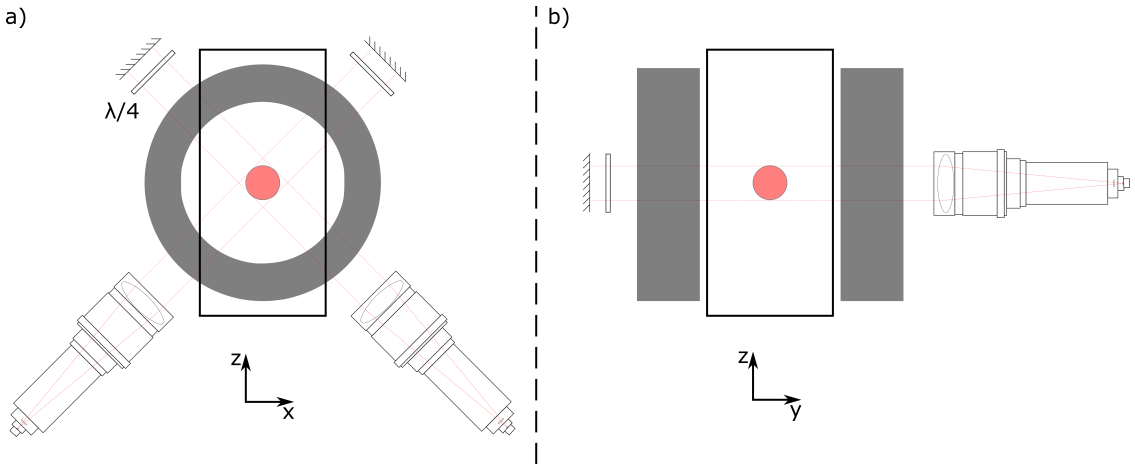

Figure 4.1: Schematic showing the arrangement of the trapping optics. The z-axis of the experiment is defined as the vertical axis, and it is the direction along which the interferometer pulses are applied. In the x-z plane, shown in part a), two trapping beams are aligned to be mutually perpendicular. Each beam is produced by a packaged fiber collimator with built-in quarter waveplates (Schäfter-Kirchoff 60FC). Each beam passes through the walls of the glass vacuum chamber to illuminate the trapping region before being retro-reflected through a quarter-wave plate. This generates the required σ^+—σ^- beam for trapping/cooling. Part b) shows the y-z plane, with the third collimator. In both parts of the diagram, the trapping coils are shown, as well as the region in which the trap will form.

2. Each trapping beam is circularly-polarized, and retro-reflected through a

quarter waveplate so as to produce a so-called σ^+–σ^- configuration along each of three mutually-perpendicular directions, as shown in fig. (4.1). The fiber collimators used in this setup (Schäfter-Kirchoff 60FC), have an integrated quarter-wave plate, which allows them to produce a 36 mm $1/e^2$ diameter beam with the required circular polarization. These collimators greatly simplify the optical design of the trapping setup compared to previous experiments in this group using only free-space optics. In this work, we will limit discussion to the trapping and manipulation of ^{87}Rb, the less-common naturally-occurring isotope of rubidium. This isotope is chosen due to its larger hyperfine level splitting compared to the more common isotope ^{85}Rb, as shown in fig. (4.2). This reduces the effect of off-resonant stimulated excitation of a transition during AI experiments.

3. The cooling force is generated by detuning the light below the $|F = 2\rangle \rightarrow |F' = 3\rangle$ transition by an amount roughly equal to 2.5Γ. Thus, any atom that moves away from the centre of the trapping volume will see a velocity-dependent force acting towards the centre, as shown in fig. (4.3). Confinement of the atoms in the trapping volume is accomplished by adding a magnetic field gradient with its zero at the centre of the trapping volume. This can be done with a set of "trapping coils", configured in the "anti-Helmholtz" configuration, which creates a constant gradient in the vicinity of the trapping volume. This gradient transforms a velocity-dependent force (absorption of a Doppler-shifted beam) into a position-dependent force. The anti-Helmholtz configuration is achieved by a pair of circular current-carrying loops with the same radius, R, positioned a distance R from each other. The two coils must

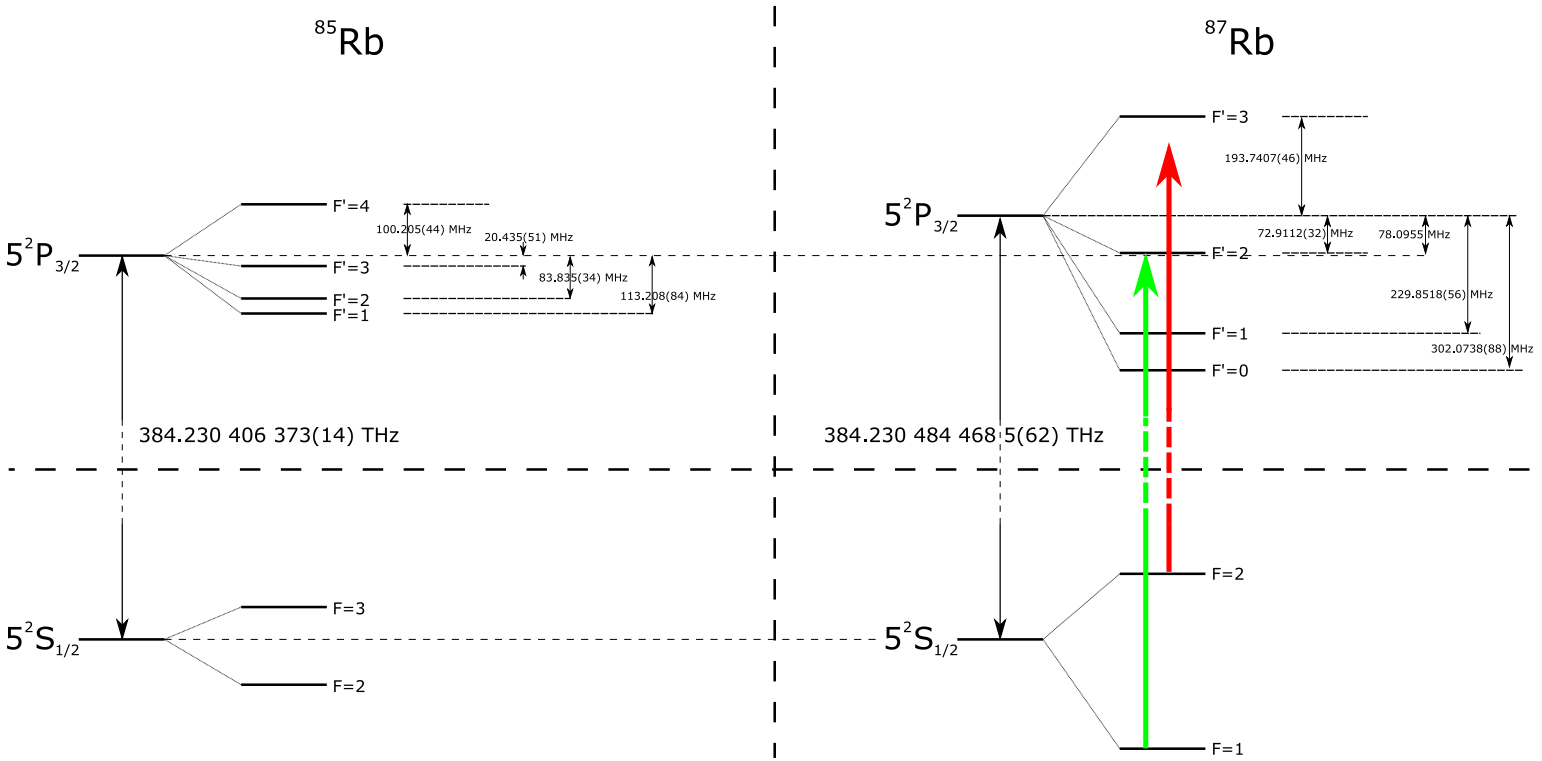

Figure 4.2: Level diagrams for the D2 line in ^{87}Rb and ^{85}Rb. The hyperfine splittings of the ground and excited states are shown to scale, but the scales are different for the ground and excited states and the separation between the excited and ground states are not to scale. The data in this figure are taken from refs. [88, 89], and represent the result of experimental measurements. On the right-hand side, the trapping (red) and repump (green) frequencies are marked. Note how the trapping light is red-detuned $\sim 2.5\Gamma$ from the $|F = 2\rangle \rightarrow |F' = 3\rangle$, which ensures optimal cooling of the atoms in the MOT.

be parallel, and the current must flow in opposite directions in each.

4. If the trapping beam polarizations are chosen correctly, the trapping force will have a magnitude proportional to the distance of the atom from the field zero. This is due to the Zeeman effect shifting the magnetic sublevels of the trapping transition closer to resonance for the beam that is directed towards the centre of the volume. The $|F = 2\rangle \leftrightarrow |F' = 3\rangle$ transition is called a cycling transition because atoms that are excited to $|F' = 3\rangle$ can only transition back to the $|F = 2\rangle$ ground state due to dipole selection rules, whether the transition is spontaneous or stimulated. It is, however, possible to excite atoms $|F = 2\rangle \rightarrow |F' = 2\rangle$ and even $|F = 2\rangle \rightarrow |F' = 1\rangle$. The other excited states are off-resonant, but have non-zero transition probabilities. Since it requires many thousands of absorption/emission cycles to significantly reduce an atom's velocity, even a very small transition probability can result in an atom being excited to the non-resonant excited state. Atoms that make this transition can be lost to the trapping cycle, since dipole selection rules allow for the spontaneous transitions $|F' = 2\rangle \rightarrow |F = 1, 2\rangle$ and $|F' = 1\rangle \rightarrow |F = 1, 2\rangle$. An atom that makes a spontaneous transition into the "dark" $|F = 1\rangle$ state will only be able to re-enter the trapping cycle if a repump beam is provided to resonantly excite the $|F = 1\rangle \rightarrow |F' = 2\rangle$ transition. This allows atoms to then spontaneously decay $|F' = 2\rangle \rightarrow |F = 2\rangle$ whence they can re-enter the trapping cycle. The repump beam doesn't need any particular polarization state, as it does not participate in the trapping/cooling. This is because in standard implementations of a 3D MOT, the trapping and repump light is combined into the same fibre using a polarizing beam splitting cube and sent

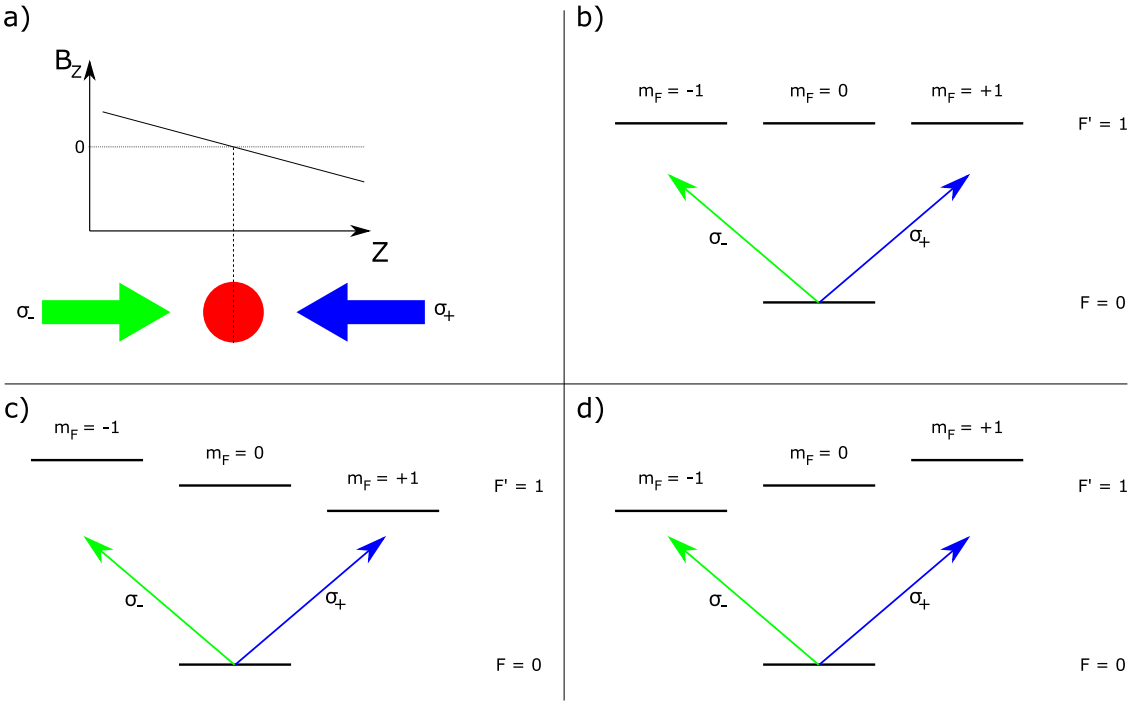

Figure 4.3: Generation of a position-dependent trapping force for MOT generation. This diagram uses a simplified atomic level structure for clarity. Part a) of the diagram shows one of the three sets of mutually-perpendicular beam pairs that are used to form the MOT. A plot of the uniform magnetic field gradient is shown above. In this system, the trapping volume (the space within the vacuum chamber in which the MOT will form) is illuminated from the right by a σ_+ polarized beam, and from the left by a σ_- polarized beam. Circularly-polarized beams are used to limit which magnetic sublevels participate in the trapping transitions. When the atom absorbs from the σ_+ beam, $m'_F = m_F + 1$. Conversely, when the atom absorbs from the σ_- beam, $m'_F = m_F - 1$. Part b) shows a simplified level diagram, wherein we consider a $|F = 0\rangle$ ground state, having a single magnetic sublevel, and a $|F' = 1\rangle$ excited state, which has three. Part c) shows the case where the atom has moved to the right (ie: z > 0), causing the $m_F = +1$ state to be Zeeman-shifted down, closer to resonance with the σ_+ beam. As a result, the atom will become more likely to absorb from the blue beam, resulting in a force that acts towards z $= 0$, pushing the atom back towards the field zero. In part d), the opposite case is shown, where the atom has moved to the left (ie: z < 0), causing the Zeeman shift to move the green beam closer to resonance, producing a force that pushes the atoms back towards z $= 0$.

to the trap, where only the trapping beam can have the correct polarization for trapping/cooling. At the end of the trapping process, a cloud of $\sim 10^9$ atoms has been collected into a volume of a few mm^3.

5. Due to the competing effects of cooling from the Doppler cooling effect and heating from spontaneous emission by the atoms, for a two-level atom, there is a limit to which we can cool the atoms using this process: $T_D = \hbar\Gamma/(2k_B)$. Here, T_D is the so-called Doppler limit temperature. For Rb, $T_D = 145\ \mu K$, which is too-high a temperature for long-timescale atomic interferometry experiments. Fortunately, polarization gradient cooling [90] allows for reduction of the atomic sample temperature below the Doppler limit. Fortuitously, for $\sigma_+ - \sigma_-$ trapping beams such as are used here, polarization gradient cooling occurs "for free", resulting in a theoretical lower limit of the recoil temperature, which is $\sim$ 184 nK for Rb. This limiting temperature is not achievable in practice due to the heating effect of the trapping beams having non-zero intensity without further molasses cooling. The equilibrium temperature of the trapped sample for typical conditions is closer to 11 μK [59].

6. It is possible to cool the sample below the above-mentioned equilibrium temperature by carefully controlling the final stages of trap loading. We refer to this as the "molasses cooling" step, a timing diagram for which is shown in fig. (4.4). To achieve this, after the MOT has grown to the required size, the trapping coils are turned off by use of a high-current transistor switch. The field is fully turned-off in around 100 μs. At this point, we perform 6 ms of molasses cooling which consists of chirping the trapping light further away

from atomic resonance (from 15 MHz detuning to 58 MHz detuning). As the frequency is chirped, the intensity of the light decreases due to the finite bandwidth of the AOM that is used to generate this beam. This is actually a desirable characteristic, as decreasing the light intensity further limits the amount of heating that is caused by the trapping force. The trap temperature is linearly-dependent on the light shift parameter, Ω^2/Δ, so decreasing the intensity allows temperatures below the Doppler limit, to a level as low as ~ 1 μK [87].

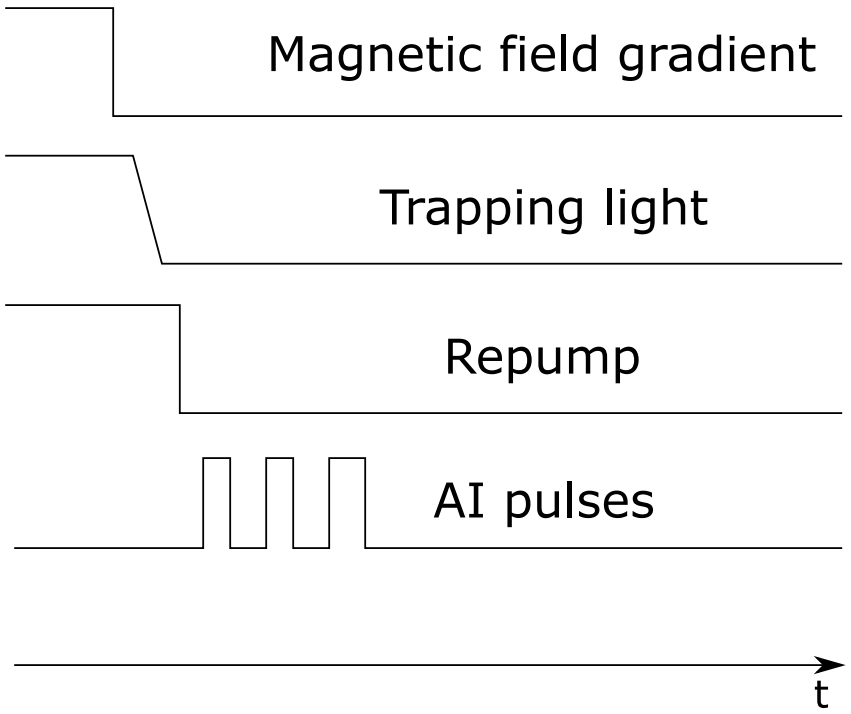

Figure 4.4: Trap timing diagram. Once the trap has been fully loaded, the trapping field is turned off, taking about 100 μs. Once the field is fully off, the frequency of the trapping light is ramped, as indicated by the slanted line, over 6 ms. At this point, the repump is left on for a further 50 μs to ensure that all atoms are optically pumped into the correct ground state. It is only then that the AI pulses can be applied to the sample.

7. Once the molasses cooling step has been completed, the trapping beams are turned-off. As soon as this is done, the AI pulses are applied to the atoms. Because of the nature of the grating-echo AI (see section 2), the atomic state populations that contribute to the final interference that is read-out at the end of the experiment will travel along different trajectories in space-time. The result is that there will be a differential accumulation of atomic wavefunction phase for each of the trajectories, resulting in a non-canceling phase term that will manifest itself at the time of the read-out. The total phase of the signal at the read out will contain terms that are proportional to the atomic recoil frequency, as well as components that depend on the inertial forces applied to the atoms, such as accelerations (due to gravity or movement of the apparatus) and rotations.

Since the AI signal is generated by application of SW excitations to the atoms, some consideration of the interaction probability for such interactions is warranted. For single photon-atom interactions, the probability of the interaction is determined by the combination of the atomic transition linewidth (assuming the existence of a real excited state) and the linewidth of the light. The centre of mass (COM) motion of the atom must be considered as well, as this will Doppler-shift the atomic line center one way or the other depending on the direction of movement between the atom and the light source.

For non-resonant two-photon interactions, the atomic lineshape is no longer considered. The interaction can be thought of as as excitation of the atom into a virtual state, followed by an immediate stimulated emission, which returns the atom to its ground state. In such interactions, the Doppler shift becomes much

more important. This is because the interactions of interest to the echo AI are those in which the atom absorbs from one TW component of the SW and then emits into the other. The frequency difference as seen by the atom between the two TW components will be equal to twice the Doppler shift produced by its motion. As such, for fixed-frequency beams, we are benefiting from the non-zero linewidth of the laser and the dominant effect of pulse-bandwidth broadening, which produces linewidths that overlap to a significant extent even for atoms that are traveling at a few m/s, as will be shown below.

Another important consideration involving Doppler shifts and two-photon resonance relates to the choice of atomic species used in the interferometer. As the atoms fall and the Doppler shift increases, the detuning of the SW light from one-photon resonance will decrease for one of the TW components. This means that for increasing experimental timescale, there will be an increasing number of atoms lost from the interferometer due to resonant interactions with one of the TW components of the SW. This contributes to the choice of ^{87}Rb as our atomic species of interest, as the increased hyperfine state splitting makes resonant excitation of the atoms less likely as the Doppler effect moves one of the components of the SW beam closer to atomic resonance.

To extend the experimental timescale, one must either increase the bandwidth of the laser by employing shorter pulses, or one must chirp the frequency of each component of the SW AI excitation beam in opposite directions. That is, the beam that is directed downwards must be chirped towards higher frequency, and the one directed upwards must be chirped towards lower frequency. In this manner, the atoms experience a true SW potential during the entire experiment. Both of these

approaches have their advantages and disadvantages. In particular, it is trivial to shorten the SW pulses when approaching long experimental timescales, within the limit of available laser power. There is an associated loss of pulse area, however, which can only be compensated for by increasing the power in the beam that is pulsed to create the SWs. In addition, the interaction probability will decrease as the experiment progresses, adding unwanted variability to the interferometer. Since it is usually not feasible to increase beam power, the beam-chirping approach is desirable. It requires no modification of the pulse duration, thereby preserving the pulse area, and the interaction probability remains constant for the entire duration of the experiment. The only downside is the increased experimental complexity: this approach requires arbitrary waveform generators.

4.2 Pulse-bandwidth effect

As the atoms fall and the centre frequency of the atomic line becomes shifted away from two-photon resonance with the SW light, the interaction of the atoms with the SW light field will become less effective. It is possible to maintain some amount of interaction strength by decreasing the width of the pulses used for the excitation. This exploits the uncertainty principle, which states that the product of the uncertainty in the bandwidth of a pulse, FWHM_ω, with its temporal width, FWHM_t, is a constant. The value of this constant depends on the shape of the pulse. For a pulse whose temporal power profile is shaped like a Gaussian, $\mathrm{FWHM}_\omega \mathrm{FWHM}_t = 0.441$, which is a property of Fourier transforms. In our case, the time profile of the pulses can be modeled as a piecewise function with Gaussian rising and falling edges sep-

arated by an interval of constant power. Some example pulse shapes can be seen in fig. (4.5a). Applying a Fourier transform to the electric field amplitude of the pulse shapes in this figure, yields the electric field frequency distribution of the pulse, which is then squared to yield the curves seen in fig. (4.5b). As shown, the FWHM of the frequency-space distribution scales inversely with the time-width of the pulse.

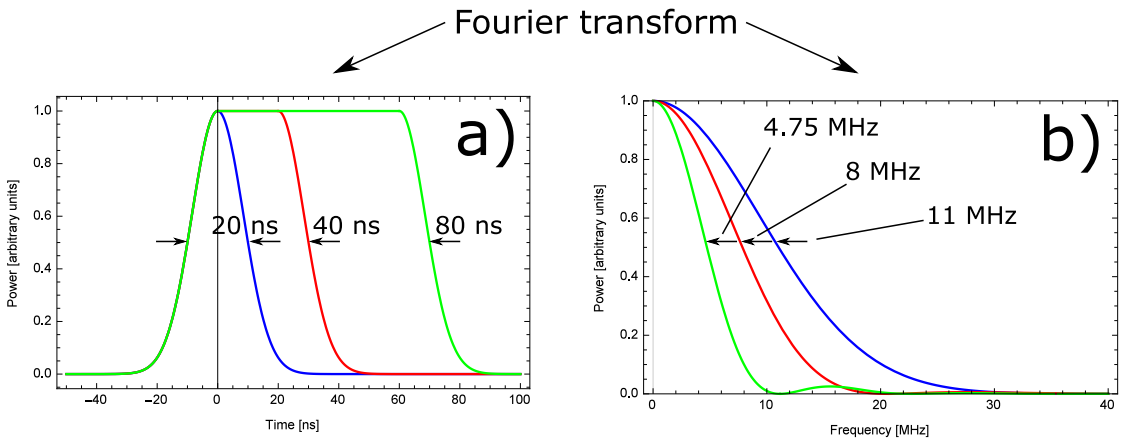

Figure 4.5: Plots illustrating the pulse-bandwidth effect for different pulse widths. The modeled pulse shape (part a) and corresponding Fourier transforms (part b), assuming 20 ns rise time. This is a realistic rise time for the AOMs and beam diameters used in these experiments. The blue curve represents the minimum pulse width (20 ns FWHM) that doesn't alter the shape of the rising and falling edges, known as a bandwidth-limited pulse. For part b), the HWHM of the blue, red and green curves are 11 MHz, 8 MHz and 4.75 MHz, respectively. Multiplying the FWHM in frequency- and time-space for each configuration gives the corresponding time-bandwidth product, which are 0.44, 0.63 and 0.76 for the blue, red and green curves, respectively. As expected, the minimum-width, Gaussian pulse has a time-bandwidth product of 0.44.

As shown in fig. (4.5), the maximum bandwidth that can be achieved by shortening the SW pulses is approximately 22 MHz FWHM, which is limited by the rise time of the AOMs used in this experiment. In the following calculations, the

non-relativistic Doppler effect is used, as the speed of the atoms with respect to the source of the light beams is only on the scale of a few m/s throughout the experiment. Therefore, for each component of the SW beam, the optical frequency seen by the atoms will be $\nu_{1,2} = (1 \pm gt/c)\nu_0$. We also assume a constant Doppler-broadened atomic lineshape. For an experiment that lasts 50 ms, the magnitude of the Doppler shift of the line center is then ~ 629 kHz.

To calculate the linewidth of the two-photon interaction, we consider the fact that the excited state that the atoms pass through during their interaction with the SW light field is a virtual state. As such, the linewidth of the interaction is due entirely to the overlap of the laser lineshape and the atomic Doppler lineshape. This assumes that there is no real atomic excited state in the vicinity of the virtual state, which would act as a potential loss mechanism from the AI scheme.

Adopting a simplified atomic structure which has only a single ground state and no real excited state, the lineshape of the two-photon interaction process is simply the multiplication of the respective lineshapes of the two traveling-wave components of the SW excitation with the Gaussian lineshape due to the Doppler width of the atoms. Here, it is important to realize that for non-zero temperature, there will exist some velocity class in the atomic velocity distribution that will be on two-photon resonance with the SW components. This velocity class will comprise a vanishingly-small fraction of the total atom number for narrow pulse bandwidths, or for very fast-moving atomic samples.

We assume that each beam has a CW Lorentzian linewidth of 1 MHz, which can be broadened according to the time-bandwidth product by shortening the pulse width. As such, we can expect a 20 ns pulse to have a bandwidth of around 22

MHz. In a frame stationary with respect to the laser source, the lineshape of the laser is given by:

$$L_{laser} = \frac{\gamma}{\pi((\nu - \nu_0)^2 + \gamma^2)}, \tag{4.1}$$

where γ is the laser bandwidth and ν_0 is the central wavelength of the laser.

The Gaussian lineshape of the atoms is:

$$L_{atom}(T) = \sqrt{\frac{mc^2}{2\pi k_B T \nu_0^2}} \exp\left(-\frac{mc^2(\nu - \nu_0)^2}{2k_B T \nu_0^2}\right). \tag{4.2}$$

For the complete lineshape of the two-photon process, we simply multiply the individual Doppler-shifted lineshapes of the two traveling-wave components of the SW excitation field together with the Gaussian lineshape of the atoms. The resulting lineshape function is:

$$L_2 = \sqrt{\frac{mc^2}{2\pi k_B T \nu_0^2}} \times \frac{\gamma^2}{\pi^2((\nu - \nu_0(1 - v/c))^2 + \gamma^2)((\nu - \nu_0(1 + v/c))^2 + \gamma^2)} \times \exp\left(-\frac{mc^2(\nu - \nu_0)^2}{2k_B T \nu_0^2}\right), \tag{4.3}$$

where we have allowed each of the traveling wave components to be Doppler-shifted in opposite directions, as would happen for a falling atom in a vertically-oriented SW excitation field. The integral of this lineshape is then proportional to the probability of the atom interacting with the Doppler-shifted SW field.

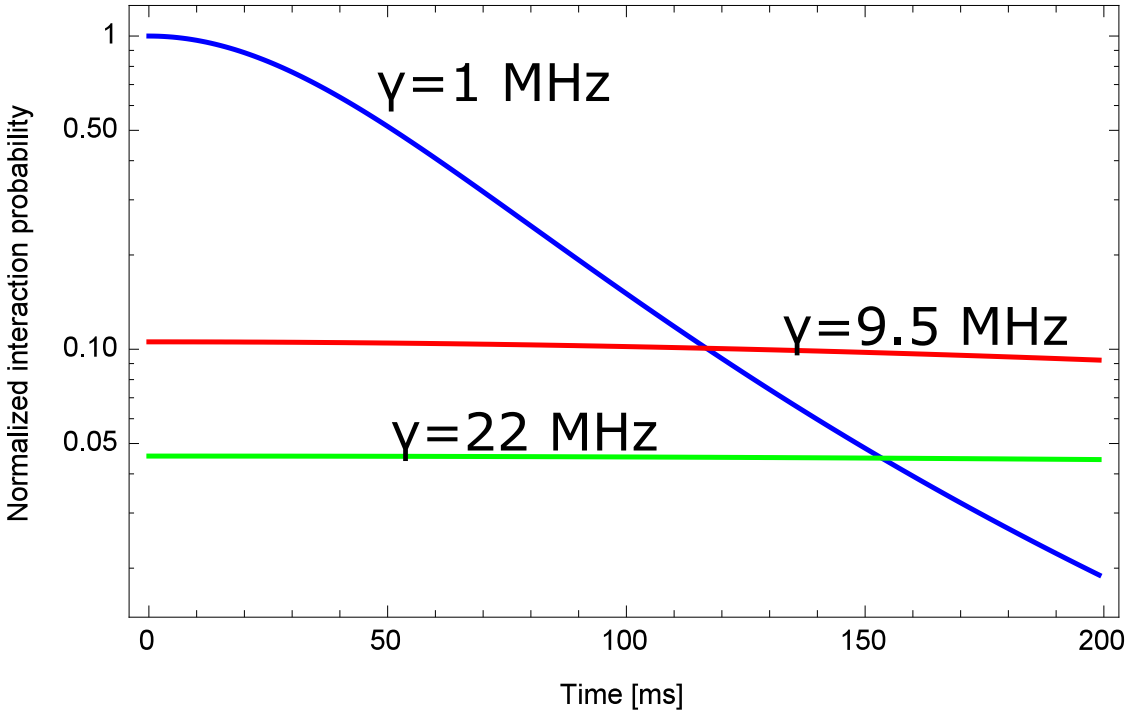

Figure 4.6: Interaction strength of the two-photon interaction for different pulse bandwidths plotted against experimental timescale. Here, all of the curves are the result of integrating eq. (4.3) and plotting the result against T, the elapsed time after the atoms begin falling. For each curve, the frequencies of the laser beams in the lab frame are not changed. The atom center of mass is assumed to be accelerating at 9.81 m/s^2 throughout, and the atomic temperature is 10 μK. Each curve is marked with its corresponding pulse bandwidth. The bandwidths of the blue, red and green curves correspond to $\sim$ 1 μs, 80 ns and 20 ns pulse FWMH values, respectively. The decay of the interaction probability for the blue curve is the result of the two SW components being Doppler-shifted outside of the thermal distribution of the atoms due to the COM velocity of the atoms. For the higher-bandwidth pulses, the laser bandwidth is sufficient to maintain a relatively constant interaction probability, but the overall probability is much lower due to the poor overlap of the atomic and laser linewidths.

As shown in fig. (4.6), the interaction strength decreases strongly as the pulses are shortened, due to the poor overlap of the atomic lineshape with the Doppler-shifted SW component lineshapes. As the spectral width of the pulses increases, the relative fraction of the total light intensity that is on-resonance with the atoms decreases. Therefore, there is a trade-off present in this technique, where one exchanges available pulse intensity for a more long-lived interaction. Chirping the SW beams completely removes the effect of the Doppler shift on the interaction probability, producing a constant interaction strength for all experimental timescales, limited only by atomic transit time effects.

One additional reason for chirping the SW beams is that unnecessarily-high SW bandwidths (due either to the pulse-bandwidth effect or to the laser linewidth) lead to increased uncertainty in the magnitude of $\mathbf{q} = 2\hbar k$, which can be problematic for high-precision measurements due to its effect on the uncertainty in the measured value of ω_q. For properly-chirped SW beams, the interaction strength becomes independent of the experimental timescale, allowing for longer experiments, up to the hard limit imposed by the transit time of the atoms due to their non-zero thermal velocity. The transit time can be extended by using larger-diameter excitation beams, colder atoms, or fountain techniques.

4.3 Proposed experimental layout

To counteract the Doppler effect, the SW components must be chirped at a rate of ± 12.58 MHz/ms, which is easily-achieved with off-the-shelf arbitrary waveform generators. The experimental layout required to produce chirped-beam SW excita-

tions is shown in fig. (4.7). This experiment uses a commercially available arbitrary waveform generator, marked as "Sweep generator" to produce an RF signal with a linearly-ramping frequency. This chirped RF is mixed with a carrier frequency near the operating frequency of the AOMs, producing sum and difference frequencies in the output of the mixer. This output is split, with one half being high-pass filtered, and the other being low-pass filtered. The high-passed signal will produce the "up-chirp", or the SW component whose frequency increases with time. The opposite is true for the low-passed signal. Each of these signals is input to a separate AOM, labeled as "K1 AOM" and "K2 AOM" in the figure. The diffracted beams from these two AOMs have orthogonal linear polarizations, so they are combined on a polarizing beam splitting cube. The combined K1-K2 beam is combined with a read-out beam from a separate AOM on a normal 50:50 beam splitter, before the beam is sent into a fibre. The output of the fibre is aligned vertically up through the atom cloud, before being retro-reflected from a flat mirror. This mirror produces a node in the SW potential at its surface, which determines the frame of reference for the experiment. For this reason, the position of this mirror will be monitored by a sensitive accelerometer, so that the motion of the experimental apparatus with respect to Earth's gravity field can measured. A post-correction can then be performed on the data to remove the effect of the motion of the apparatus.

The beam that reflects off the mirror double-passes through a $\lambda/4$ plate. This creates two super-imposed SW fields, one of which will be chirped the correct direction to stay on two-photon resonance with the atoms, the other being chirped the wrong way and quickly going off-resonance.

Simultaneous to the above, there will be a probe beam, generated in a different

light colour, which is used as an interferometric gauge of the length of the analytical arm of the AI. The variations in the length of this arm are corrected for by varying the phase of the driving RF with a fast feedback circuit and voltage-controlled phase shifter.

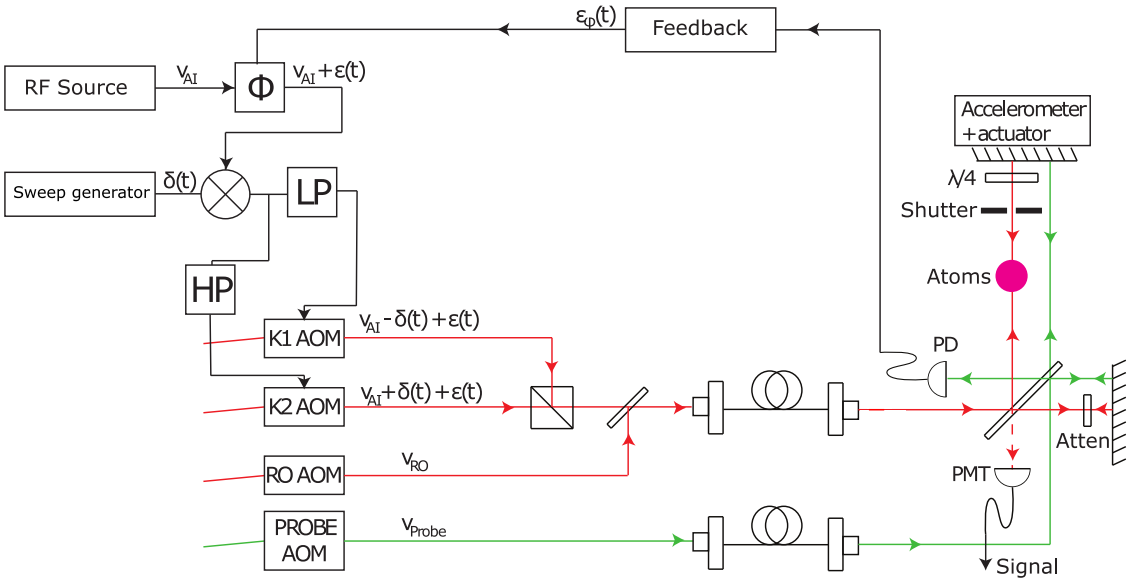

Figure 4.7: Proposed experimental layout for a gravity measurement with chirped excitation beams.

It is possible to circumvent the requirement for chirped AI beams somewhat for the case of short experimental timescales. For example, in an experiment that has poorly-controlled magnetic fields in the experimental volume, the maximum timescale for AI experiments may be as short as 10 ms. In this case, the added experimental complexity required for chirped excitation beams is not worth the required effort. To account for the Doppler shift of the AI beams during the experiment, the SW pulses are kept very short (< 1 μs), which broadens the pulse bandwidth of the beams enough to maintain a two-photon resonance condition for

at least some of the photons involved in the SW pulse. This has the notable drawback of requiring significantly-higher pulse energies for the excitation beams than the case of the chirped beams, for the same effective pulse area.

4.4 Magnetic field coil design

Magnetic fields and field gradients both contribute to the decoherence of the atoms during the experiment. The available timescale will be strongly limited by the presence of any ferromagnetic materials in the vicinity of the MOT, due to the pulsed trapping coil causing gradual magnetization of these materials over time. Additionally, some care must be taken to develop a set of magnetic field/gradient canceling coils which can be used to reduce this effect on the AI.

In previous gravimeter experiments in this research group, a stainless steel vacuum cell was used to generate and contain the MOT. Built around this cell was a set of three mutually-orthogonal cancellation coil systems. Each coil system was comprised of an overlapped Helmholtz-style and anti-Helmholtz style coil pair. In reality, only the coils oriented along the vertical axis were anywhere near the true Helmholtz spacing, which for round cross-section coils is simply the radius of the coils. The other two pairs of coils were separated by a distance much greater than the Helmholtz spacing, resulting in incomplete cancellation of the fields/gradients inside the cell. This is because the coils were arranged as the sides of a rectangular box, which makes it impossible for more than one pair of coils to be in a Helmholtz arrangement.

Further complicating matters was the fact that none of these coil pairs were

actually round in cross-section, which means that a true Helmholtz configuration was actually impossible. The best that one can do is a pseudo-Helmholtz arrangement, with square-cross-section coils separated by 0.55 times their side length. A comparison of the axial field at the centre of the coil pair for separations of half the side length (as in the old setup) and for a separation of 0.55 times the side length (as in the new setup) is shown in fig. (4.8). The proper coil separation results in a close approximation to the ideal situation, with a nearly flat field in the vicinity of the coil centre.

The final issue with the old apparatus was that the pulsed trapping coils needed for MOT generation, would slowly magnetize the chamber walls over time. The magnetization of the chamber would result in a gradual drift of the field/gradient value at the centre of the cell, changing the required coil currents needed for the cancellation coils to minimize the field/gradients at the MOT position, thereby optimizing the experimental timescale.

To combat these effects, a borosilicate glass cell was built, which removes all magnetizable materials from the vicinity of the MOT volume. The field/gradient cancellation coils were redesigned to be true pseudo-Helmholtz pairs, and the separation of the coil pairs was correctly set to allow for maximum zero field/gradient along the experimental axis. To achieve this, the coils had to be nested inside one-another, as shown in fig. (4.9). Each square coil was actually composed of two electrically-independent coils wound on top of one another, one to produce a field-canceling effect (Helmholtz arrangement), and the other to produce a gradient-canceling effect (anti-Helmholtz arrangement). Each sub-coil was composed of 76 turns of 12-gauge magnet wire, potted in place inside a plastic form using a

commercially-available two-part magnet potting epoxy (Electrowind, inc.).

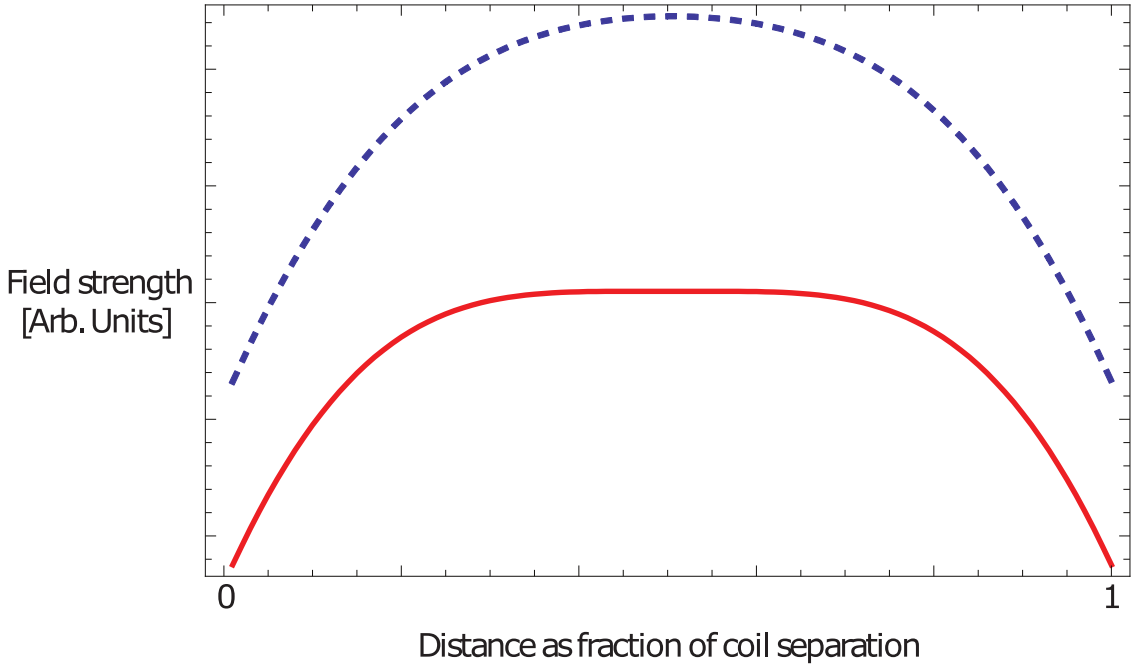

Figure 4.8: Comparison of the axial component of the field generated by a pair of square coils separated by 0.5 times their side length (dashed line), and by 0.55 times their side length (solid line). Having the coils closer together results in a higher field at the centre, but shortens the range over which the gradient is at an acceptable level.

The trapping coils in the glass-cell experiment were wound on cylindrical plastic forms with 60 mm $\times$ 20 mm cross section. Each coil had 319 turns of 12-gauge magnet wire, again potted in place with the same two-part epoxy as the cancellation coils.

4.5 Apparatus construction

The experiment was built on a standard optics table, equipped with laminar flow pneumatic isolation legs, as shown in fig. (4.10). On top of this was placed a sub-Hz

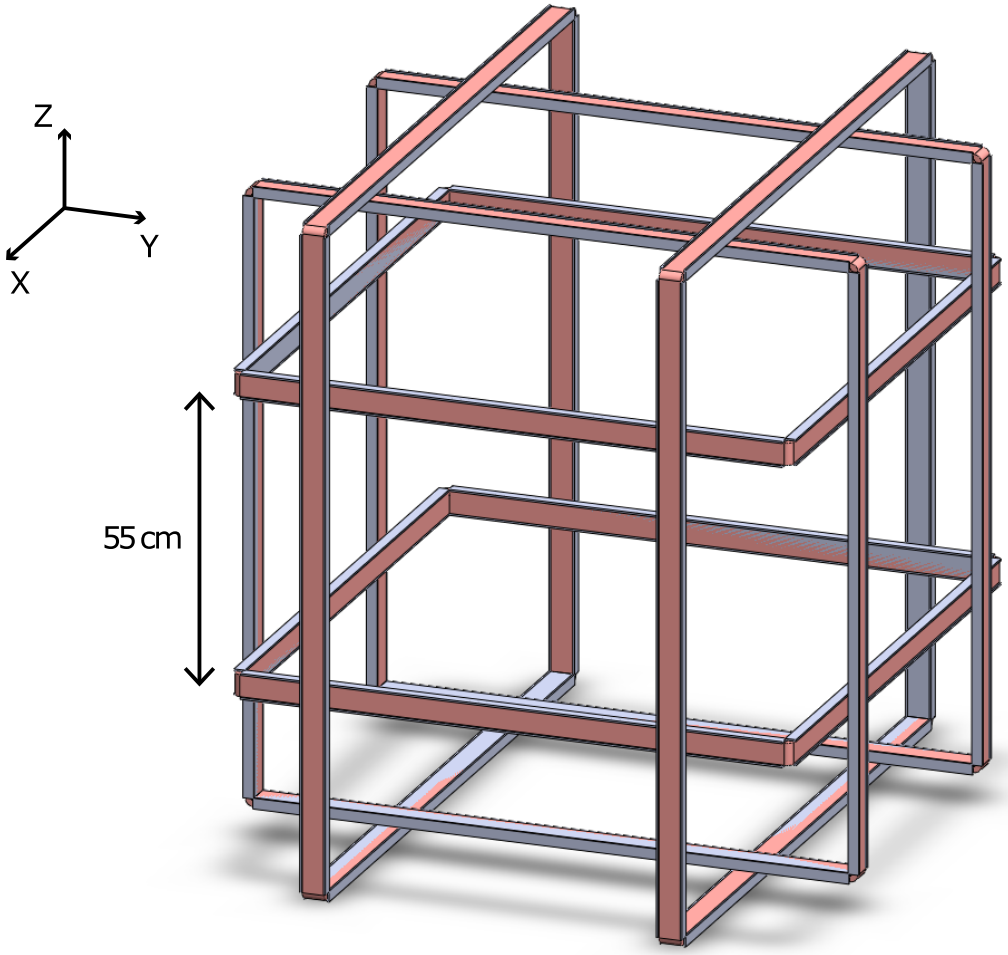

Figure 4.9: Plan view of the field/gradient cancellation coil box. The smallest set of coils, with their common axis along Z, have a side length of 1 m. The ones with their axes along X have a side length of 1.1 m, while the largest pair, with axes along Y, have a side length of 1.2 m.

tuned spring isolation platform, which was intended to damp vibrations in the 1-10 Hz region, which the optics table is still vulnerable to, due to its vertical resonance at 1 Hz. Due to load limitations on the tuned-spring platform, the field/gradient cancellation coils were mounted directly to the main optics table. The AI optics and vacuum chamber were placed on the sub-Hz platform. To prevent unwanted oscillation of the main optics table due to the high centre of mass (COM) of the system, ballast weights in the form of concrete slabs are meant to be placed on the support structure running under the table. The ion pump is also hung under the table, so as to minimize the effect of its fringing field on the experiments, and to help lower the COM further.

The apparatus layout shown in fig. (4.10) includes a series of elevated breadboards onto which the trapping optics were mounted. Due to load restrictions on the sub-Hz vibration platform, these breadboards had to be made from lightweight plastics. The resulting system had poor mechanical and thermal stability, requiring frequent realignment of the trapping optics. As a result, a simplified trapping optics layout was attempted, which used three integrated fiber output collimators as shown in fig. (4.1). The notable drawback of this approach is the retro-reflected trapping beams, which, due to the glass cell having no AR coating, causes a power loss in the retro-reflected beam of at least 16% for the normal-incidence beam pair due to Fresnel reflection coefficients. For the beams incident at 45°, the losses would be even worse. In addition, the effect of the glass cell's perturbation to the beams' spatial profiles has not been explored. For best results, it will probably be necessary to add three more fiber output collimators so as to achieve a full 6-beam trap. These effects will have to be explored in future work.

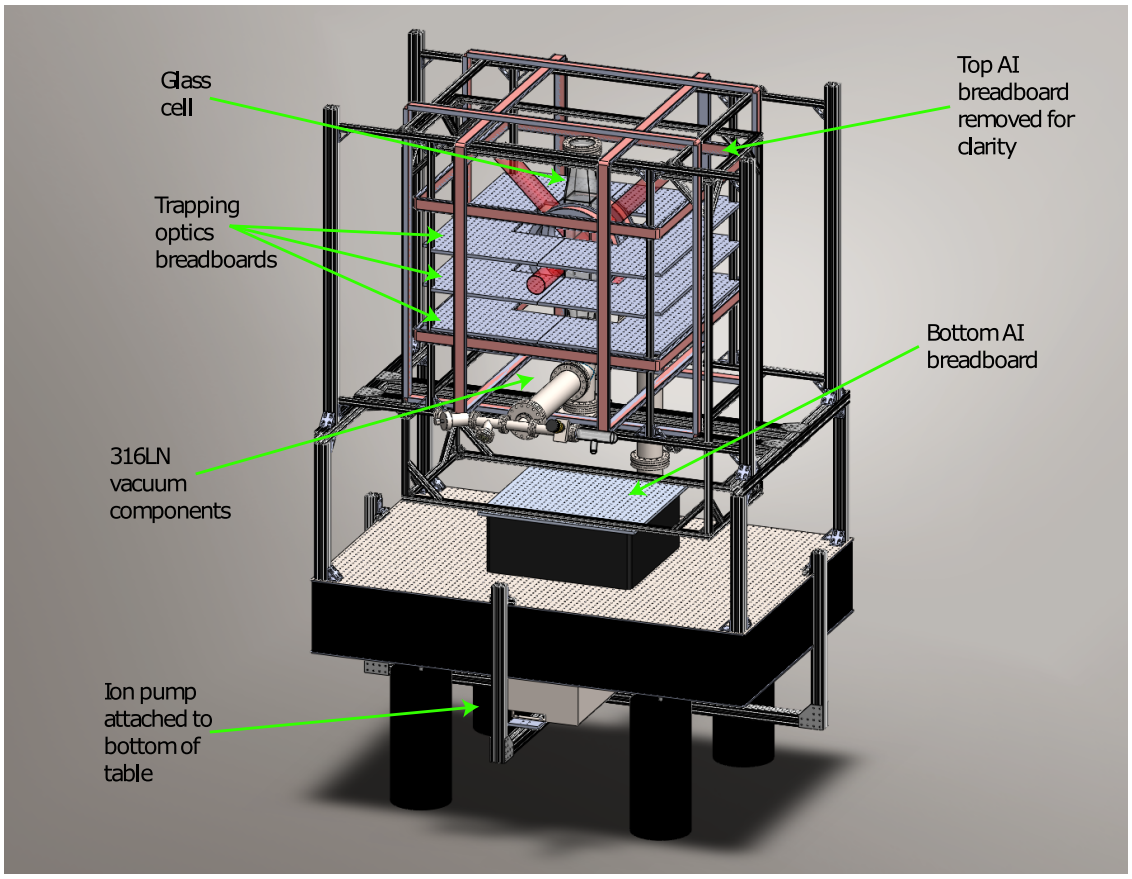

Figure 4.10: CAD model of the vacuum chamber with its associated support structure and vibration isolation setup.

5 Hermetically-sealed, interference filter stabilized, low cost, external cavity diode laser

High stability, narrow linewidth laser sources are an important tool for spectroscopy and inertial sensing applications. Ever since the development of the external cavity diode laser (ECDL), which started as a general development in techniques related to dye laser frequency stabilization [91, 92], low cost solid state lasers have become the light source of choice for many experimentalists and manufacturers. The ECDL's small form factor, coupled with the relative ease of producing ~ 1 MHz bandwidths make them very appealing. For certain applications, including long timescale atom-interferometric experiments, however, the stability of a commercially-available ECDL may not be high enough. In these cases, it may be necessary to isolate the laser from the influence of various external effects, such as vibrations, temperature changes and atmospheric pressure variation. The standard method to protect a laser from vibrations is to place it on an air cushion or tuned spring suspension. Temperature changes are usually avoided by actively stabilizing the temperature of the diode and external cavity with solid-state heat pumps and PID feedback loops. By placing the laser within a hermetically sealed chamber and evacuating the air, the last effect can be completely eliminated. Additionally,

removing the air from the laser cavity aids stability by eliminating any effect of pressure-related variation of refractive index on the laser's frequency. Presented here is a design for a laser that achieves sub-MHz linewidth while simultaneously achieving high frequency stability through the implementation of these and other techniques. This development project was crucial to the work being done in our lab, as we lost the use of our Ti:sapph laser in 2016.

5.1 External Cavity Diode Lasers

Laser diodes, operated without any external optical feedback or stabilization, are said to be free-running. A free-running laser diode has multiple uses in the physical sciences, including illumination of fluorescent markers in specialized microscopy, broadband pumping of laser cavities, and telecommunications. For these applications, the $\sim$ 1 nm bandwidth of a free-running diode does not pose a problem. However, such a light source does not allow for probing of atomic transitions, as the natural linewidths of atomic transitions ($\sim$ 1 MHz) are much smaller than the linewidths of free-running diodes ($\sim$ 500 GHz).

By taking advantage of their susceptibility to optical feedback, it is possible to greatly narrow the linewidth of a diode through use of external optics. The original implementation of what is now referred-to as the Littman-Metcalf configuration uses a diffraction grating and a mirror to couple the first-order diffracted beam back into the diode as shown in fig. (5.1). It is important to note that the gratings used in these lasers are of a particular type known as "blazed" diffraction gratings. These gratings are optimized to diffract strongly only at a single wavelength. Today,

by far the most common configuration is the so-called Littrow style ECDL [93–95], which has reduced complexity and higher output power compared to the Littman-Metcalf style. The main reasons for the continued use of the Littman-Metcalf configuration are that: a) the light makes two passes over the grating, resulting in a lower linewidth and b) the Littman-Metcalf design can achieve a much higher mode-hop-free tuning range than the Littrow configuration. The Littrow configuration is also shown in fig. (5.1). The difference between the configurations is in how each produces the feedback light. The Littman-Metcalf design uses a fixed grating and movable mirror, while the Littrow design uses a moving grating. The Littrow design as presented here has a major drawback in that the orientation of the output beam changes as the grating angle is varied. This is counteracted in practice by attaching a mirror to the diffraction grating, such that the 0th-order beam strikes the mirror after the grating. If the mirror is attached in such a way that its angle with respect to the grating is fixed, the output beam pointing will not change as the grating angle is varied. However, there will be a small offset shift of the beam with varying grating angle, which can be minimized by keeping the lever arms in the optics setup as small as possible.

In either implementation of the ECDL, an external optical cavity is established, in which only light of the specific frequency which satisfies eq. (5.1) and eq. (5.2) is allowed to resonate. All other frequencies of light are diffracted at slightly different angles from the grating, and are lost from the cavity. The slight increase in round-trip gain for the first-order beam results in the other modes of the free-running diode being suppressed, producing a reduction in the linewidth to the level of ~ 1 MHz.

There are drawbacks to the ECDL, however. Tuning of the output frequency

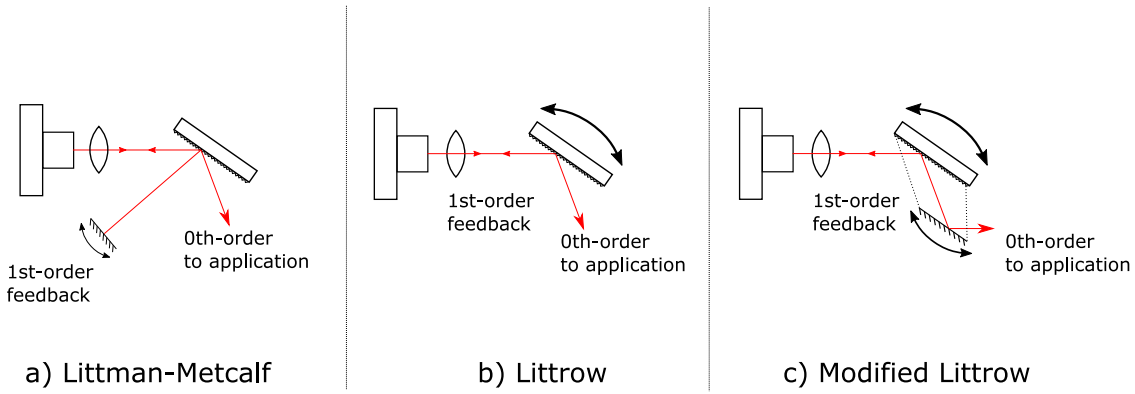

Figure 5.1: Schematic comparison of different optical configurations for grating-stabilized external cavity diode lasers. The Littman-Metcalf (a) and Littrow (b) style ECDLs differ in how they feed-back the diffracted light from the grating to the diode. In both cases, the 1st-order beam (usually, although higher-order beams can be used if needed) is created by a blazed diffraction grating. The modified Littrow (c) ECDL adds a mirror which is attached to the rotation mount of the grating such that it is held parallel to the grating surface. This has the advantage of converting angular shifts in the beam steering due to grating angle adjustment into small linear offsets of the beam position.

requires varying the angle between the diode output and the grating surface. Thus, the frequency of the light reflected into the first order beam is changed, and the output of the diode changes to match it due to feedback. The limitation here is that as the grating is tilted, there is often a coincident shift in the length of the external cavity. This is a result of the axis of rotation of the grating not necessarily passing through the point of reflection of the beam, as shown in fig. (5.2). For there to be stable lasing in the external cavity at wavelength λ_0, two conditions must be met. Firstly, the length of the cavity defined by the rear facet of the diode and the surface of the grating must satisfy:

$$l = n\lambda_0/2, \tag{5.1}$$

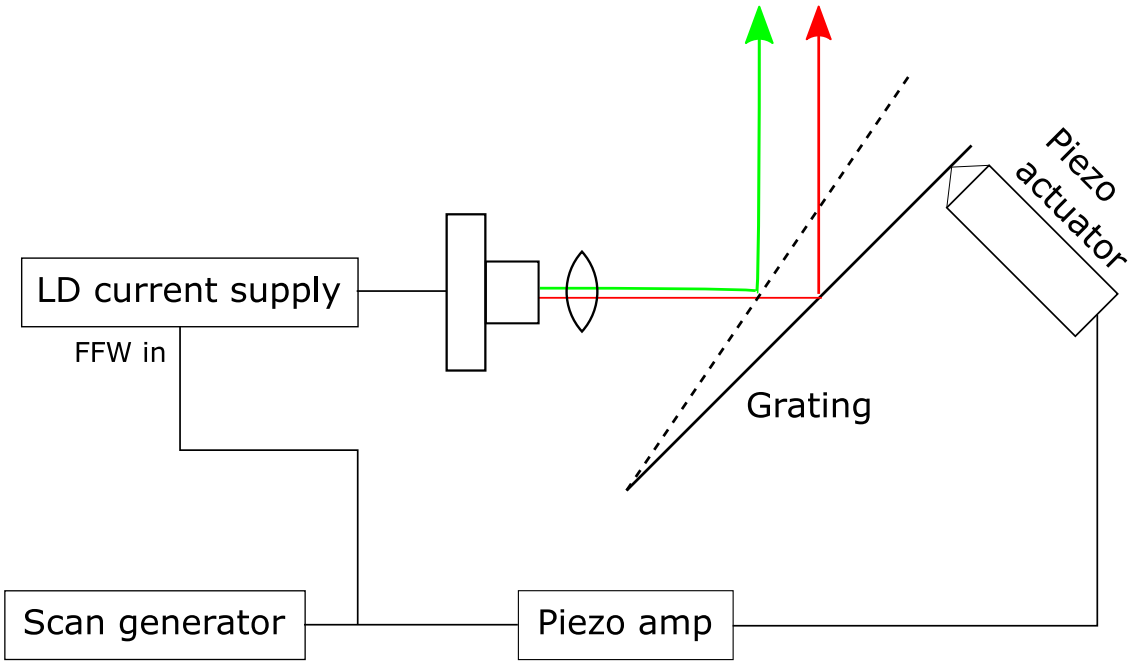

Figure 5.2: Demonstration of the need for feed-forward in a Littrow-style external cavity diode laser. The length of the external cavity (horizontal green and red lines) varies as the grating rotates. The amount of rotation here is greatly exaggerated for clarity, but a change of the external cavity length of $\lambda/2 \simeq 0.5$ μm can cause the laser to jump to a different mode. This effect is reduced by varying the LD current synchronously with the grating angle scan, which can significantly extend the mode-hop-free scan range of the ECDL.

where n is an integer. Secondly, the grating must be oriented at an angle:

$$\theta_g = \arcsin \lambda_0/2a, \tag{5.2}$$

where a is the grating line spacing, to ensure that the Littrow configuration supports this wavelength. As the grating is rotated, the values of both l and θ will vary, but not necessarily in a way that continues to support the conditions of eq. (5.1) and eq. (5.2). As soon as it is energetically favourable to do so, the laser will "mode hop" into an adjacent longitudinal mode, amounting to a change in n of ± 1. This will result in a shift in the output frequency of the laser by the free spectral range (FSR) of the external cavity, which, for a planar cavity is given by:

$$\nu_{FSR} = c/2l. \tag{5.3}$$

For a typical external cavity length of ~ 1 cm, the FSR is around 15 GHz. This amount of discontinuous shift in the output frequency of the laser is not acceptable when one wishes to examine features with spectral widths around 1 MHz.

To get around this issue, it is possible to include an electronic "feed forward" (FFW) circuit in the laser's control electronics. This circuit applies a current modulation to the laser diode's current, as shown in fig. (5.2), compensating partially for the changing length of the cavity by altering the center wavelength of the diode's output. By ramping the current of the diode, the principle mode of the external cavity remains the most energetically favourable mode for a larger range of grating angles. In this way, it is possible to scan the output frequency of the laser over several GHz without mode hops. This is called the mode hop-free tuning range of

the ECDL.

While it is possible to produce a low-cost Littrow-style ECDL (a laser very similar to the one described in ref [95] can be built from off-the-shelf optics for only a few hundred dollars), there is a significant fabrication challenge involved in reducing the need for a feed-forward and improving the stability of the laser. Modern, commercially fabricated ECDLs use more sophisticated grating mounting mechanics to ensure that the axis of rotation of the grating coincides with the point of reflection from the grating surface. Additionally, advanced systems will also incorporate piezo-based external cavity length adjustment as an additional feed-forward option, meaning that some systems are capable of < 1 MHz linewidth scanning across the diode's entire gain envelope (up to 100 nm [96, 97]).

5.2 The interference filter stabilized ECDL

Band-pass interference filters (IFs) offer an alternative to the diffraction grating for use in external cavities. There are substantial benefits from the IF-based approach, including the decoupling of the optics responsible for wavelength selection and establishing cavity dimensions. In the grating-based ECDL (GECDL), both of these roles are performed by the grating, whereas in the interference-filter-based ECDL (IFECDL), the IF performs the wavelength selection, while the cavity length is controlled by a piezo-actuated cat's eye retro-reflector.

Modern capabilities allow for the manufacture of filters whose pass-band widths are sub-nanometre. The pass-band of a filter can be varied over ~ 1 nm by adjusting the angle of incidence of the light to the filter [49]. The filter's pass band will move

towards lower wavelengths as the incident angle moves away from normal incidence. This is due to the simple geometric principle under which these filters operate. Each filter is a coated piece of optical glass. This coating is composed of layers of alternating index of refraction, whose thicknesses are set to produce constructive interference in the forward direction for a specific wavelength. Essentially, each layer acts as a low-finesse etalon, whose transmission can be modeled by:

$$T(R, n, d) = \frac{(1-R)^2}{1 - 2R\cos\frac{4\pi nd\cos\theta}{\lambda} + R^2}, \tag{5.4}$$

where n is the index of the layer in question, d is its thickness, θ is the angle of incidence and $R = (n_1 - n_2)^2/(n_1 + n_2)^2$ is the intensity reflection coefficient for the interface. This function neglects absorptive losses in the etalon. For typical dielectric materials with refractive indices ranging from around 1.5 - 2.0, $R < 15\%$, meaning the finesse of a single dielectric layer is usually less than 1. For such a low finesse cavity, the transmission function has a very broad spectrum. Therefore, it is necessary to use multiple layers to reduce the width of the transmission function.

By stacking many layers of alternating refractive index, the width of the transmission function peak that is nearest to the design wavelength can be reduced significantly. The width of a single peak in the transmission function is given by:

$$\lambda_{FWHM} = \frac{2\arcsin(1/\sqrt{F})\lambda^2}{2nd\cos\theta + \lambda}, \tag{5.5}$$

where $F = 4R/(1-R)^2$ is the finesse of a single layer etalon. As we add dielectric layers onto the filter, the finesse of this etalon grows, decreasing the width of its central transmission peak.

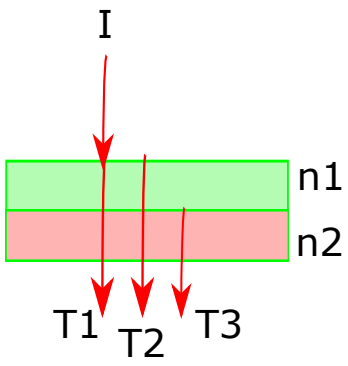

Figure 5.3: Simplified two layer interference filter coating. The layers have indices of refraction n_1 and n_2. At each interface, there is a reflected (omitted for clarity) and transmitted beam. As such, each layer acts as an etalon, transmitting strongly only the wavelength that satisfies $\lambda = d/n$, where d is the thickness of the layer and n is an integer.

The angular-dependence of eq. (5.4) makes it useful in the IF-stabilized diode laser. The position of the maximum in transmission is varied by the rotation of the filter with respect to the beam. The dependence is given by:

$$\lambda_t = \lambda_0 \sqrt{1 - \frac{\sin^2 \theta}{n_{eff}^2}}, \tag{5.6}$$

where λ_t is the transmitted wavelength, λ_0 is the normal-incidence centre wavelength of the IF, θ is the angle of incidence and n_{eff} is the effective index of the filter. This effective index varies depending on the polarization state of the input light, as well as the angle of incidence, but is usually around 2 [49]. The exact value of n_{eff} can be determined experimentally by varying θ and fitting the resulting transmitted wavelength to eq. (5.6).

The slight increase in gain at the centre of the pass band means that the diode mode that is closest to it will receive a slightly higher portion of the diode gain. This

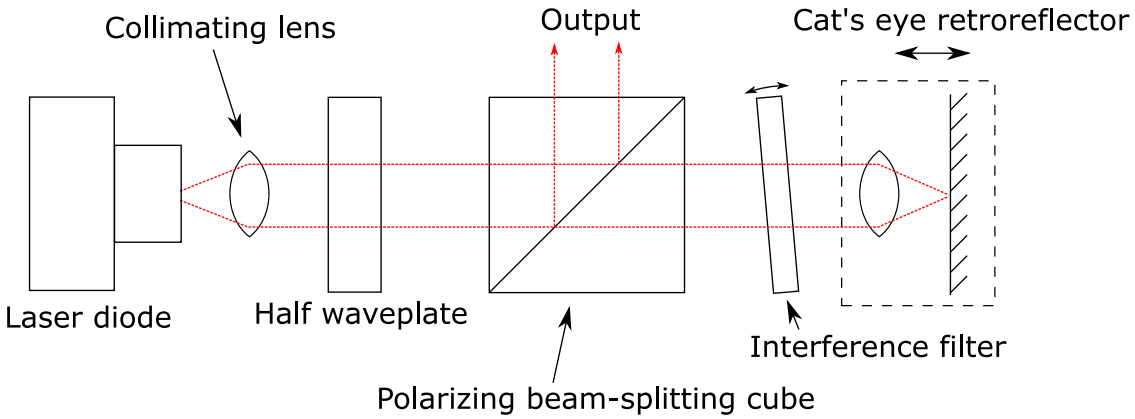

Figure 5.4: Schematic of the optics in our IFECDL. The light is emitted from the diode, through a collimating aspheric lens. The beam then passes through a variable beam splitter, composed of a half-wave plate, and a polarizing beam-splitting cube. This redirects the majority of the beam towards the output, leaving the remainder to be used for the feedback. The feedback light passes though an interference filter on a rotation mount, before being retro-reflected from a cat's eye retro-reflector on a piezo-driven translation mount. The retro-reflected feedback light makes another pass through IF, before having its original polarization state restored by its second pass through HWP.

will result in that mode "winning" the mode competition, becoming the dominant mode for the laser system, and receiving an exponentially-higher component of the available gain. This process occurs in exactly the same way as for the GECDL.

5.3 Design considerations

Our design for the IFECDL is shown in fig. (5.4). The variable beam splitter in the cavity makes use of high-quality polarization optics to ensure that: a) the polarization state of the feedback light after double-passing through the half-wave plate/polarizing beam splitter combination will match the output polarization state of the diode, and b) the light passing though the IF is possessed of only one linear polarization component, which eliminates the possibility of there being two feedback beams of opposite linear polarization which experience a different wavelength shift for a given position of the IF. As such, the waveplate used is an air-spaced pseudo-zero-order model with $< \lambda/300$ retardance accuracy, while the PBS has a 1000:1 extinction ratio. Switching to a narrow-band PBS would result in a better extinction ratio (as high as 3000:1), which is a potential future improvement for this design.

The prototypes manufactured so far are machined from aluminum 6061, a common structural alloy. Aluminum is used for its easy machining properties and its high thermal conductivity. Early versions of the laser had a welded stainless steel cover, which was chosen to allow welding of the vacuum half-nipple to the lid. By moving the vacuum fitting to the laser base, and swapping it for a threaded fitting, the lid was no longer required to be welded, and was replaced with a machined alu-

minum design. The mounts for the optical components were designed to simplify alignment by removing adjustable degrees of freedom.

The primary advantage of the IFECDL over the GECDL comes from the fixed alignment of the feedback beam from the retro-reflector, which uses a cat's eye arrangement to minimize the effect of mirror tilt during translation. The cat's eye arrangement uses a short focal length lens to focus the incoming beam onto the surface of the mirror. Upon reflection, the same lens re-collimates the beam. If the feedback beam is properly collimated, the system has reduced sensitivity to the axial position of the cat's eye. Additionally, any angular shift of the mirror in the retroreflector will produce only a lateral shift of the retro-reflected beam, rather than an angular misalignment, which for the extended cavity can result in complete loss of feedback. As such, once the alignment of the feedback light is established manually, the feedback condition is satisfied for small axial movements of the cat's eye around its nominal position.

The cat's eye retroreflector has to be aligned separately from the laser itself, using a collimated laser source to adjust the lens position for optimal performance. Unfortunately, the current revision of the prototype lacks lateral adjustment of the cat's eye retroreflector, reducing the quality of the optical feedback somewhat. This is because if the beam comes in off-axis to the cat's eye lens, it will be reflected back along the same direction, but with a lateral offset. This can result in a reduction of feedback coupling into the laser diode, harming the stability of the laser. Future versions of the laser will incorporate a 4-axis, X-Y-tip-tilt kinematic mount for the retroreflector, which should improve ease of alignment and stability.

5.4 Stability Measurements

To measure the stability of the laser, we employed an optical beat-note method, wherein the test laser's output beam is overlapped with the beam of an independent (reference) laser of the same wavelength. If the reference laser is of a known stability, it is possible to infer the stability of the test laser from the stability of the frequency of the beat note. This is because independent noise sources add in quadrature. Therefore, if the test laser and reference laser have respective noise levels of σ_{test} and σ_{ref} for a given record length, the resulting noise in the beat note frequency will be $\sqrt{\sigma_{ref}^2 + \sigma_{test}^2}$. Therefore, knowledge of the beat note noise can lead to knowledge of the laser's intrinsic noise.

If there isn't a reference laser of known wavelength stability available for testing, it is possible to produce a beat note from two identical lasers of unknown stability, in which case the total noise is: $\sqrt{2\sigma_{test}^2}$, assuming that each of these lasers has identical instability. This can also be done with a single laser, using a fibre-optic loop that has a length longer than the coherence length of the laser. The output of the laser is split, with half being sent through the fiber. The rest of the laser's output is combined with the output of the fiber, producing a beat note whose frequency stability is dependent on the frequency stability of the laser.

To measure the noise in the beat note, the Allan deviation method is used, as outlined in ch.6. In this case, the two signals being compared are the beat note frequency from the two lasers, where each has been locked to a different saturated-absorption peak in a Rb gas cell, resulting in a 31.7 MHz beat note, as shown in fig. (5.5). The beat note frequency stability contains the information about the

test lasers' stability.

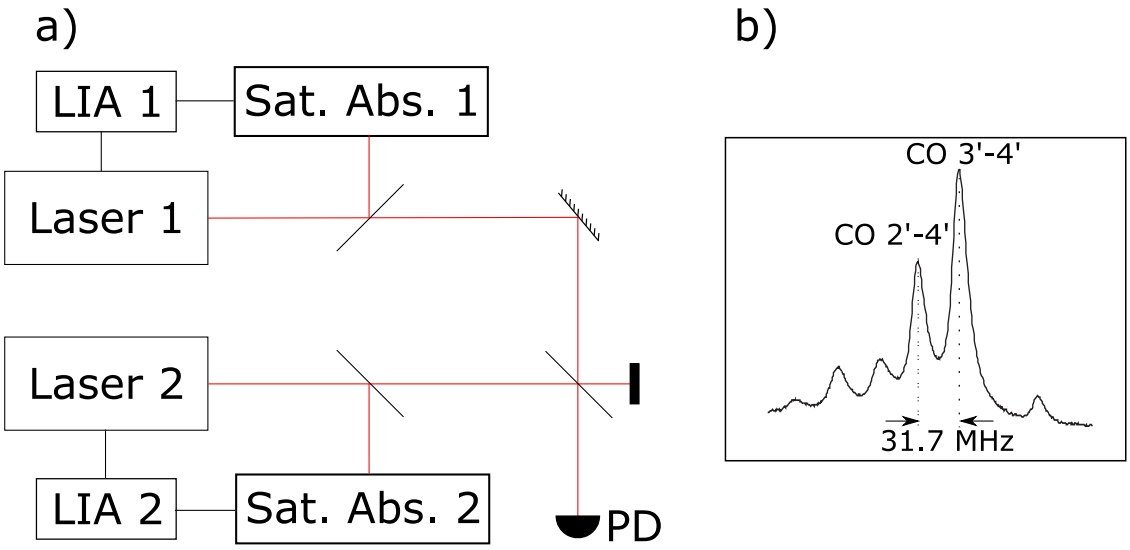

Figure 5.5: Optical layout for beatnote frequency stability measurements of lasers. Part a) shows a laser stability measurement setup. A fixed frequency beat note is generated by locking each laser to a separate feature in the Doppler-free saturated absorption spectrum of ^{85}Rb. This is done through the use of saturated absorption spectrometers (labeled as Sat. Abs. 1,2) and lock-in amplifiers (LIA 1,2). For this beat note, one laser is locked to the $F' = 3 - F' = 4$ crossover peak, and the other is locked to the $F' = 2 - F' = 4$ crossover peak, as shown in the saturated absorption spectrum in b). The difference in frequency between these two peaks is 31.7 MHz, which is also the frequency of the beat note.

5.5 Atmospheric Pressure Isolation

In early tests of a similar laser design that was built without the provision for Hermetic isolation of the laser cavity from atmosphere, it was shown that there was a correlation between the output wavelength of the laser and the ambient atmospheric pressure [52, 98]. This effect is ascribed to some combination of index changes in the cavity, and pressure deformation of laser cavity elements including

the laser diode can, which itself is a Hermetic volume containing some gas. To eliminate both effects at once, we redesigned the laser enclosure to be a robust vacuum chamber, capable of being evacuated to the mTorr level, as shown in fig. (5.6).

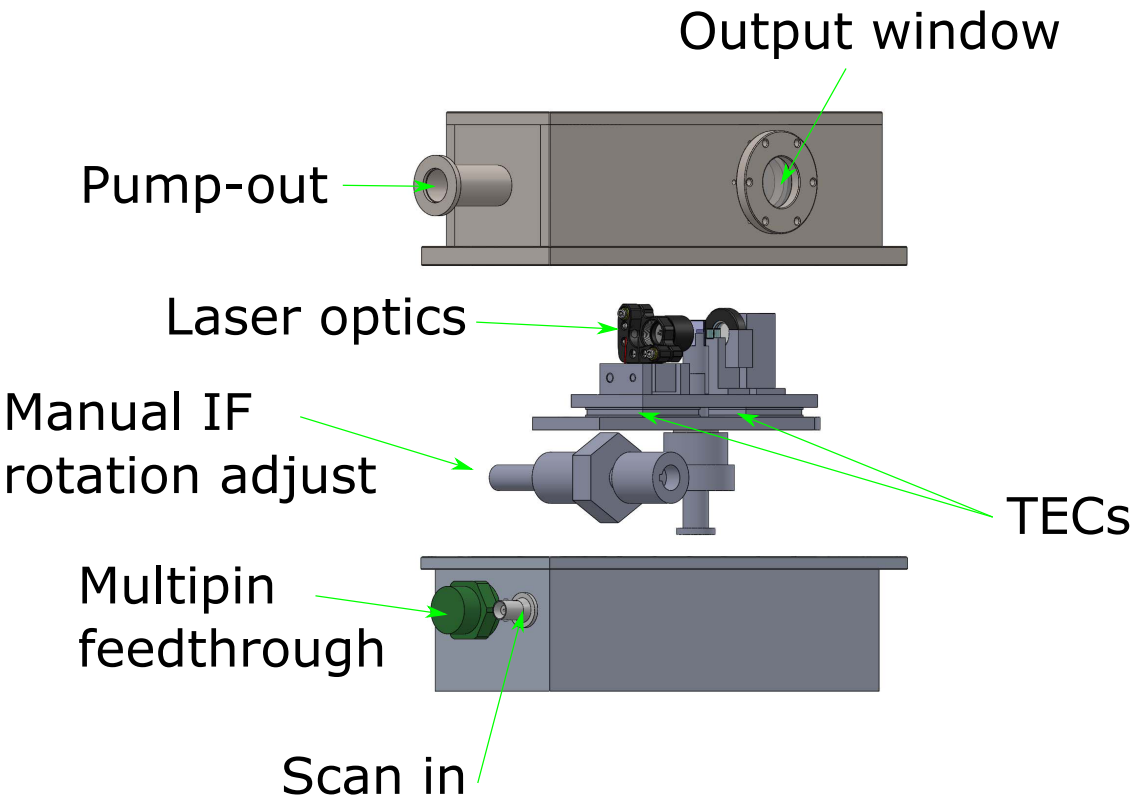

Figure 5.6: Exploded view of CAD model of laser. The top and bottom of the case are joined together with an elastomer o-ring seal, and all electrical/mechanical feedthroughs are vacuum rated. The output window is AR coated and wedged by 30 arcminutes, to eliminate unwanted optical feedback into the diode from the window. The IF angle inside the cavity is controlled by a simple gear linkage with 100:1 reduction ratio to improve manual tunability. The pump-out port is a standard vacuum flange (KF16) which easily interfaces with common vacuum pumps.

The effect of evacuating the laser cavity is shown in fig. (5.7), which is a plot

of the Allan deviation of the same laser head in both atmospheric pressure and pumped-out (1.6 torr) configurations [54]. The Allan deviation for the pumped-out laser is higher at all averaging times, which indicates that the laser is less stable (in the short-term) under these conditions. This effect is ascribed to mechanical stresses in the cavity that are introduced during the pump-out.

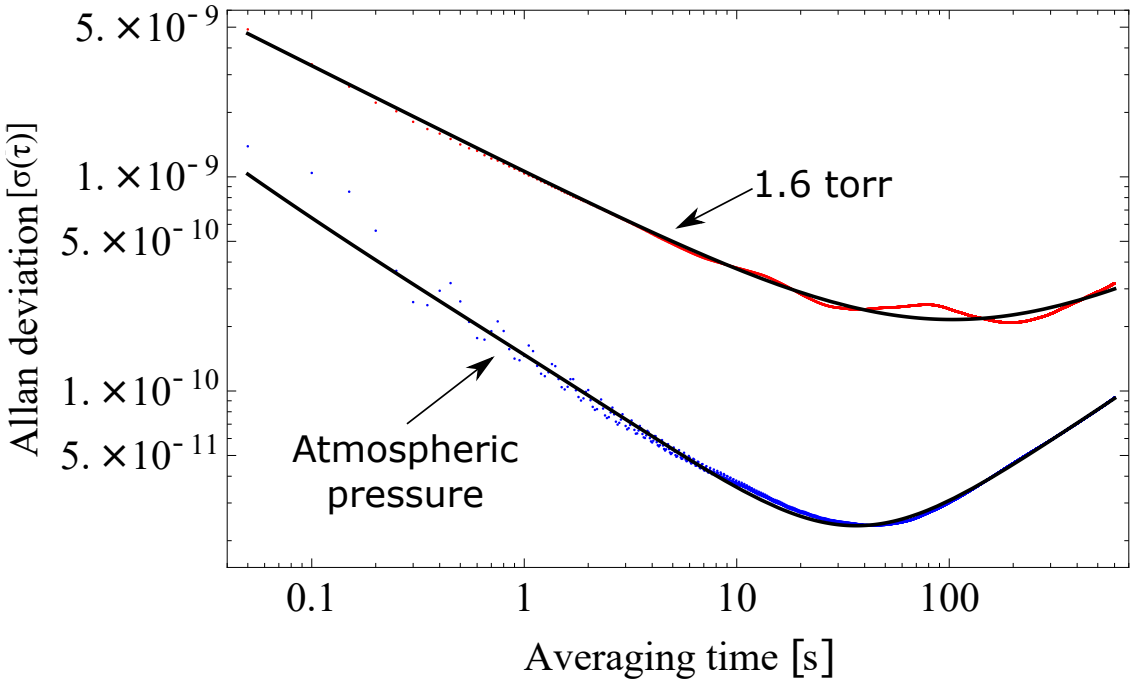

Figure 5.7: Allan deviation curves for the hermetically-sealed laser at atmospheric pressure and at 1.6 torr. These curves are produced by generating a beat note between two identically-constructed lasers, each locked to a neighboring atomic resonance. The fit function is given by eq. (5.7). For the blue curve, we have: $a = (-2.78 \pm 0.02) \times 10^{-11}$, $b = (1.87 \pm 0.05) \times 10^{-11}$, $c = (1.525 \pm 0.007) \times 10^{-10}$, $d = (4.07 \pm 0.02) \times 10^{-12}$ and $e = (2.37 \pm 0.05) \times 10^{-14}$. For the red curve, the coefficients are: $a = (1.11 \pm 0.06) \times 10^{-11}$, $b \simeq 0$, $c = (1.040 \pm 0.002) \times 10^{-9}$, $d = (1.005 \pm 0.002) \times 10^{-11}$, and $e \simeq 0$. We speculate that the near-zero terms in the fit to the red data are the result of the comparatively large amplitude of the non-zero terms in the fit.

The fit function used in fig. (5.7) is a simple power-law expression of the form:

$$\sigma(\tau) = a + b\tau^{-1} + c\tau^{-0.5} + d\tau^{0.5} + e\tau, \tag{5.7}$$

where a, b, c, d, and e represent the contributions from flicker frequency modulation, flicker phase modulation, white noise, random walk frequency modulation, and frequency drift, respectively [99].

The mechanical stability of the laser is sufficient for the initial testing phase, but improvements are needed to realize the ultimate performance of this design. In particular, the rotation mechanism for the IF is a weak point in the design, due to the backlash in the gear linkage. One easy way to improve this would be to pre-load the mechanism with a spring, or to use a flexure design [100]. Similarly, the cat's eye retroreflector is constructed with precision-ground rails supporting an optics carriage that runs on linear bearings. A flexure mechanism is the best way to improve robustness of this subsystem. Both of these improvements, among others, are planned for the next generation of the laser design. As it stands, this laser has comparable performance to typical commercial ECDL models, in terms of output power, linewidth and long-term frequency stability.

6 Radiofrequency synthesizer

6.1 Background

In frequency domain experiments involving atom interferometers [12, 55] or optical lattices [84, 101, 102], it is often necessary to generate two or more RF signals that differ in frequency. In these experiments, optical pulses with the appropriate central frequency are usually generated by driving a resonant accousto-optic modulator (AOM) with a RF source. The RF sources have to be designed to match the central frequency and tuning range of the AOM.

RF signals for such experiments must have high frequency stability. Additionally, the difference frequency must be tunable over a large range, in frequency steps small enough to satisfy the experimental requirements. Another desirable feature is that the device producing these signals should be easily interfaced with a computer.

To prepare for these types of experiments, we developed a frequency synthesizer, based on commercially-available components. There is a high level of flexibility inherent in this design, with key control parameters such as the number of outputs, tuning range, centre frequency and the operating bandwidth all being easily configurable at the outset. The synthesizer allows digital control of frequency while maintaining frequency stability with respect to a master reference oscillator. Addi-

tionally, a simple measurement and correction loop ensures that the outputs maintain a fixed phase relationship when the frequency is changed, a feature absent from most commercially-available frequency synthesizers. This implementation makes it possible to derive multiple outputs with these properties, making it a cost-effective solution for optical lattice experiments.

6.2 Experimental

6.2.1 Circuit Layout

We review a convenient setup for driving two AOMs at frequencies 250 MHz and 250 MHz + δ, where δ is a small frequency. The block diagram of the circuit is shown in fig. (6.1). In this circuit there are two phase-locked loops (PLLs) operating at 250 MHz and 238 MHz. These are fractional-divide-by-N PLLs, meaning that the reference and output frequencies are not necessarily related by an integer multiple. These devices are tunable only in steps of $\sim$ 0.2 MHz, necessitating additional electronics for fine-tuning.

To tune the output at 250 MHz + δ, we incorporate an arbitrary waveform generator (AWG), tunable in steps as small as 1 μHz. The AWG is a Stanford Research Systems (SRS) DS345. This device has a maximum output frequency of 30 MHz and is set to 12 MHz + δ. We chose 12 MHz so that when the AWG's output is mixed with the PLL output at 238 MHz, we obtain a sum frequency near 250 MHz. By filtering out the difference frequency, we obtain a signal with frequency 250 MHz + δ. Filtering is accomplished using a commercially-fabricated narrow-band notch filter from Filtronetics (FWHM 8 MHz) whose centre frequency

is 250 MHz. The RF mixer (Mini-Circuits ZP-5) used in this arrangement can accept input frequencies in the range of $0.2 - 500$ MHz.

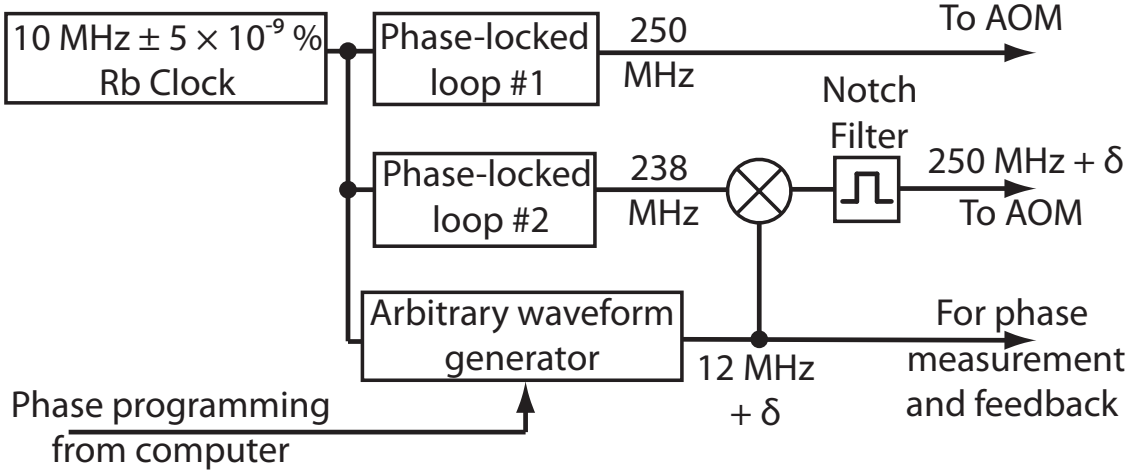

Figure 6.1: Block diagram for the synthesizer. Here, $\otimes$ is a RF mixer. A method for phase-stabilizing the outputs is explained in section (6.5).

Both of the PLLs and the AWG are referenced to a SRS PRS10 Rubidium clock operating at 10 MHz. This clock has excellent short-term stability, with a 1 s Allan deviation of 2×10^{-11}. Both the 250 MHz and 250 MHz + δ outputs are locked in phase to the stable Rb clock, and are therefore phase-locked to each other.

6.2.2 Phase-locked Loops

The PLLs used in this set-up are based on National Semiconductor's LMX2316 frequency synthesizer chip. Each chip is mounted on an evaluation board and controlled by a computer. An op-amp circuit was built to rescale and offset the output of the chip so that it could control the VCO, a Mini-Circuits ZOS-400, which is tunable from 200 MHz to 380 MHz. Figure (6.2) shows the block diagram of the PLLs shown in fig. (6.1).

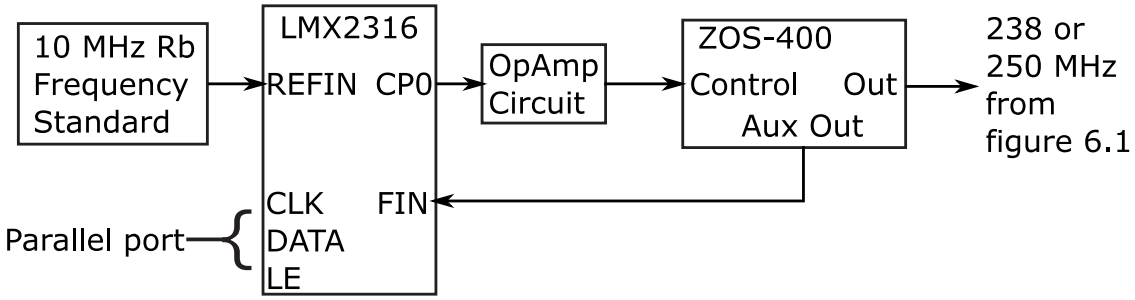

Figure 6.2: Block diagram of the PLLs used in this setup. The output labeled CP_0 provides the tuning voltage to the VCO, while the FIN input receives the feedback. The inputs CLK, DATA and LE allow programming by a computer via parallel communication. The phase detector in the LMX2316 chip is run at a sampling rate of 200 kHz, effectively re-locking the VCO to the Rb clock every 5 μs. The free-running VCO has a tuning sensitivity of 15 MHz/V. For a typical power supply ripple of 1 mV, this would result in a 15 kHz frequency fluctuation. As a result, a high-stability power supply or line filtering is needed for small frequency difference tuning.

6.3 Performance

6.3.1 Stability Measurement

In order to verify that the stability of the Rubidium clock is transferred to the outputs of the synthesizer, a reference oscillator of superior stability to the clock is required. Since no such oscillator was available, a comparison was made between the frequency synthesizer and a temperature-controlled quartz crystal oscillator of lower stability (test oscillator), which provides the internal timebase for the arbitrary waveform generator. This oscillator has an accuracy of $\pm$ 5 ppm, and ages by 5 ppm/year, with a 1 s Allan deviation of $\sim 1 \times 10^{-9}$. By confirming the Allan deviation of the crystal oscillator, we can state that the frequency synthesizer is at least as stable as the crystal oscillator.

To measure a stability of the test oscillator, its output is mixed with the output

of a much more stable reference oscillator. The resulting signal is then low-pass filtered to derive a beat note frequency. Since we measure frequency deviations with a frequency counter, it is necessary to introduce a small difference in frequency between the test and reference oscillators, so that the beat note has a non-zero nominal frequency.

In the measurement presented here, we used a PLL locked to the stable 10 MHz rubidium clock as the reference oscillator (1 s Allan deviation of 2×10^{-11}), and the test oscillator was an independent PLL locked to the temperature-controlled crystal oscillator. The test oscillator's PLL was run at 251 MHz, and the reference was operated at 250 MHz, giving a 1 MHz beat note. The beat note was input to a frequency counter, an Agilent 53131A. The counter runs at a rep rate of ~ 100 Hz, ensuring minimal dead time between measurements. A series of frequency measurements was recorded over several hours through a computer interface, and the Allan deviation [103, 104] of the test oscillator was calculated from this data.

6.3.2 Allan deviation

The Allan deviation is a statistical measure of an oscillator's stability as measured with respect to a more stable reference clock. For a given sampling time, τ, the Allan deviation is the average difference between the oscillator's frequency averaged over time τ, and its frequency averaged over an adjacent interval of time, τ. A low value of the Allan deviation implies high frequency stability over a given time interval [103].

The fractional frequency, $y_i^{\tau_0}$, is calculated from the series of frequency measurements using:

$$y_i^{\tau_0} = \frac{\nu_i^{\tau_0} - \nu_0}{\nu_0}, \tag{6.1}$$

where $\nu_i^{\tau_0}$ is the frequency at time t and ν_0 is the nominal frequency of the signal. Due to the discrete nature of the frequency sampling there exists a minimum time between measurements which we call τ_0. For this measurement, the value of τ_0 was 1.1 s. For a given value of $\tau = n\tau_0$, the Allan deviation, $\sigma_y(\tau)$, can be calculated using the following equation:

$$\sigma_y^2(\tau = n\tau_0) = \frac{1}{2(M - 2n + 1)} \sum_{k=1}^{M-2n+1} (\overline{y}_{k+n}^{\tau} - \overline{y}_k^{\tau})^2. \tag{6.2}$$

where M is the number of points in the data set. The Allan deviation is a function of the index, n, which corresponds to the time interval $\tau = n\tau_0$. For a given n, the $\overline{y}_k^{\tau}$ are calculated according to:

$$\overline{y}_k^{\tau} = \frac{1}{n} \sum_{i=k}^{k+n-1} y_i^{\tau_0}, \tag{6.3}$$

which is the mean of n measured fractional frequencies, $y_i^{\tau_0}$, separated by τ_0.

Figure (6.3) compares the measured Allan deviation of the test oscillator to the Allan deviation of the reference oscillator. The plot exhibits the expected shape for high performance oscillators. We note that the 1 s Allan deviation of the test oscillator as measured by our method matches with its reported specifications. The Allan deviation plot for the rubidium clock is taken from the manufacturer's literature [105].

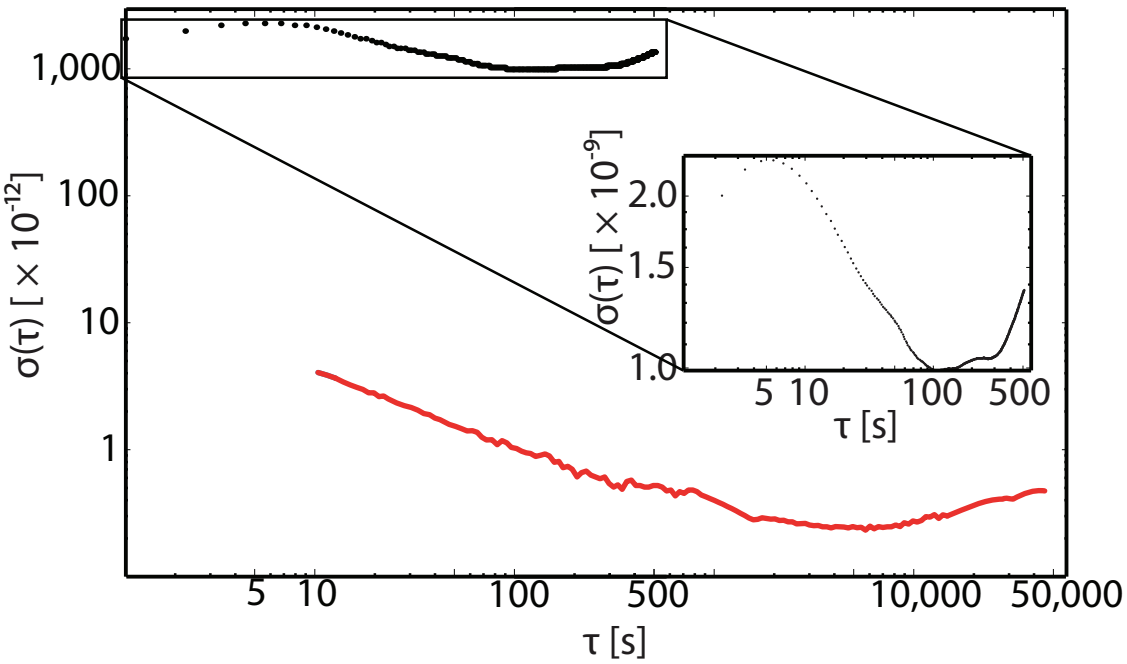

Figure 6.3: Comparison of Allan deviation of quartz oscillator (upper curve and inset) with Allan deviation of the PLL-based frequency synthesizer (lower curve).

6.4 Tunability

The tunability of the synthesizer was tested by beating the 250 MHz + δ output with the 250 MHz output. The results shown in fig. (6.4) demonstrate that the device is widely and finely tunable. The output of the synthesizer can be tuned in steps as fine as 1 μHz over a range of $\sim$ 8 MHz around 250 MHz.

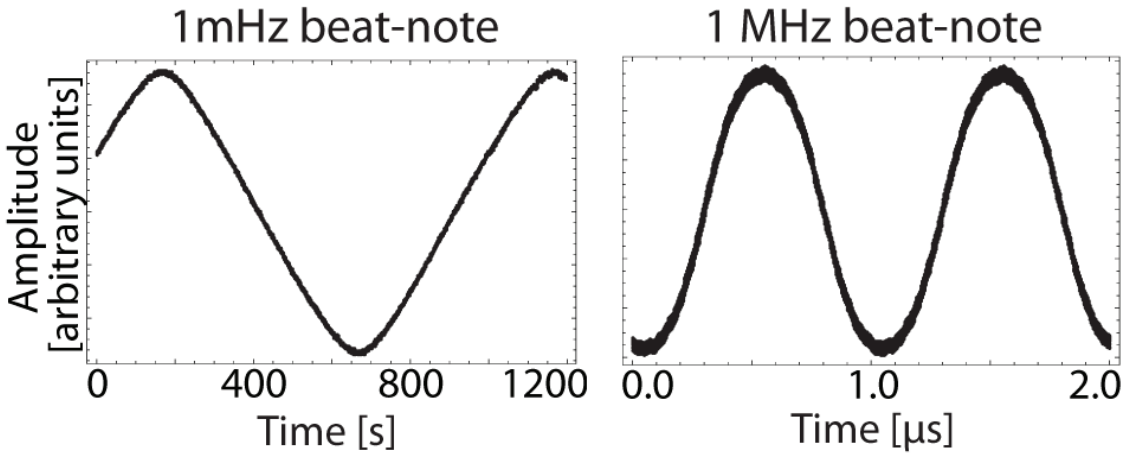

Figure 6.4: High and low frequency beat notes demonstrate tunability of the PLL-based frequency synthesizer over 9 orders of magnitude.

6.5 Phase Measurement and Correction

One of the primary drawbacks of using the AWG to tune the frequency is that every time the frequency of the AWG is changed, its output acquires a random phase. To ensure that the synthesizer's output has the same phase, the phase of the AWG has to be corrected every time the frequency is changed. The phase is measured by tapping-off a part of the 12 MHz + δ and sending it to the circuit shown in fig. (6.5). There, its phase is compared with the signal of interest, which, in this

case, would be a frequency-domain interferometric signal, which would oscillate at the same frequency. This is done by splitting the two signals, and introducing a $\pi/2$ phase shift into one component of the experimental signal with a delay line, which in this case is a coaxial cable (RG58) of length 4.12 m. The in-phase and in-quadrature components of the signal are then measured by passing the mixed signals through low-pass RF filters, and calculating the arctangent of the resulting DC levels. The resulting DC level is proportional to the phase difference between the two signals, and is sent to the arbitrary waveform generator to correct the difference to any arbitrary value to within 1.7×10^{-5} radians.

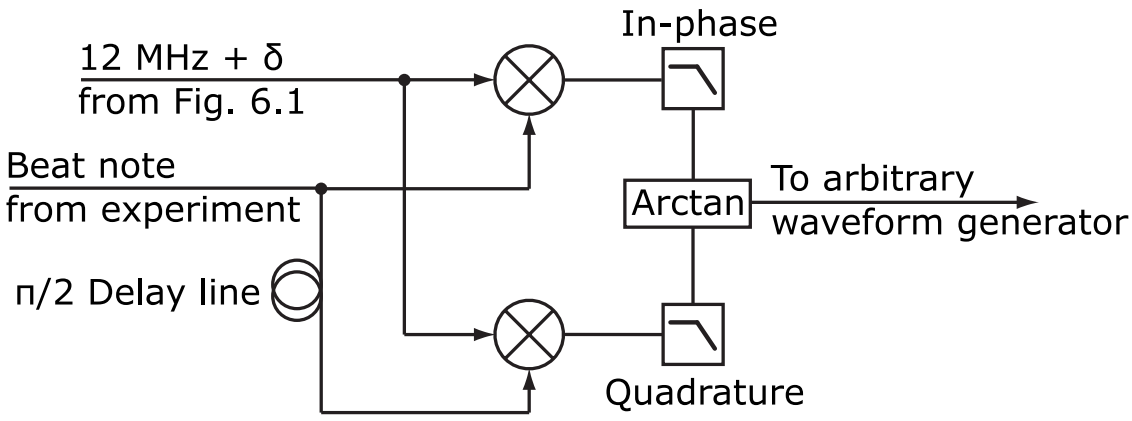

Figure 6.5: Schematic of a phase-correction circuit for the PLL-based frequency synthesizer. This circuit is designed to measure the phase of the AWG output at a time specified by an external trigger. The boxes marked "In-phase" and "Quadrature" are low-pass filters whose knee frequency is well below 12 MHz. The box marked "arctan" is a calculation of the inverse tangent of the signals, performed in software.

6.6 Frequency domain rotation measurement

One intended application of the RF synthesizer presented in this chapter was an atom-interferometric rotation measurement. Here, a frequency-domain approach was considered optimal, due to the constrained geometry of the experiment.

6.6.1 Frequency domain vs. time domain

Recall from ch.1, that the Sagnac phase is proportional to the area enclosed by the interferometer, multiplied by the component of the rotation vector that is perpendicular to the enclosed area. In other words, $\phi_S \propto \mathbf{A} \cdot \mathbf{\Omega}$, where ϕ_S is the Sagnac phase, $\mathbf{A}$ is the area vector of the interferometer, and $\mathbf{\Omega}$ is the rotation vector. To maximize the value of A, the atoms are launched along $\mathbf{x}$ with the AI beams applied along $\mathbf{y}$. This creates an enclosed area that generates a phase proportional to the rotation about $\mathbf{z}$. Since our apparatus was limited to a single AI beam due to constraints in the vacuum chamber, we chose to fix the timescale of the experiment so that the first SW pulse would be applied as the atoms entered the region illuminated by the beam, the second as the atoms cross the center of the beam, and the RO would be applied as the atoms reached the opposite side of the beam, as shown in fig. (6.6). For this reason, a frequency-domain experiment is ideal. In such an experiment a frequency difference is introduced between the TW components of the SW excitation pulses, transforming the pulses into pseudo-SW potentials. Additionally, the second pseudo-SW pulse is applied at a fixed value of T_{21} after the first. Due to the fixed pulse spacing, the geometry of the atomic trajectories is constant, allowing for maximum enclosed area for Sagnac phase generation. This

also has the benefit of reducing possible issues related to varying pseudo-SW intensity seen by the atoms as they transit the spatial profile of the SW beams. To generate the frequency-domain signal, the technique of [55] is used, which will be explained in more detail in section 6.6.3.

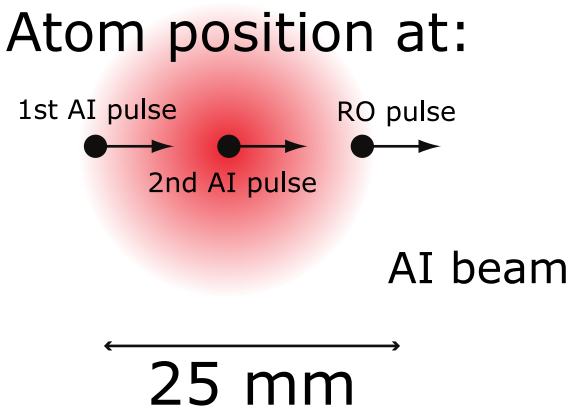

Figure 6.6: View along the axis of the AI beam for the proposed rotation experiment. The atoms are launched rightwards across the beam, which is pulsed on and off to generate a AI sequence, which separates the atoms into and out of the page.

Since the atoms travel across the spatial profile of the AI beam, at each pulse, they sample a different pulse intensity. It is for this reason that it is desirable to fix the time-separation of the AI pulses, so that at least the difference in pulse area for the first and second pulses can be fixed. It is then also possible to reduce the intensity of the second AI pulse, which occurs near the maximum of the spatial profile, by detuning an upstream AOM at the moment of this pulse. This technique was planned, but never tested.

6.6.2 Trap launch

The trap launch scheme we employed is shown in fig. (6.7). This technique produces a maximum trap launch speed of 60 cm/s. Note that this differs from a moving molasses trap launch as used in atomic fountain experiments in that the trap light is not ramped to launch the atoms. As such, we expect that the trap temperature would be greatly increased by this launch scheme. The actual extent of the temperature increase was never measured.

6.6.3 Ground state Ramsey fringe measurement of rotation

In order to measure the effect of rotation in a frequency domain experiment, the technique of ground-state Ramsey fringes can be used. This technique is so-named due to its conceptual similarity to the separated oscillatory field method developed by Ramsey [6]. In this technique, the ground-state population acquires an oscillatory phase, $\phi = 4\delta T$, which depends on the detuning between the TW components of the pseudo-SW potential, δ. A schematic of the steps needed to record the signal is shown in fig. (6.8). In brief, for each value of δ, an echo is recorded. The plot of the echo amplitude vs. δ yields a sinusoid, which is a Fourier component of the fringe pattern. Additional Fourier components are found by varying T_{21} and recording the corresponding echo amplitude vs. δ curves. The sum of these Fourier components produces the GSRF. According to [55], the central fringe of the GSRF will shift by an amount proportional to the Sagnac phase, and this is the effect that we had planned to measure.

The rotation measurement apparatus schematic is shown in fig. (6.9). The

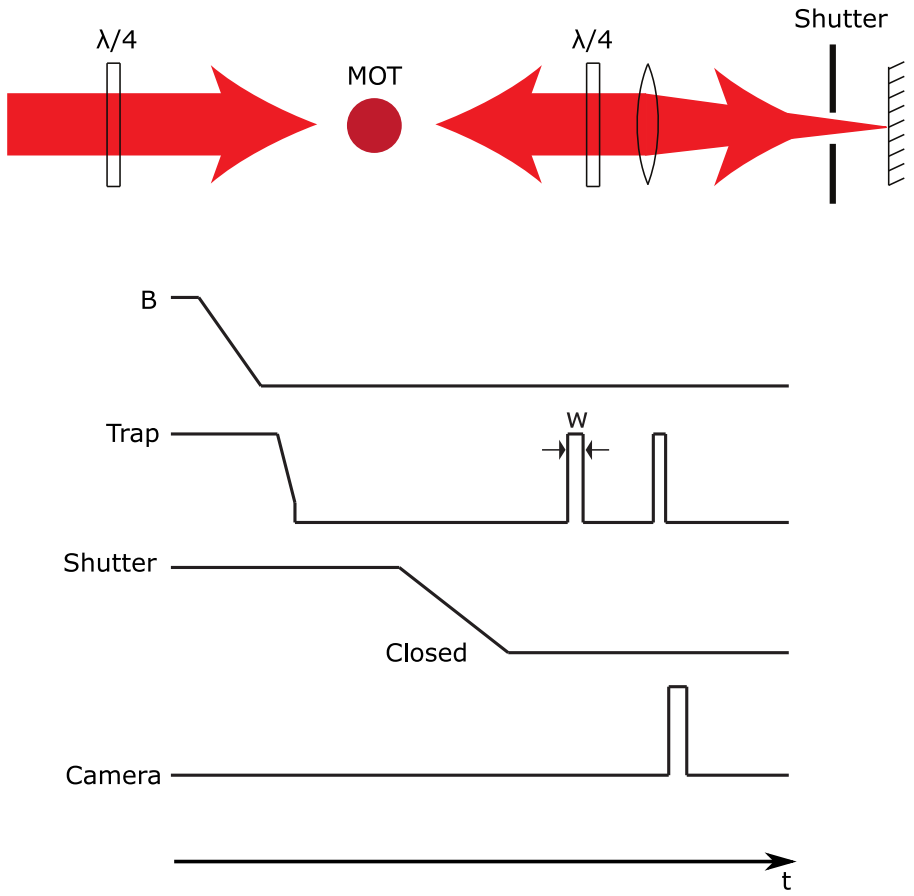

Figure 6.7: Timing diagram for the simplified trap launch scheme. The MOT is formed as usual, and the magnetic field is pulsed off. At this point, the trapping light is ramped and then pulsed off during the molasses cooling step. The shutter closes ~ 5 ms, which unbalances the radiation pressure force acting on the atoms when the trap light is pulsed on for the duration "w". The unbalanced pressure causes the atoms to travel to the right, and their position is measured by pulsing the trap light again and imaging the resulting fluorescence with a triggered camera. By varying the delay between the "launch" pulse and the"imaging" pulse, the velocity of the atoms can be measured.

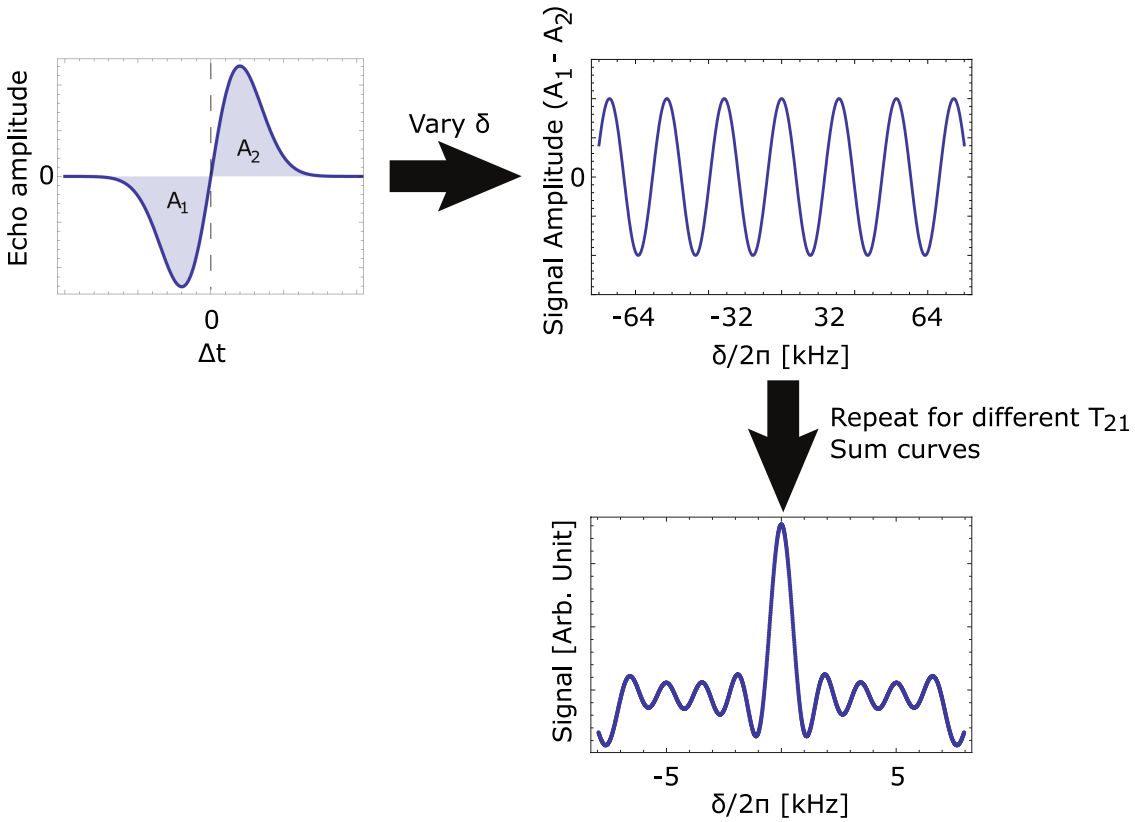

Figure 6.8: Process for generation of a ground-state Ramsey fringe from frequency-domain echo data. To generate the ground state Ramsey fringe, start at the top left, where a single echo signal is shown. By varying δ, the detuning between the two SW components of the AI beams, the relative sizes of the lobes of the echo oscillate sinusoidally, as shown in the second plot. The frequency and amplitude of this oscillation are both dependent on T_{21}. To generate the final plot, the value of T_{21} is varied and the curves summed. The characteristic shape of the GSRF shows a large central fringe, whose position for no force acting on the atoms is at $\delta = 0$. When a force (such as the Coriolis force) acts on the atoms, the position of the central fringe shifts, and the symmetry of the fringe pattern is broken.

crucial shortcoming of this experiment, and the reason that it was not able to produce a rotation measurement is the fact that the two halves of the optics, marked by the dashed boxes in the figure, had to be positionally stable with respect to one-another on the scale of the wavelength of the light in order for the Sagnac phase to not cancel. We quickly found that the vibration stability of this setup was so poor that the GSRF signal could not be created.

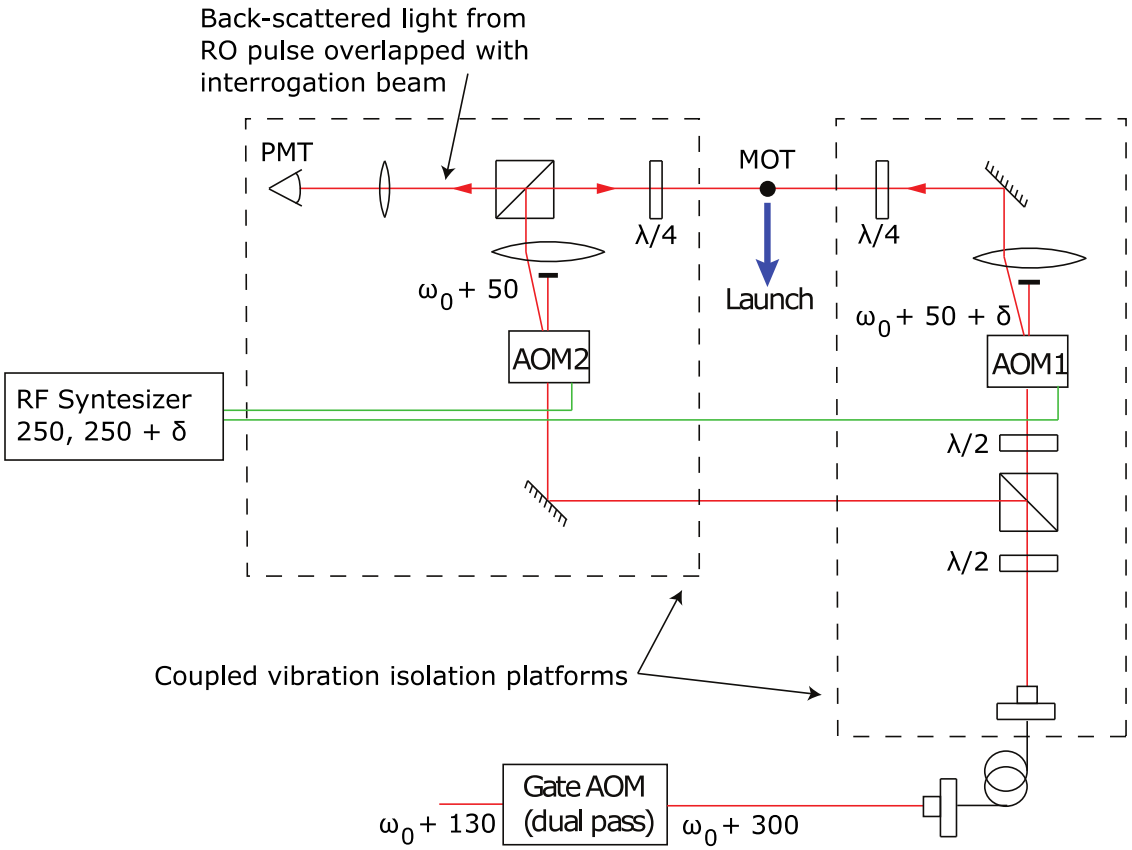

Figure 6.9: A schematic diagram of the rotation measurement experiment. The detuning of the two SW components from each other, δ, is created with our PLL-based frequency synthesizer. At the time of the read-out, a small amount of light is sent counter-propagating to the applied read-out, such that the back-scattered read-out light is co-propagating with this "interrogation" beam. Thus, a beat note of a desired frequency (10 MHz in our case) is formed at the detector.

7 Conclusions and future work

7.1 Conclusions

The work presented in this thesis lays the groundwork for a future precision measurement of gravitational acceleration using the grating echo technique. I have built the experimental cell and characterized the magnetic field coils. The interference filter-stabilized diode laser will be a valuable tool for the planned gravity and rotation experiments, based on the results of characterization. I have also shown that the predictions of Brynle Barrett's numerical simulations regarding density grating formation show good qualitative agreement with experimental results, indicating that lattice pre-loading of the atoms before the interferometric pulse sequence is a promising method for improving signal-to-noise. I have also laid the groundwork, in the form of building a frequency synthesizer and attempting the measurement myself, for a cold atom rotation measurement in the frequency domain, which can be completed by a future grad student.

7.2 Future work

7.2.1 Gravitational acceleration measurement

The AI optics need to be refined following the proposed experimental layout in ch.4. In conjunction with this effort, the electronics needed for frame motion detection and active stabilization need to be developed. Some improvement to the trapping optics are also possible, by moving from three retro-reflected MOT beams to 6 independent beams. This will reduce the intensity imbalances introduced by passage through the borosillicate glass.

7.2.2 Laser development

Whereas the current performance of the laser design is adequate, a further round of physical prototypes is warranted. I would suggest implementing a flexure-based cat's eye retroreflector mount to avoid some troubling mechanical issues with the linear bearing-based mount currently employed. Similarly, a piezo-driven rotation stage with preloading spring could eliminate the need for a mechanical feedthrough for the interference filter adjustment. Further work will also be needed to figure out why the laser's stability is worse for all averaging times when the cavity is evacuated.

7.2.3 Density modulation formation and lattice preloading

The findings of the simulations of AI signal amplitude for lattice pre-loaded atoms are tantalizing enough to warrant a dedicated experiment to measure this effect.

This exploration could be performed on the gravity apparatus described in this thesis, with only minor modifications to the optics.

7.2.4 AI rotation measurement

Similarly, the Sagnac gyroscope experiment could be implemented in the gravity apparatus, albeit with a more significant modification to the optics than for the lattice experiment. The measurement is a tricky one, due to the engineering challenge of building phase-stable platforms on either side of the experimental cell, and also due to the need for controlled and repeatable atom launch. Still, this could be a satisfying and interesting experiment.

A Finite element modeling of permanent magnet stack for Faraday isolators

In many applications of laser spectroscopy, back-reflections of the light output from a laser cavity must be avoided. While it is necessary for narrow linewidth operation of a diode laser to establish a small amount of optical feedback, strong optical feedback into the laser cavity will ultimately damage the diode laser oscillator, usually irreparably. It is even more important to avoid back-reflections of output light into tapered amplifiers, which are much less tolerant to optical feedback. While it is possible to avoid back reflections through careful selection of optics and ensuring that the beam is never at normal incidence to a reflective surface, there still remains the possibility that some feedback will occur during alignment or due to an unforeseen reflection. A common solution to this problem is a device called a Faraday isolator. The Faraday isolator can be thought of as an optical diode, allowing transmission of light in one direction, but strongly absorbing the counter-propagating beam. Since diode lasers and tapered amplifiers are keystone components of many atomic interferometry experiments, including ours, it is vital to utilize Faraday isolators.

A.1 Components of the device

All Faraday isolators use a total of three optical elements in series, as shown in fig. (A.1). In all cases, the middle element is called a Faraday rotator, and is responsible for the asymmetrical transmission properties of the device. The underlying physics of this element will be examined in the next section. The other two elements are either both polarizers, or both birefringent crystals, which determine the light's polarization. There are cases where the polarization of the light is well controlled, in which case, polarizers are used. In cases where the light polarization varies, or for high power beams, the birefringent crystals are preferred. For simplicity, and because they were the devices used in the experiments presented in this work, we will limit the discussion to the case of polarizer-based Faraday isolators.

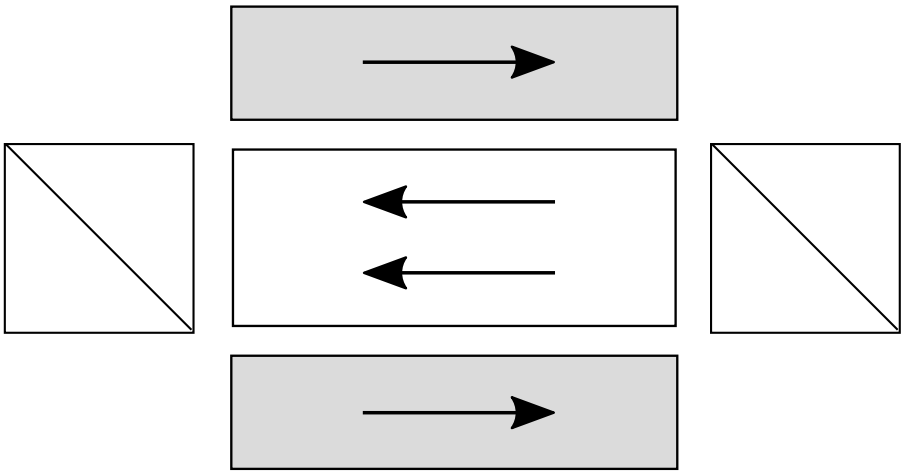

Figure A.1: Schematic diagram of a polarizer-based Faraday isolator. The arrows represent the average magnetic field within the body of the permanent magnet (shaded region) and the optical glass. This isolator would allow transmission of light in the right-to-left direction (along the field vectors in the optical glass), while blocking light from propagating in the opposite direction.

A.2 The Faraday effect

The physical effect that makes the Faraday isolator possible is named after Michael Faraday. He discovered that, for certain types of optically-transmissive glasses, applying a magnetic field to the glass would affect the polarization of the light passing through it. As it happens, the Faraday effect is not limited to glasses, and can be seen in many optical materials, including certain fluids.

The Faraday effect is a form of circular birefringence, wherein one circular polarization state propagates with a higher group velocity through the material than the other. Since linearly-polarized light can be described as a superposition of two circularly-polarized states, the result is a relative phase shift between the two circular components of the light, creating a rotation of the linear polarization vector. As this is a magneto-optical effect, it requires an externally-applied magnetic field. The "Verdet constant" of the material determines the amount of polarization rotation per unit length and per unit applied magnetic field. The total rotation of the polarization vector caused by a material with Verdet constant V and constant magnetic field $\vec{B}$ applied along the axis of light propagation is given by:

$$\theta = V\vec{B} \cdot \vec{d}, \tag{A.1}$$

where $\vec{d}$ is a vector with magnitude equal to the length of the material, pointing in the direction of light propagation.

We will first consider the case of a simple Faraday isolator, composed of a linear polarizer, followed by the Faraday rotator and then by a second linear polarizer. The polarizers can be of any common design, but, as the rejection ratio of the isolator

is simply the rejection ratio of one of the polarizers, a high quality polarizer such as a Glan-Taylor prism is best. The rotator consists of a cylindrical piece of optical glass, inserted into a permanent magnet stack. The glass must be chosen to have high transmissivity at the operating wavelength, and the required Verdet constant to generate a 45° rotation of the input polarization based on the axial field strength produced by the magnet. The magnet stack must be designed to have a central bore into which the optical glass is inserted, and along which it produces a very strong field.

Assume that the light input to the isolator from the laser is polarized along $\hat{x}$ (ie: $\vec{E}_0 = (E_\circ, 0, 0)$). If the isolator is properly aligned to the input beam, after the first polarizer (P_1), the polarization vector of the light will not have changed ($\vec{E}_1 = \vec{E}_0$). After passing through the optical glass in the Faraday rotator, the light will have had its polarization vector rotated by 45°, so that its polarization vector is now $\vec{E}_2 = 1/\sqrt{2}\,(E_\circ, E_\circ, 0)$. Again, for a properly-aligned isolator, the second polarizer (P_2) will do nothing to the light at this point, and the final result will be a beam with polarization $\vec{E}_3 = \vec{E}_2$, which is at 45° to the diode laser's original output. Due to absorption in the optics of the isolator, a few percent of the light's power is usually lost.

Now we consider a retro-reflected beam. As the optics downstream of the isolator can have unpredictable interactions with the light, the polarization of the retro beam cannot be easily predicted. We will consider the worst-case scenario, where the light reaches P_2 with the polarization $\vec{E}_4 = 1/\sqrt{2}\,(E_\circ, E_\circ, 0)$. That is, the polarization of the retro-reflected beam is aligned with the transmission axis of P_2. The polarization vector of the light transmitted through P_2 is then unchanged

(ie: $\vec{E}_5 = \vec{E}_4$). Because the light is now propagating in the direction opposite the magnetic field, there will be a negative sign in the final result of eq. (A.1). This means that the light's polarization vector after the Faraday rotator will be: $\vec{E}_6 = (0, E_\circ, 0)$, which will not transmit through P_1, as this polarizer's transmission axis is aligned to transmit light polarized along $\hat{x}$. If P_1 has a rejection ratio of $10^5 : 1$, a reasonable value for commercially-available polarizers, the amount of light that will be transmitted through P_1 will be very small indeed, considering that the retro-reflected beam is usually very weak to begin with.

Equation (A.1) makes several assumptions. The first is that the Verdet constant of the material is constant along its length. This is a reasonable assumption to make for a piece of manufactured optical glass or crystal, as the Verdet constant is a bulk property of the material and is usually well controlled. It is worth pointing out that the Verdet constant is wavelength and temperature dependent, so proper Faraday rotation requires both monochromatic light and a constant temperature [106]. The requirement for monochromaticity can be lessened by using glasses that have Verdet constants that are only weakly dependent on wavelength. The second assumption is that the magnetic field is also constant, both in time and in space. This is a much more difficult condition to achieve, requiring careful design of a permanent magnet stack to produce the required field. While it is in principle possible to produce a Faraday rotator with an electromagnet, it is rarely seen in practice. A permanent magnet stack requires no power supply or electrical connections.

Any radial variation in the magnetic field will produce a radial variation in the Faraday rotation angle. For any straight line path through the rotator, the rotation is more generally given by:

$$\theta(r) = \int_0^d \nu(r,z)B(r,z)dz. \tag{A.2}$$

This equation assumes cylindrical symmetry.

A.3 Analytical model

An analytical model of a homebuilt Faraday isolator [107] will be discussed to contrast with the more sophisticated finite element model to be presented in the next section. This analytical model represents each magnet in the isolator as a pair of magnetic dipoles. The disadvantage of this model is that it does not take into account the magnetic susceptibility of the magnet, which allows for the direction of the magnetization vector within the magnetic material to vary spatially. As such, this analytical model is of limited usefulness when designing or optimizing the magnet inside a Faraday isolator. An example of the fit of this model to experimental data is shown in fig. (A.2). For comparison, the same data fit to the finite element model of the same magnet is shown in fig. (A.3). As shown, the finite element model represents the experimental data much more closely than does the dipole model.

A.4 Finite element model of commercial Faraday isolator magnet stack

Finite element modeling of a permanent magnet stack in a commercial Faraday rotator was performed to assist a manufacturer in improving their design, both in

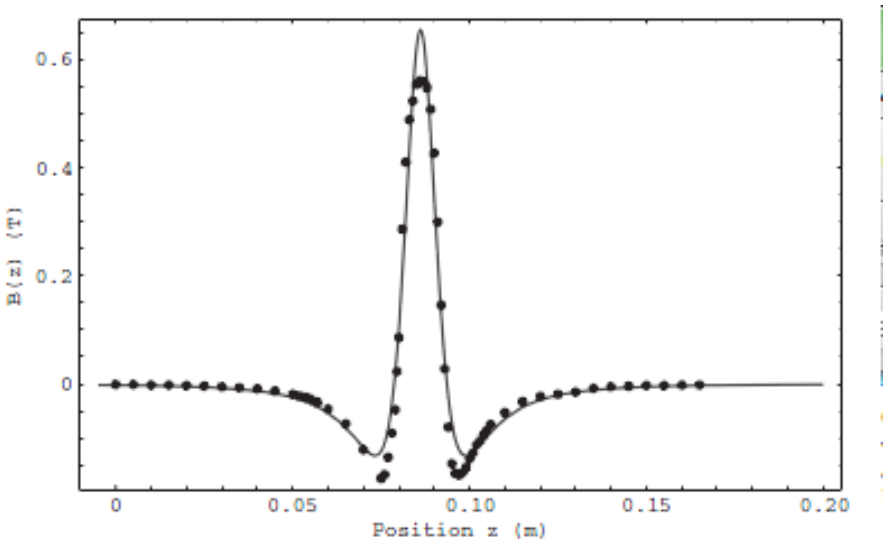

Figure A.2: Experimentally measured data from ref. [107], with the associated fit to the simple dipole-based model. This is the axial magnetic field of a single magnet, modeled as a dipole. Note the poor agreement between the model and the data near the extrema.

terms of magnetic field uniformity and in terms of temperature stability. An outline drawing of the magnet stack assembly is shown in fig. (A.4). The freely-available "Radia" package for Mathematica was used to perform the calculations. Models of the magnets were produced and segmented to allow relaxation effects to modify the magnetization vector within the bulk of the magnet.

In a finite-element model of magnetic materials, the material is initially assigned a bulk magnetization for the entire object. The object is then cut into small segments, each of which can have an independent magnetization vector. Since the magnetic field due to each segment can influence the magnetization of its neighbors, an iterative relaxation process has to occur, where the magnetization vector of each element is allowed to change due to the field of the other segments and the coercivity of the material. The coercivity of a material characterizes how easy it is

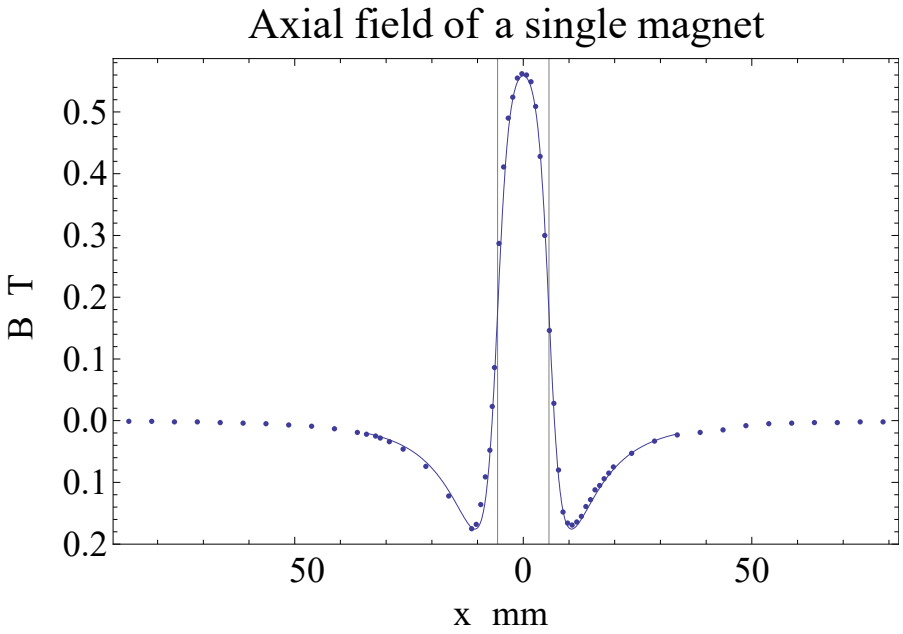

Figure A.3: The same experimental data as in fig. (A.2), this time overlaid with the result of a finite-element model of the magnet. The initial magnetization of the magnet was used as a free parameter to find the best fit to the existing data.

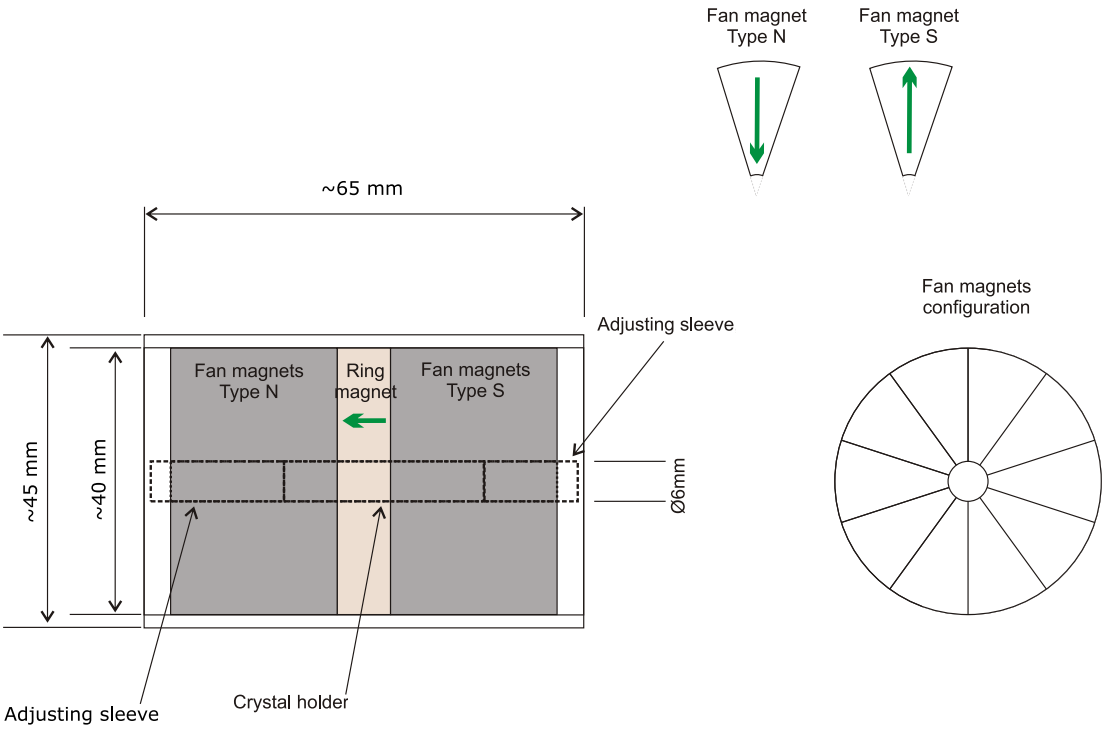

Figure A.4: Diagram of the magnet assembly from the manufacturer. Note that the green arrows represent the bulk magnetization vector of each element of the assembly.

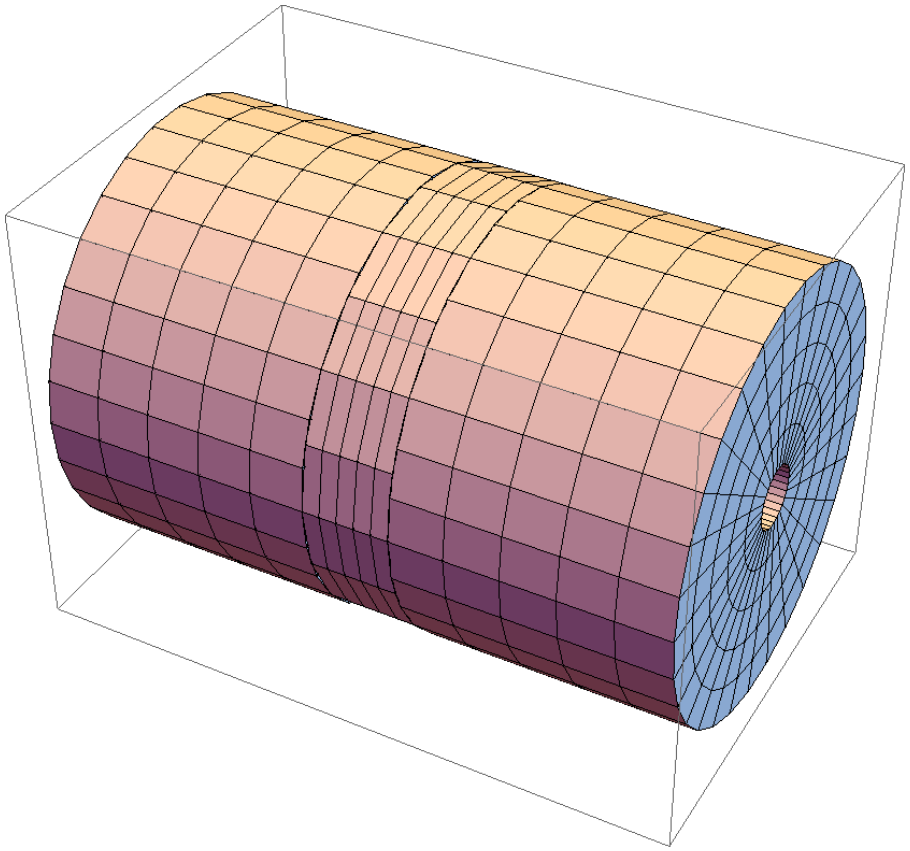

Figure A.5: Complete magnet stack as modeled in the finite element analysis, showing segmentation volumes. The central portion of the stack is a ring-shaped magnet with axial magnetization, while each of the ends is composed of a collection of radially-magnetized wedge pieces

to alter the material's magnetization with an external applied field. The relaxation is repeated until the change in the magnetization falls below a threshold, at which point the model is used to calculate external fields.

With the magnets segmented as shown, the initial, uniform magnetization was allowed to relax until the absolute change in the magnetization per step in the relaxation became less than 0.1 mT. To accomplish this, the software considers each element in turn, and recalculates its magnetization due the influence of the fields of the adjacent sections. The resulting variation of the axial field along both axial and radial directions are shown in fig. (A.6) and fig. (A.7), respectively. Note that we only concern ourselves with the axial field component, as it is the only field component that contributes to the Faraday effect.

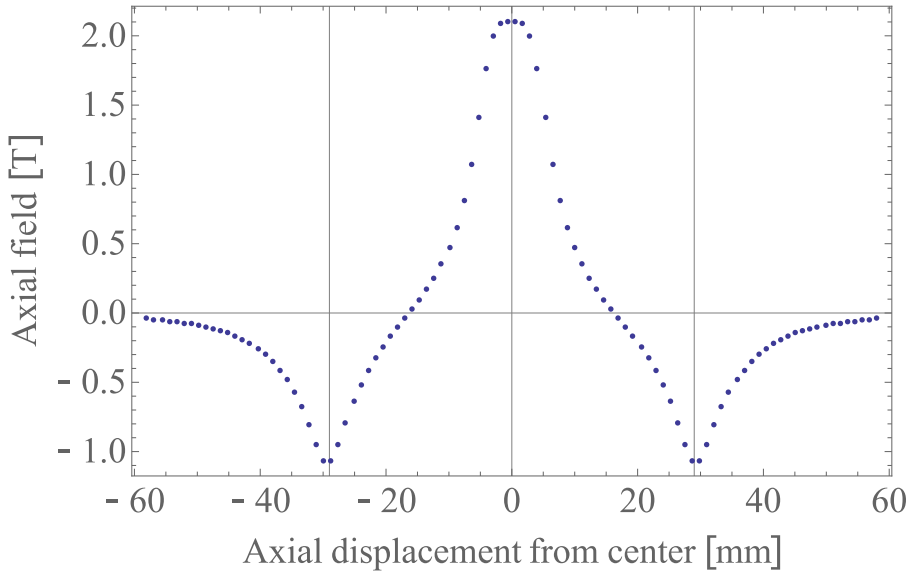

Figure A.6: Axial field along the axis of symmetry of the magnet stack after relaxation. The vertical gridlines at ±30 mm represent the ends of the magnet stack.

With a known magnetic field in the central bore of the magnet stack, the Faraday rotation for a given glass (with a known Verdet constant) was calculated. The result

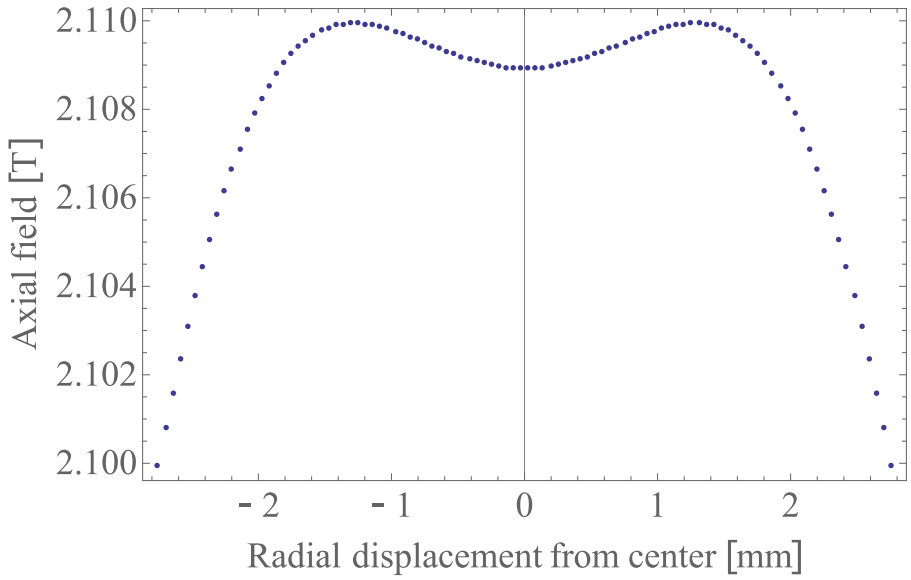

Figure A.7: Axial field at the axial center of the stack as a function of radius after relaxation. Note that the radius of the axial hole through the center of the stack is 2.9 mm

is plotted in fig. (A.8). Using this data, we were able to show that the design of the magnet would allow for tuning of the rotation angle in the vicinity of 45° by carefully moving the glass along the axis of the magnet assembly. This is important in a manufacturing environment, as the properties of the magnets, including physical dimensions and bulk magnetization, as well as the Verdet constant of the glass may all vary from batch to batch. Being able to tune the isolator to produce the required isolation is then necessary for production of a uniform product.

Finally, a comparison of the temperature stability of Faraday rotators made of two different magnetic materials was performed. Since the manufacturer was interested in high-power applications, the temperature stability was of primary importance. The two materials that were compared were NdFeB and SmCo. The first material, NdFeB, is a very common rare-earth magnetic material used widely

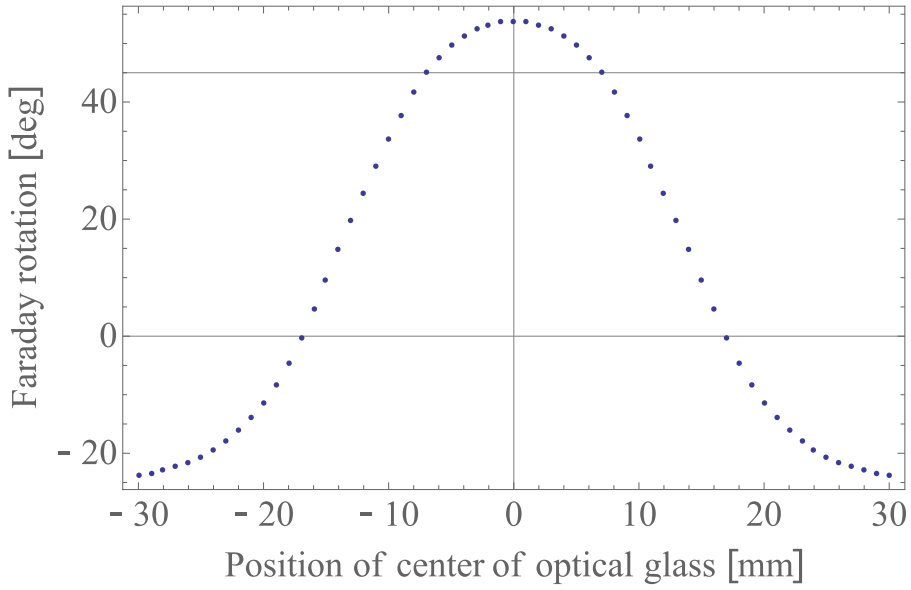

Figure A.8: Faraday rotation produced by a certain glass of known Verdet constant placed in the magnet assembly plotted against its axial displacement from the centre of the stack. A gridline is plotted at 45°, the desired rotation. Here, the optical glass is 30 mm long.

in industry. It has a residual induction of 1.26 T, and a reversible temperature coefficient of induction of -0.12 %/°C. Residual induction is a measure of the strength of the field that can be produced by a magnetic material. It is equal to the magnetic flux density at zero applied external field for a saturated magnet. The reversible temperature coefficient of induction determines the material's sensitivity to changes in temperature. The effect of this coefficient on the operation of the Faraday rotator is shown in fig. (A.9). By way of contrast, the SmCo material has a somewhat lower residual induction of 1.12 T, but a much lower reversible temperature coefficient of induction of -0.035 %/°C, making it better suited for applications demanding high temperature stability, as shown in fig. (A.10). The one major drawback of the SmCo material is its reduced coercivity value, meaning

that its magnetization vector would be more strongly affected by externally-applied fields.

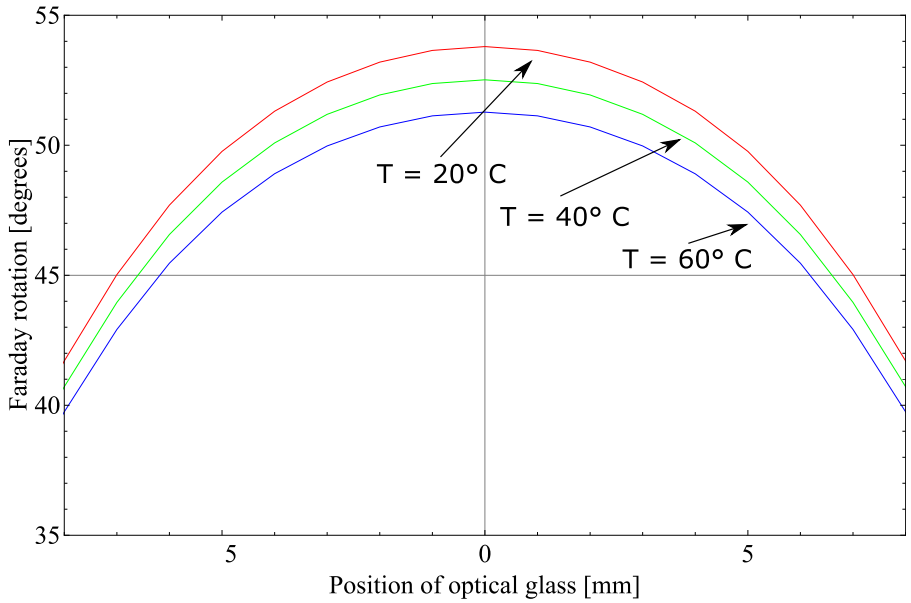

Figure A.9: Variation in Faraday rotation produced by the magnet stack composed of NdFeB for a selection of temperatures. As the temperature increases, the rotation changes for a given position of the optical glass. Here, the optical glass is 30 mm long.

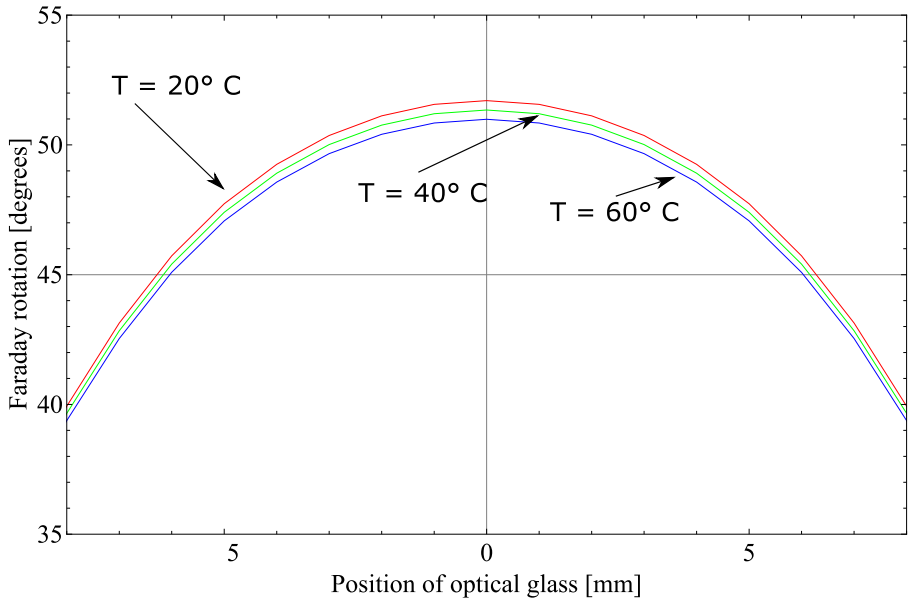

Figure A.10: Variation in Faraday rotation produced by the magnet stack composed of SmCo for a selection of temperatures. Note the reduced width of the band of curves, reflecting the higher stability of this material under changes in temperature. Here, the optical glass is 30 mm long.

B Bragg scattering simulations results by Brynle Barrett

The content of this appendix excerpts a longer document written by Brynle Barrett in January 2014 [86]. The document was meant to supplement my experimental results related to Bragg scattering and grating formation effects (presented in ch.3).

B.1 Modified grating-echo formalism

In order to discuss the topic of Bragg scattering from a periodic structure generated in a cloud of cold atoms, we must modify the theory of the grating formation so as to take into account the limited spatial extent of the atomic probability distributions. The standard derivation of the echo signal uses atomic plane waves, which have infinite extent. Therefore, this signal derivation differs somewhat from the one presented in ch.2.

We start with the initial wavefunction:

$$\phi_0(p) = \left(\frac{1}{\sqrt{\pi}\sigma_p}\right)^{1/2} \exp\left[-\frac{1}{2}\left(\frac{p-p_0}{\sigma_p}\right)^2\right], \tag{B.1}$$

where p_0 is the average momentum of the atoms (usually zero) and $\sigma_p = \sqrt{2Mk_BT}$

is the $1/e$ width of the momentum distribution. Here, we have made the initial wavefuntion equal to the square-root of the momentum-space probability distribution, which ensures that $\int |\phi_0(p)|^2 dp = \int |\psi_0(z)|^2 dz = 1$.

The initial wavefunction in position space is then:

$$\psi_0(z) = \frac{1}{\sqrt{2\pi\hbar}} \int_{-\infty}^{\infty} \phi_0(p) e^{ipz/\hbar} dp = \left(\frac{1}{\sqrt{\pi}\sigma_z}\right)^{1/2} \exp\left[-\frac{1}{2}\left(\frac{z}{\sigma_z}\right)^2\right] e^{ip_0 z/\hbar}, \quad \text{(B.2)}$$

where $\sigma_z = \hbar/\sigma_p$. The position-space and momentum-space probability distributions for both $T = 10$ μK and $T = 10$ nK are shown in fig. (B.1).

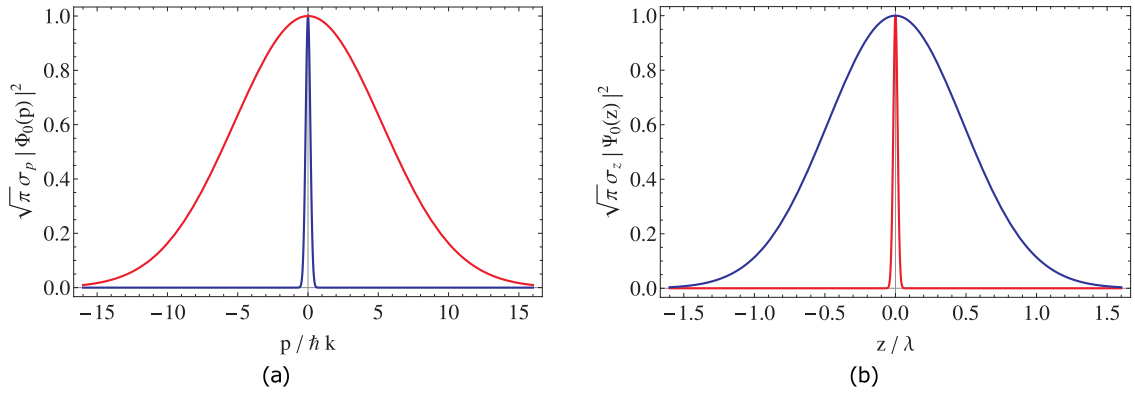

Figure B.1: Initial probability distributions in momentum- (part a) and position-space (part b). The red curves correspond to $T = 10$ μK and the blue curves to $T = 10$ nK.

We now apply the usual SW excitations to the initial wavefunction. After a SW pulse at $t = 0$, the position-space wavefunction is:

$$\psi_1(z,0) = e^{iu_1\cos(qz+\phi)}\psi_0(z) = \left(\frac{1}{\sqrt{\pi}\sigma_z}\right)^{1/2} e^{-\frac{1}{2}(z/\sigma_z)^2}\sum_n A_n e^{in\phi_1}e^{i(p_0+n\hbar q)z/\hbar}, \tag{B.3}$$

where $q = 2k$, $u_1 = \Omega_{eff}\tau_1$ is the pulse area of the SW, ϕ_1 is the phase of the SW and $A_n = (-i)^n J_n(u_1)$ is a coefficient from the Jacobi-Anger expansion.

To compute the evolution of the wavefunction after SW pulse, we perform the same steps used in the derivation of the echo signal in ch.2. Namely, we transform to momentum-space, apply the free-space operator: $e^{-ip^2t/2M\hbar}$, and then transform back to position-space. The resulting wavefunction is:

$$\psi_1(z,t) = \left(\frac{1}{\sqrt{\pi}\sigma_z F(t)}\right)^{1/2} e^{\frac{1}{2F(t)}(z/\sigma_z)^2}\sum_n A_n e^{in\phi}e^{-\frac{i}{F(t)}(p_0+n\hbar q)^2 t/2M\hbar}e^{\frac{i}{F(t)}(p_0+n\hbar q)z/\hbar}. \tag{B.4}$$

The function $F(t) = 1 + i\omega_D t$ describes the spatial dispersion of the wave function due to the Doppler width $\omega_D = \sigma_p^2/M\hbar = 2k_BT/\hbar$. This models the physical reality which has the faster atoms move away from the center of the sample.

Figure (B.2) shows the probability distribution at $t = 0.5\tau q$, where $\tau_q = \pi/\omega_q$ is the recoil period, and $\omega_q = \hbar q^2/2M$ is the recoil frequency. Note that at this early time, only the low-temperature sample shows any evidence of spatial modulation.

To extract the back-scattered read-out light that comprises our signal of interest, we take the $2k$-Fourier harmonic of the probability distribution:

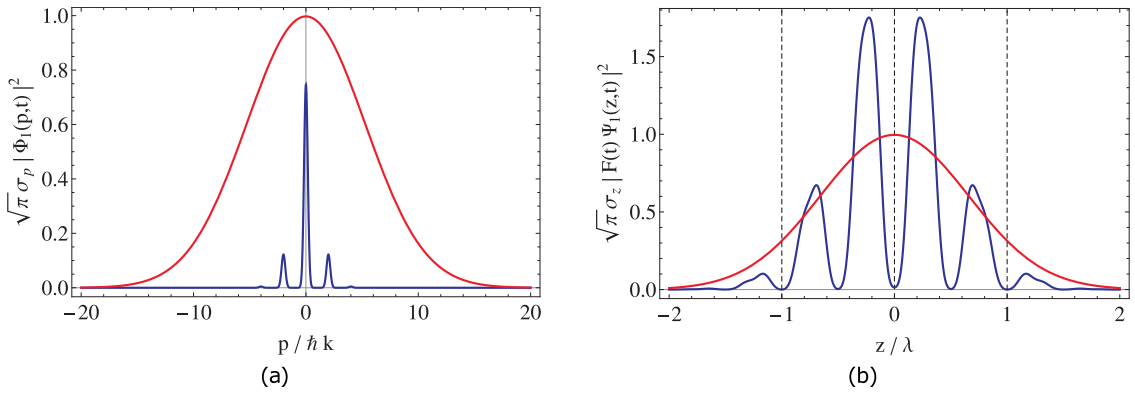

Figure B.2: Probability distributions at $t = 0.5\tau_q$ in momentum- (part a) and position-space (part b). The red curves correspond to $T = 10$ μK and the blue curves to $T = 10$ nK. As shown, the spatial modulation in the 10 nK sample is significant, while there is no discernible effect on the 10 μK sample.

$$E_1(t) \propto E_{RO} e^{iqp_0 t/M} e^{-(t/\tau_{coh})^2} \sum_{n,n'} A_n A_{n'}^* e^{i(n-n')\phi_1} e^{-[(n-n'+1)\hbar q/2\sigma_p]^2} e^{i(n+n')\omega_q t}, \quad \text{(B.5)}$$

where the timescale for the signal decay is $\tau_{coh} = 2M/q\sigma_p$, the coherence time. For the low-temperature (10 nK) sample, $\tau_{coh} \sim 90$ μs, allowing a few recoil oscillations in the signal to be observed, as shown in fig. (B.3). It is important to note that the signal shown here is produced by a single atom, with the decay in amplitude produced by the expansion of the probability distribution due to the temperature of the sample. The coherence time of the 10 μK sample is around 2.8 μs, which means that the signal decays away before any recoil oscillations can be seen.

To generate the expression for the back-scattered signal after two SW pulses, we apply the usual technique, as was used in ch.2. First, we set $t = T_{21}$ in eq. (B.4), which represents the evolved wavefunction at the time of the second SW pulse. We

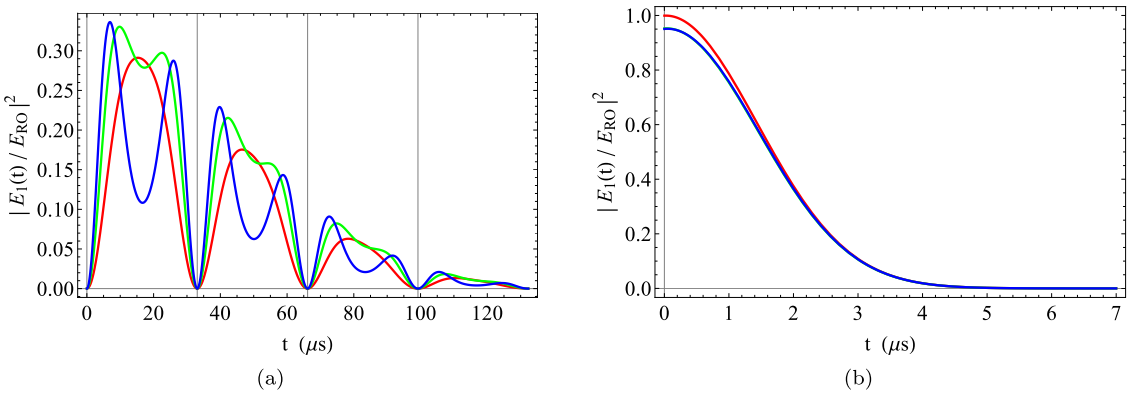

Figure B.3: Back-scattered field after one SW excitation for different sample temperatures. Part a) has $T = 10$ nK while part b) has $T = 10$ μK. For both plots, the red curves correspond to $u_1 = 0.75$, the green curves to $u_1 = 1.13$ and the blue curves to $u_1 = 1.5$. The blue and green curves are mostly overlapped in part b).

then apply a second instance of the SW interaction operator, followed by switching to momentum space, a free-space evolution operator until $t = 2T_{21}$, and then a switch back to position-space. The wavefunctions are then:

$$\phi_2(p,t) = \left(\frac{1}{\sqrt{\pi}\sigma_p}\right)^{1/2} e^{-ip^2(t-T_{21})/2M\hbar} \sum_{n,m} A_n B_m e^{i(n\phi_1+m\phi_2)} \times \left(e^{-\frac{1}{2}[F(T_{21})p-P_{n,m}]^2/\sigma_p^2} e^{-ip_n^2 T_{21}/2M\hbar}\right)^{1/F(T21)}, \quad \text{(B.6)}$$

$$\psi_2(z,t) = \left(\frac{1}{\sqrt{\pi}\sigma_z F(t)}\right)^{1/2} e^{-\frac{1}{2F(t)}(z/\sigma_z)^2} \sum_{n,m} A_n B_m e^{i(n\phi_1+m\phi_2)} \times \left(e^{-ip_n^2 T_{21}/2M\hbar} e^{-\frac{1}{F(t)}P_{n,m}^2(t-T_{21})/2M\hbar}\right)^{1/F(T_{21})} e^{\frac{i}{F(t)}P_{n,m}z/\hbar}. \quad \text{(B.7)}$$

Here, $A_n = (-i)^n J_n(u_1)$, $B_m = (-i)^m J_m(u_2)$, $p_n = p_0 + n\hbar q$ is the momentum

excited by the first SW pulse, $P_{n,m} = p_0 + [n + mF(T_{21})]\hbar q$ is the complex total momentum excited by the two pulses, and ϕ_i is the phase of the i$^{\text{th}}$ pulse.

One interesting feature of eq. (B.6) is the presence of the dimensionless scaling factor $F(T_{21})$ in the momentum wavefunction. This differs from the case of one SW pulse, implying that the motion of the atoms during the inter-pulse time plays a role in the interference seen at the read-out time.

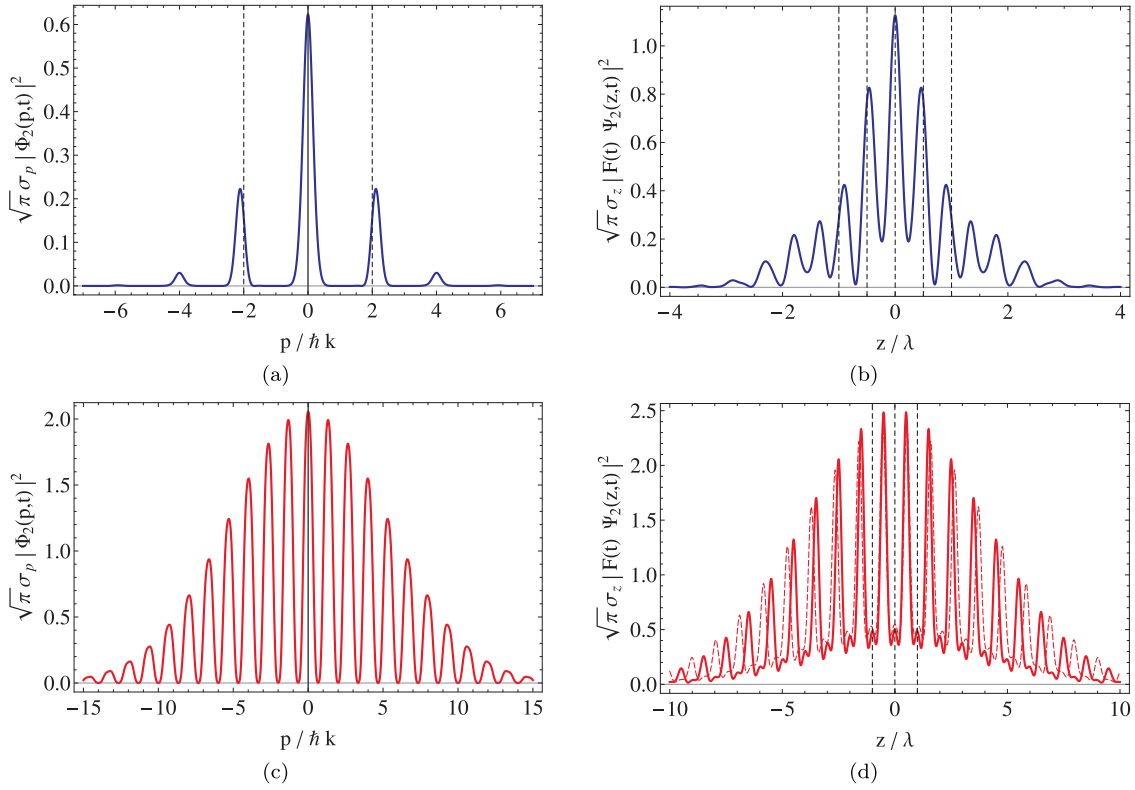

Figure B.4: Two-pulse probability distributions for 10 nK (parts a and b) and 10 μK (parts c and d) samples. Here, $T_{21} = 1.5\tau_q \sim 49$ μs. For all figures, $u_1 = u_2 = 0.75$ and the read-out time is $t = 2t_{21}$ except for in the dashed curve in part d), which is at $t = 2T_{21} + 6$ μs. Parts a) and c) are momentum-space wavefunctions, and parts b) and d) are position-space. Note the strong $\lambda/2$-periodic nature of the spatial distribution for the 10 nK sample. For the 10 μK sample, the strongest Fourier component has period λ at $t = 2T_{21}$. By time $t = 2T_{21} + 6$ μs, the period of the modulation has grown to $\sim 1.06\lambda$, which no longer satisfies the Bragg condition.

Within the time window between $2T_{21} - 2\tau_{coh}$ and $2T_{21} + 2\tau_{coh}$, the Bragg condition for back-scattering is closely satisfied by the modulation in the density of the sample. Outside of this window, the grating spacing will not support coherent back-scattering. Note that this statement neglects higher-order gratings that can form at greater integer multiples of the inter-pulse time, but our simulation is not interested in these effects. The back-scattered field from the ηq-Fourier harmonic component of the grating is given by:

$$E_2^{(\eta)}(t, T_{21}) = E_{RO} e^{-\eta^2[(t-T_{21})/\tau_{coh}]^2} e^{-i\eta^2\omega_q T_{21}} \sum_{n,n',m,m'} A_n A^*_{n'} B_m B^*_{m'} e^{i[(n-n')\phi_1 + (m-m')\phi_2]}$$
$$\times e^{i\eta q(P^*_{n',m'} + P_{n,m+\eta})(t-T_{21})/2M} \Big(e^{-[i(p_n^2 - p_{n'}^2) + (p_n^2 + p_{n'}^2)\omega_D T_{21}]T_{21}/2M\hbar}$$
$$\times e^{\left[(P_{n,m+\eta} - P^*_{n',m'})^2 + i\left(P^2_{n,m+\eta} - (P_{n',m'})^2\right)2\omega_D T_{21} - (P_{n,m} + P^*_{n',m'})^2(\omega_D T_{21})^2\right]/(2\sigma_p)^2} \Big)^{1/|F(T_{21})|^2}. \tag{B.8}$$

To calculate the signal from the λ-periodic component of the grating, we would use $E_2^{(1/2)}(t, T_{21})$, while the $\lambda/2$-periodic component's signal would be $E_2^{(1)}(t, T_{21})$. Figure (B.5), shows the predicted signals from eq. (B.8) for both λ- and $\lambda/2$-periodic components of the grating. It is interesting to note that the temporal width of the signal from the λ-periodic structure is a factor of 2 larger than that from the $\lambda/2$-period component, as shown in part a). This is a result of the coherence time being dependent on the Fourier component chosen: $\tau_{coh} = 1/\eta q \sigma_v$. Also, the λ-periodic signal oscillates at twice the atomic recoil frequency, while the $\lambda/2$-periodic signal oscillates at the recoil frequency. Both of these effects suggest that the λ-periodic signal is not the one that we see in experiment.

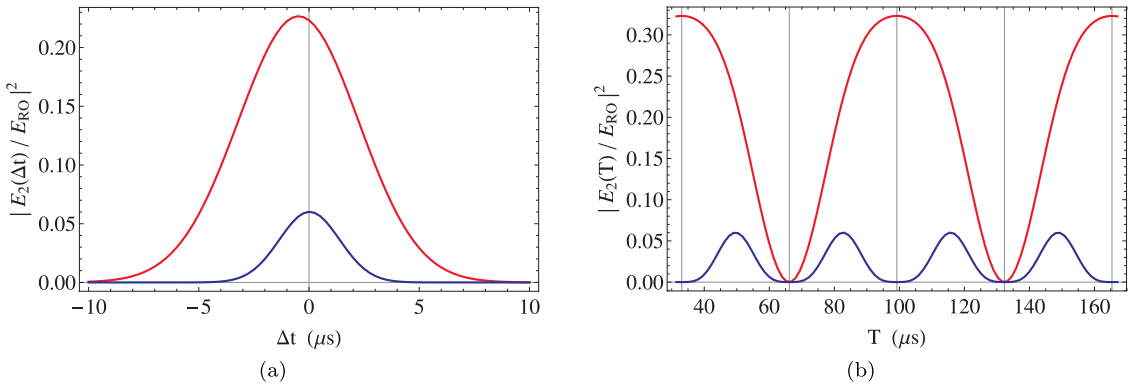

Figure B.5: Back-scattered signals $|E_2^{(1)}|^2$ (blue curves) and $|E_2^{(1/2)}|^2$ (red curves), produced by $\lambda/2$- and λ-periodic grating components, respectively. Part a) shows the signal as a function of $\Delta t = t - 2T_{21}$ for $T_{21} = 1.5\tau_q$. Part b) shows the signal as a function of T_21 for $\Delta t = 0$. For both figures, $u_1 = u_2 = 0.75$

B.2 Bragg scattering from a periodic structure

This section describes the theory used in Brynle's simulations of Bragg scattering from an arbitrary density distribution, $\rho(z)$. In the simulation, the Bragg-scattered electric field, E_{scat} is related to the input electric, E_{in} field by:

$$E_{\text{scat}}(t) = r(t)E_{\text{in}}, \quad I_{\text{scat}} = R(t)I_{\text{in}}, \tag{B.9}$$

where $r(t)$ is the complex reflectivity of a given time-evolving density distribution, $\rho(z,t)$ and $R(t) = |r(t)|^2$ is the reflection coefficient associated with $\rho(z,t)$. For the simulation to succeed, we must have a way of calculating $R(t)$ for an arbitrary density distribution.

The Fresnel reflectivity for a dielectric boundary is:

$$r = \frac{n_1 - n_2}{n_1 + n_2}. \tag{B.10}$$

For a dilute atomic gas, the refractive index is given by:

$$n \simeq 1 - \rho \frac{\mu_{eg}^2}{2\epsilon_0 \hbar \Gamma} \frac{\Delta/\Gamma}{1 + (\Delta/\Gamma)^2}. \tag{B.11}$$

The reflection coefficient for a spatially-modulated density is proportional to the contrast of the density modulation: $R \sim (n_{\max} - n_{\min})^2 \sim (\rho_{\max} - \rho_{\min})^2$.

For atoms confined to a periodic harmonic potential, the power of the Bragg-scattered light along the backwards direction can be approximated as [26, 85]:

$$P_{\text{scat}} = 2 \left(\frac{\pi \beta}{\lambda_B D} \right)^2 N^2 \left| \frac{\alpha(\Delta_B)}{\epsilon_0} \right|^2 I_{\text{in}}, \tag{B.12}$$

where λ_B and Δ_B are the wavelength and detuning of the applied Bragg beam, respectively. D is the diameter of the atomic sample transverse to the Bragg beam, N is the number of atoms, and β is the Debye-Waller factor, which contains information about the contrast of the sample:

$$\beta = e^{-\frac{1}{2}(\Delta \mathbf{k})^2 \omega_z^2} = e^{-2k_B^2 \omega_z^2}, \tag{B.13}$$

where $\Delta \mathbf{k} = 2\mathbf{k}_B$ is the lattice vector, and ω_z is the width of each wave packet confined to the lattice sites. The frequency-dependent, complex atomic polarizability, α is given by:

$$\alpha(\Delta_B) = \frac{3\epsilon_0 \lambda_B^3}{4\pi^2} \frac{S_{F,F'}}{i + 2\Delta_B/\Gamma}, \tag{B.14}$$

where $S_{F,F'}$ is an oscillator strength corresponding to the atomic transition, $F \rightarrow F'$. The polarizability describes the degree to which the atoms are polarized by

the incident Bragg beam. Going one step further, we can describe the reflection coefficient by combining eq. (B.12) and eq. (B.14):

$$R(\Delta_B) = \frac{18}{\pi^3}\left(\frac{\lambda_B}{D}\right)^4 \beta^2 N^2 \frac{S_{F,F'}^2}{1+(2\Delta_B/\Gamma)^2}. \tag{B.15}$$

Equation (B.15) assumes that the intensity of the back-scattered light is $I_{\text{scat}} = 4P_{\text{scat}}/\pi D^2$. Due to the $\propto N^2$ scaling of the reflection coefficient, we can more accurately classify this type of scattering as resonant elastic Rayleigh scattering from atoms in a periodic distribution satisfying the Bragg condition [85]. The field produced by this coherent scattering process is many orders of magnitude larger than that produced by diffuse atomic scattering, which scales $\propto N$, due to the phase-matching condition along the backwards direction.

An estimate of the reflection coefficient can be generated by substituting reasonable values for the parameters in eq. (B.15). If we choose $N = 10^9$ atoms, $D = 1$ cm, $\Delta_B = 30$ MHz, $S_{F,F'} = 7/10$ ($F = 2 \rightarrow F' = 3$ transition in ^{87}Rb [108]), and $\omega_z = \lambda/6$ (corresponding to $\beta = 0.11$), we obtain $R \sim 1.3 \times 10^{-3}$. This is comparable to the values we see in experiments.

B.2.1 Transfer matrix formalism

A more complete picture of Bragg scattering from density-modulated atomic samples is provided by considering a transfer matrix approach [29, 73, 109, 110]. This technique is identical to the ray tracing transfer matrix technique taught in introductory optics courses. We consider a thin section of the atomic sample, which extends from z to $z + \delta z$, having a constant density and polarizability. The electric

field along the Bragg axis can then be decomposed into forward- and backwards-traveling components, E^+ and E^-, respectively. The field is written as:

$$E(z) = E^+(z)e^{ik_B z} + E^-(z)e^{-ik_B z}. \tag{B.16}$$

To find the field components in the adjacent layer of the atomic distribution, we use a matrix, $\mathcal{M}$, called the transfer matrix, such that:

$$\begin{pmatrix} E^+(z+\delta z) \\ E^-(z+\delta z) \end{pmatrix} = \mathcal{M}_{\delta z}(z) \begin{pmatrix} E^+(z) \\ E^-(z) \end{pmatrix}. \tag{B.17}$$

To calculate the entirety of the transmitted and reflected beams, we must multiply together K copies of the transfer matrix, where $K = L/\delta z$ is the number of layers that we consider in the atomic sample of length L. For a single layer of the sample, the transfer matrix can be written as:

$$\mathcal{M}_{\delta z}(z) = \begin{pmatrix} 1+i\zeta(z) & i\zeta(z) \\ -i\zeta(z) & 1-i\zeta(z) \end{pmatrix} \begin{pmatrix} e^{ik_B\delta z} & 0 \\ 0 & e^{-ik_B\delta z} \end{pmatrix}. \tag{B.18}$$

The first matrix corresponds to the interaction of the light with the medium, where ζ is the position-dependent single-layer reflectivity:

$$\zeta(z) = -\rho(z)\delta z \frac{k_B}{2} \frac{\alpha(\Delta_B)}{\epsilon_0}, \tag{B.19}$$

where $\rho(z)$ is the atomic density at position z, and α is the polarizability from eq. (B.14). The second matrix in eq. (B.18) describes a free propagation over length δz. The complete transfer matrix of the atomic sample is then given by:

$$\mathcal{M} = \prod_{j=1}^{K} \mathcal{M}_{\delta z}(j\delta z). \tag{B.20}$$

To compute the reflectivity and transmissivity of the complete atomic sample, we can extract elements of the transfer matrix:

$$r = \frac{\mathcal{M}_{12}}{\mathcal{M}_{22}}, \quad t = \frac{1}{\mathcal{M}_{22}}. \tag{B.21}$$

As usual, the reflection and transmission coefficients are found by squaring the relevant quantities: $R = |r|^2$, $T = |t|^2$. It is important to realize that unlike traditional Fresnel coefficients for reflection and transmission, the atomic coefficients do not necessarily add to one. This is due to the complex nature of the polarizabiltiy of a two-level atom: the complex parts of r and t represent the incoherent scattering processes that the atoms can undergo.

B.2.2 Atomic density distribution

For the simulations to proceed, we require a mathematical description of the density profile of the atomic sample at the time of the Bragg beam. If we assume that the atomic sample's initial distribution of positions is a Gaussian,

$$p(z, 0) = \rho_0 e^{-(z/\sigma_z)^2}, \tag{B.22}$$

where ρ_0 is the peak density, and σ_z is the $1/e$ radius of the atom cloud, each atom within the sample will have a final probability distribution function which will depend on that particular atom's interactions with the SW AI pulses. As a

result, the overall final density distribution of the atom sample will simply be the sum of each atom's individual density function, weighted according to the overall distribution of atoms. For our system, the final probability distribution of the atoms will be $|\psi(z,t)|^2$, where the final wavefunction $\psi(z,t)$ will be given by either eq. (B.4) or eq. (B.7), depending on whether we have applied one or two SW pulses to the atoms, respectively. Since we are dealing with identical atoms, the overall density distribution of the sample is then the addition of N_L copies of $|\psi(z,t)|^2$, each centered at a z that is determined by $P(z_j)$:

$$\rho(z,t) = \rho_0(\sqrt{\pi}\sigma_z|F(t)|)\sum_{j=1}^{N_L} e^{-(z_j/\sigma_z)^2}|\psi(z-z_j,t)|^2, \tag{B.23}$$

where N_L is the number of atomic layers being summed:

$$N_L = L/s \sim 4\sigma_s\rho_0^{1/3}, \tag{B.24}$$

where $s \sim \rho_0^{-1/3}$ is the average atomic spacing, and $L \sim 4\sigma_s$ is the length of the sample. This quantity can equally be expressed as the number of SW nodes across the sample length, $N_{\lambda/2} = 2L/\lambda$, multiplied by the filling fraction of the atoms within the nodes, $r_{fill} = \lambda/2s$. For our typical experimental parameters are $\sigma_s = 0.5$ cm and $\rho_0 = 10^{10}$ atoms/cm^3, the number of nodes is $N_{\lambda/2} \sim 50,000$, the filling fraction is $r_{fill} \sim 0.08$ and therefore $N_L \sim 4200$.

B.3 Simulation results

Figure (B.6) shows an example of the density distribution for a 10 nK sample with a length $L \sim 20\lambda$, summed over 1000 atomic layers. Part a) shows the

density modulation of a single atom, which approaches 100%. When summing over a large number of atoms, if one does not take into account the change in SW phase seen by each atom at its starting position, $\phi_{z_j} = qz_j$, the summed density distribution's contrast tends to wash-out, as shown in part b). To correct this, we simply substitute ϕ_{z_j} for ϕ_1 in eq. (B.4), which restores the near 100% contrast that we expect, as shown in c).

The result of using smaller sample length and larger density (to speed calculation) conditions in eq. (B.23) is shown in fig. (B.6d and e). The reflection coefficient shows some unexpected modulation in the first few recoil periods, which are expected to dissipate for larger sample size and atom number.

Figure (B.7) shows an analogous set of figures as in fig. (B.6), except for a two-pulse simulation. In part a), we see a single 10 μK atom's probability distribution, while parts b) and c) show the summed probability distribution for the ensemble where each atom sees the same SW phase, and where the SW phase depends on the atom's position at the time of the SW pulse, respectively. It is notable that b) and c) are almost identical here, whereas in fig. (B.6), there was a large difference seen between the cases of constant and position-dependent SW phase. This effect is only seen when summing a large number of atoms, and it is due to the non-zero temperature of the sample. The atoms disperse along z at a rate that is not proportional to a multiple of the recoil velocity, which means that the modulation periods after two SW pulses of any two atoms in the sample will not necessarily align at the read out time.

Parts d) and e) of fig. (B.7) show the signal as functions of Δt and T_{21}, respectively. The two-lobed echo signal that we see habitually in experiment is faithfully

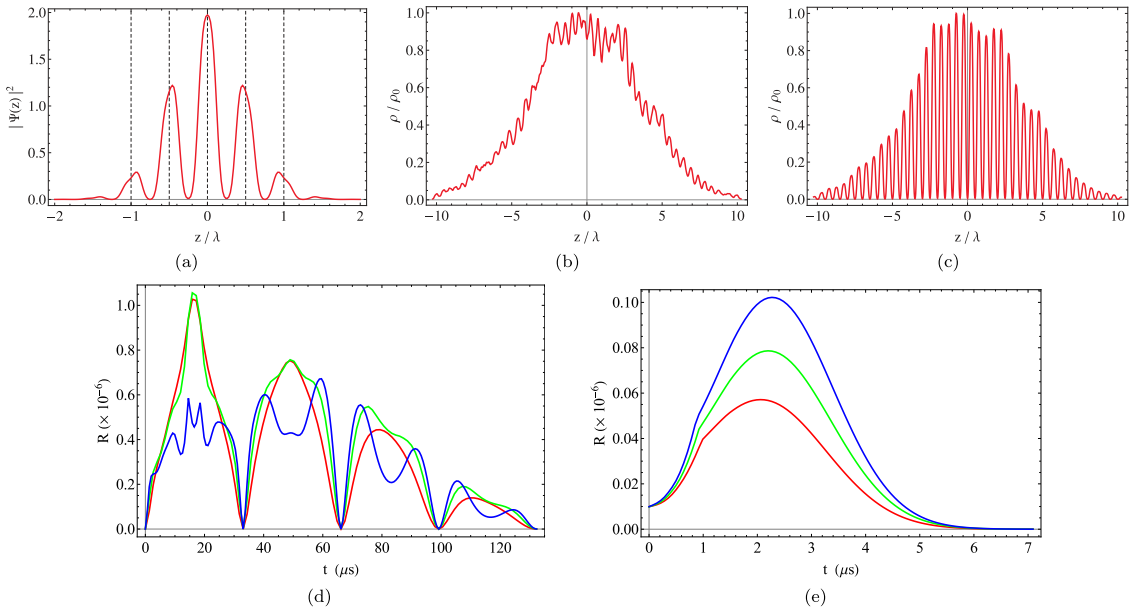

Figure B.6: One-pulse simulation results. Part a) shows an example of a the spatial probability distribution $|\psi_1(z)|^2$ of a 10 nK wave packet at time $t = 0.5\tau_q$ after a single SW pulse of area $u_1 = 0.75$. Part b) is the incorrectly-summed density distribution of 1000 atoms spread randomly over $L \sim 20\lambda$, all experiencing the same SW phase. Part c) shows the corrected picture, where each of the 1000 atoms sees the correct SW phase, based on its position in the distribution. Part d) shows the reflection coefficient for a 10 nK sample, while part e) shows the same for a 10 μK sample. For parts d) and e), the curve colours correspond to different pulse areas: red has $u_1 = 0.75$, green has $u_1 = 1.13$ and blue has $u_2 = 1.5$. For these curves, the density is the result of summing 100 atoms over $L = 20\lambda$, with $\rho_0 = 5 \times 10^{10}$ atoms/cm^3

reproduced in d), while part e) shown a non-physical increasing trend in signal amplitude as T_{21} increases, which is an artifact of how the simulation handles the expansion of the wavepackets. Nonetheless, we see the typical recoil-modulated echo amplitude, as we expect.

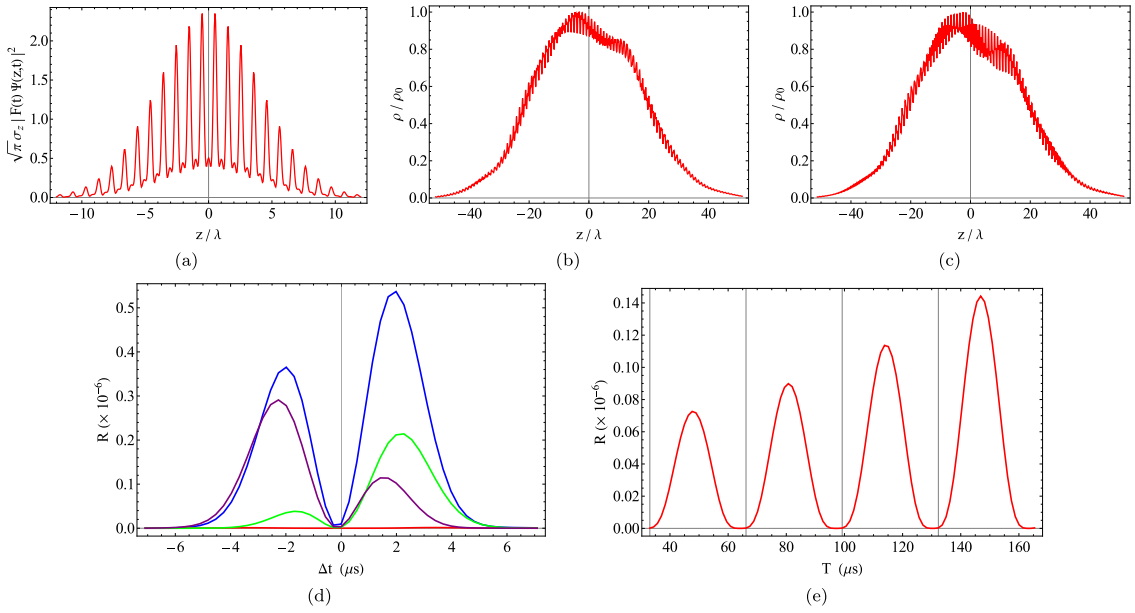

Figure B.7: Two-pulse simulation results. Part a) shows an example of the probability density of a 10 μK atom at time $t = 2T_{21}+2$ μs after two SW pulses ($T_{21} = 1.5\tau_q$, $u_1 = u_2 = 0.75$. This time corresponds to a maximum in the back-scattered signal. Parts b) and c) show the summed density distribution from 1000 atoms spread randomly over $\sim 100\lambda$ for constant and position-dependent SW phases, respectively. Part d) shows the reflection coefficient as a function of Δt, for 100 atoms spread over $\sim 100\lambda$ with density $\rho_0 = 5 \times 10^{10}$ atoms/cm^3. Each curve is coloured according to the pulse separation: red has $T_{21} = 1.0\tau_q$, green has $T_{21} = 1.25\tau_q$ blue has $T_{21} = 1.5\tau_q$ and purple has $T_{21} = 1.75\tau_q$. Part e) shows the reflection coefficient as a function of T_{21}, for $t = 2T_{21} + 2$ μs. For both d) and e), $u_1 = u_2 = 0.75$.

The response of the two-pulse signal to read out detuning is shown in fig. (B.8). Here we see some jagged features of the curve near zero detuning, which are not seen in experimental data. We believe that this is due to the integration of the

echo over Δt, which smooths-out these features. This is supported by seeing how the jagged features change as the value of Δt is changed in the simulation.

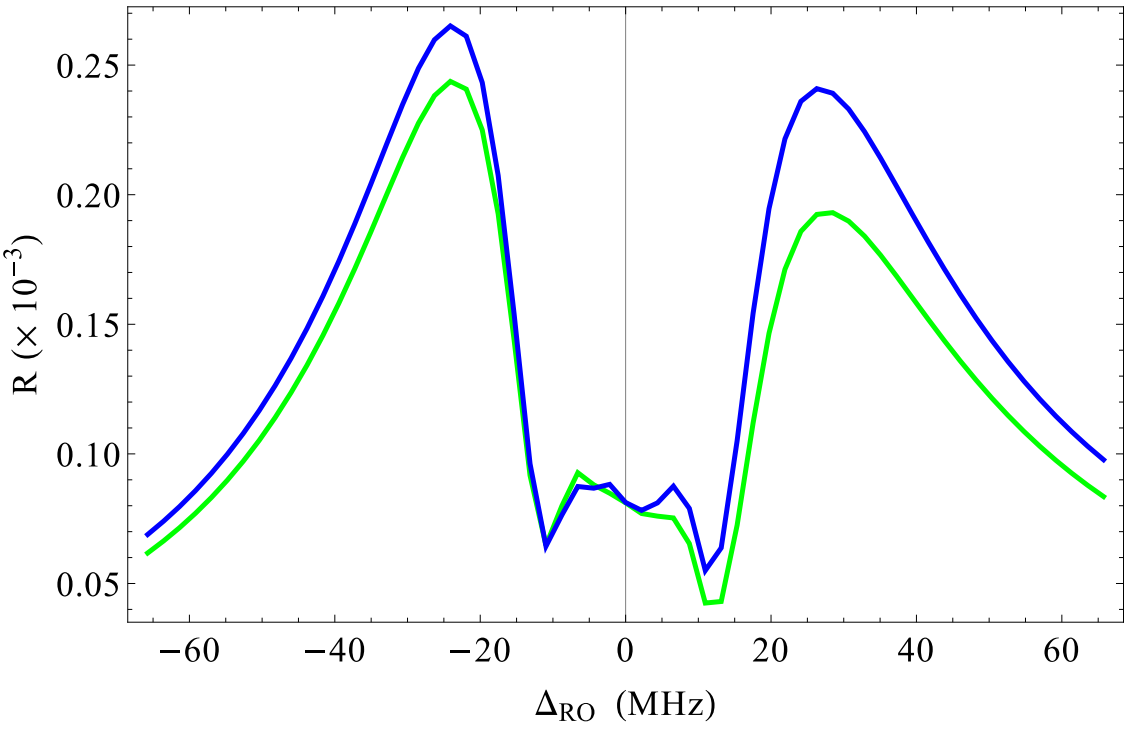

Figure B.8: Reflection coefficient as a function of read out detuning, Δ_{RO}. The two curves represent $\Delta t = -2$ μs (blue) and $\Delta t = 2$ μs (green). For both curves, $T = 10$ μK, $\rho_0 = 1.5 \times 10^{10}$ atoms/cm^3, $L = 2$ cm, $N_L = 1000$, $T_{21} = 51.5\tau_q \sim 1.7$ ms, $u_1 = u_2 = 0.75$ and $\Delta_{AI} = -380$ MHz

Bibliography

[1] K. Kim, J. Ko, Y. Cho, K. Lee, B. Kwak, and K. Park. A skew-symmetric cantilever accelerometer for automotive airbag applications. *Sensors and Actuators A*, 50:121–126, 1995.

[2] H. Seigel. A guide to high precision land gravimeter surveys. Technical report, Scintrex Ltd., 1995.

[3] T. Niebauer, G. Sasagawa, J. Faller, R. Hilt, and F. Klopping. A New Generation of Absolute Gravimeters. *Metrologia*, 32:159–180, 1995.

[4] P. Lett, R. Watts, C. Westbrook, W. Phillips, and P. Gould. Observation of Atoms Cooled Below the Doppler Limit. *Physical Review Letters*, 61(2):169–173, 1988.

[5] I. I. Rabi, S. Millman, P. Kusch, and J. R. Zacharias. The Molecular Beam Resonance Method for Measuring Nuclear Magnetic Moments. *Physical Review*, 55(1938):526–535, 1939.

[6] N. Ramsey. A New Molecular Beam Resonance Method. *Physical Review*, 76(7):996, 1949.

[7] P. Gould, G. Ruff, and D. Pritchard. Diffraction of Atoms by Light: The Near-Resonant Kapitza-Dirac Effect. *Physical Review Letters*, 56(8):827–830, 1986.

[8] D. Keith, M. Schattenburg, H. Smith, and D. Pritchard. Diffraction of Atoms by a Transmission Grating. *Physical Review Letters*, 61(14):1580–1583, 1988.

[9] O. Carnal and J. Mlynek. Young's double-slit experiment with atoms: A simple atom interferometer. *Physical review letters*, 66(21):2689, 1991.

[10] E. Raab, M. Prentiss, A. Cable, S. Chu, and D. Pritchard. Trapping of Neutral Sodium Atoms with Radiation Pressure. *Physical Review Letters*, 59(23):2631–2634, 1987.

[11] M. Kasevich, E. Riis, S. Chu, and R. DeVoe. RF Spectroscopy in an Atomic Fountain. *Physical Review Letters*, 63(6):612–615, 1989.

[12] M. Kasevich and S. Chu. Atomic Interferometry using Stimulated Raman Transitions. *Physical Review Letters*, 67(2):181–184, 1991.

[13] K. Gibble and S. Chu. Laser-Cooled Cs Frequency Standard and a Measurement of the Frequency Shift due to Ultracold Collisions. *Physical Review Letters*, 70(12):1771–1774, 1993.

[14] P. L. Kapitza and P. A. M. Dirac. The reflection of electrons from standing light waves. *Mathematical Proceedings of the Cambridge Philosophical Society*, 29:297–300, 1933.

[15] D. Freimund, A. Kayvan, and H. Batelaan. Observation of the Kapitza-Dirac Effect. *Nature*, 413:142–143, 2001.

[16] A. Peters, K. Chung, and S. Chu. Measurement of Gravitational Acceleration by Dropping Atoms. *Nature*, 400:849–852, 1999.

[17] S. Dickerson, J. Hogan, A. Sugarbaker, D. Johnson, and M. Kasevich. Multi-axis inertial sensing with long-time point source atom interferometry. *Physical review letters*, 111(8):083001, 2013.

[18] F. Riehle, T. Kisters, A. Witte, J. Helmcke, and C. Bordé. Optical Ramsey Spectroscopy in a Rotating Frame: Sagnac Effect in a Matter-Wave Interferometer. *Physical Review Letters*, 67(2):177–180, 1991.

[19] D. Schiavone, D. Capolongo, and M. Loddo. Near-Station Topographic Masses Correction for High-Accuracy Gravimetric Prospecting. *Geophysical Prospecting*, 57:739–752, 2009.

[20] J. Weber and J. Larson. Operation of LaCoste and Romberg Gravimeter at Sensitivity Approaching the Thermal Fluctuation Limits. *Journal of Geophysical Research*, 71(24):6005–6009, 1966.

[21] Micro-g Lacoste, Inc. Fg5-x absolute gravimeter specifications. http://www.microglacoste.com/fg5xspecs.php. Accessed: 29/03/2015.

[22] K.-J. Boller, A. Imamoğlu, and S. Harris. Observation of electromagnetically induced transparency. *Physical Review Letters*, 66(20):2593, 1991.

[23] M. Fleischhauer, A. Imamoglu, and J. Marangos. Electromagnetically Induced Transparency: Optics in Coherent Media. *Reviews of Modern Physics*, 77(2):633–673, 2005.

[24] S. Cahn, A. Kumarakrishnan, U. Shim, T. Sleator, P. Berman, and B. Dubetsky. Time-Domain de Broglie Wave Interferometry. *Physical Review Letters*, 79(5):784–787, 1997.

[25] P. Martin, B. Oldaker, A. Miklich, and D. Pritchard. Bragg Scattering of Atoms from a Standing Light Wave. *Physical Review Letters*, 60(6):515–518, 1988.

[26] S. Slama, C. Von Cube, B. Deh, A. Ludewig, C. Zimmermann, and Ph. W. Courteille. Phase-Sensitive Detection of Bragg Scattering at 1D Optical Lattices. *Physical Review Letters*, 94(193901):1–4, 2005.

[27] M. Kozuma, L. Deng, E. Hagley, J. Wen, R. Lutwak, K. Helmerson, S. Rolston, and W. Phillips. Coherent Splitting of Bose-Esinstein Condensed Atoms with Optically Induced Bragg Diffraction. *Physical Review Letters*, 82(5):871–875, 1999.

[28] G. Birkl, M. Gatzke, I. Deutsch, S. Rolston, and W. Phillips. Bragg Scattering from Atoms in Optical Lattices. *Physical Review Letters*, 75(15):2823–2826, 1995.

[29] S. Slama, C. von Cube, M. Kohler, C. Zimmermann, and P. Courteille. Multiple Reflections and Diffuse Scattering in Bragg Scattering at Optical Lattices. *Physical Review A*, 73(023424):1–9, 2006.

[30] M. Andersen and T. Sleator. Lattice Interferometer for Laser-Cooled Atoms. *Physical Review Letters*, 103(070402):1–4, 2009.

[31] A. Schilke, C. Zimmermann, P. Courteille, and W. Guerin. Photonic Band Gaps in One-Dimensionally Ordered Cold Atomic Vapors. *Physical Review Letters*, 106(223903):1–4, 2011.

[32] C. Mok. *Demonstration of Improved Sensitivity of Echo Atom Interferometers to Gravitational Acceleration*. PhD thesis, York University, 2013.

[33] C. Mok, B. Barrett, A. Carew, R. Berthiaume, S. Beattie, and A. Kumarakrishnan. Demonstration of Improved Sensitivity of Echo Interferometers to Gravitational Acceleration. *Physical Review A*, 88(023614):1–15, 2013.

[34] A. Firester. Design of Square Helmholtz Coil Systems. *Review of Scientific Instruments*, 37:1264, 1966.

[35] L. Aplet and J. Carson. A faraday effect optical isolator. *Applied Optics*, 3(4):544–545, 1964.

[36] S. Dimopoulos, P. Graham, J. Hogan, and M. Kasevich. General Relativistic Effects in Atom Interferometry. *Physical Review D*, 78(4):042003, 2008.

[37] P. Cheinet, B. Canuel, F. Pereira Dos Santos, A. Gauguet, F. Yver-Leduc, and A. Landragin. Measurement of the sensitivity function in a time-domain atomic interferometer. *IEEE Transactions on Instrumentation and Measurement*, 57(6):1141–1148, 2008.

[38] B. Barrett, P.-A. Gominet, E. Cantin, L. Antoni-Micollier, A. Bertoldi, B. Battelier, P. Bouyer, J. Lautier, and A. Landragin. Mobile and remote inertial sensing with atom interferometers. *Proceedings of the international school of physics "Enrico Fermi"*, 188:493–555, 2014.

[39] B. Barrett, L. Antoni-Micollier, L. Chichet, B. Battelier, P.-A. Gominet, A. Bertoldi, P. Bouyer, and A. Landragin. Correlative methods for dual-species quantum tests of the weak equivalence principle. *New Journal of Physics*, 17(8):085010, 2015.

[40] G. Sagnac. The demonstration of the luminiferous aether by an interferometer in uniform rotation. *Comptes Rendus*, 157:708–710, 1913.

[41] R. Hurst, J.-P. Wells, and G. Stedman. An Elementary Proof of the Geometrical Dependence of the Sagnac Effect. *Journal of Optics A: Pure and Applied Optics*, 9(10):838–841, 2007.

[42] T. Gustavson, P. Bouyer, and M. Kasevich. Precision Rotation Measurements with an Atom Interferometer Gyroscope. *Physical Review Letters*, 78(11):2046–2049, 1997.

[43] E. Su, S. Wu, and M. Prentiss. Atom Interferometry Using Wave Packets with Constant Spatial Displacements. *Physical Review A*, 81(4):043631, 2010.

[44] B. Young, F. Cruz, W. Itano, and J. Bergquist. Visible Lasers with Subhertz Linewidths. *Physical Review Letters*, 82:3799–3802, 1999.

[45] M. Weel and A. Kumarakrishnan. Laser-Frequency Stabilization using a Lock-In Amplifier. *Canadian Journal of Physics*, 80:1449–1458, 2002.

[46] K. Corwin, Z. Lu, C. Hand, R. Epstein, and C. Wieman. Frequency-Stabilized Diode Laser with the Zeeman Shift in an Atomic Vapor. *Applied optics*, 37(15):3295–3298, 1998.

[47] T. W. Hänsch. Repetitively pulsed tunable dye laser for high resolution spectroscopy. *Applied Optics*, 11(4):895–898, 1972.

[48] G. Hanes and J. Dobrowolski. An intracavity interference filter laser wavelength selector. *Applied Optics*, 8:482–483, 1969.

[49] X. Baillard, A. Gauguet, S. Bize, P. Lemonde, Ph. Laurent, A. Clairon, and P. Rosenbusch. Interference-Filter-Stabilized External-Cavity Diode Lasers. *Optics communications*, 266(2):609–613, 2006.

[50] M. Gilowski, Ch. Schubert, M. Zaiser, W. Herr, T. Wübbena, T. Wendrich, T. Müller, E. Rasel, and W. Ertmer. Narrow bandwidth interference filter-stabilized diode laser systems for the manipulation of neutral atoms. *Optics Communications*, 280(2):443–447, 2007.

[51] M. Fleming and A. Mooradian. Spectral Characteristics of External-Cavity Controlled Semiconductor Lasers. *IEEE Journal of Quantum Electronics*, QE-17(1):44–59, 1981.

[52] R. Berthiaume, A. Vorozcovs, and A. Kumarakrishnan. An auto-lock laser system for long term frequency stabilization. In *APS DAMOP 2010*, Houston, TX, 2010.

[53] A. Vorozcovs, V. Popovici, and A. Kumarakrishnan. An automated system for laser frequency stabilization using digital feedback. In *CAP Congress 2008*, June 2008.

[54] H. Beica, A. Carew, A. Vorozcovs, P. Dowling, A. Pouliot, B. Barron, and A. Kumarakrishnan. An auto-locked laser system for precision metrology and lidar applications. In *Laser radar technology and applications XXII*, volume Proc. SPIE 101910K-101910K-6, 2017.

[55] M. Weel and A. Kumarakrishnan. Observation of Ground-State Ramsey Fringes. *Physical Review A*, 67(6):061602(R), 2003.

[56] G. Biedermann, H. Mcguinness, A. Rakholia, and Y. Jau. Atom Interferometry in a Warm Vapor. *Physical Review Letters*, 118:163601, 2017.

[57] B. Barrett, A. Carew, H. Beica, A. Vorozcovs, A. Pouliot, and A. Kumarakrishnan. Prospects for Precise Measurements with Echo Atom Interferometry. *Atoms*, 4(19), 2016.

[58] B. Barrett, A. Carew, S. Beattie, and A. Kumarakrishnan. Measuring the Atomic Recoil Frequency using a Modified Grating-Echo Atom Interferometer. *Physical Review A*, 87:033626, 2013.

[59] B Barrett. *Techniques for Measuring the Atomic Recoil Frequency Using a Grating-Echo Atom Interferometer*. PhD thesis, York University, 2012.

[60] B. Barrett, I. Chan, C. Mok, A. Carew, I. Yavin, A. Kumarakrishnan, S. Cahn, and T. Sleator. Chapter 3 - Time-Domain Interferometry with

Laser-Cooled Atoms. In *Advances in Atomic, Molecular, and Optical Physics*, volume 60, pages 119–199. Elsevier Inc., 2011.

[61] S Beattie. *Atom Interferometric Studies of Light Scattering - A New Technique For Measuring Atomic Recoil*. PhD thesis, York University, 2009.

[62] S. Beattie, B. Barrett, M. Weel, I. Chan, C. Mok, S. Cahn, and A. Kumarakrishnan. Influence of Spontaneous Emission on a Single-State Atom Interferometer. *Physical Review A*, 77(1):013610, 2008.

[63] M. Weel. *Measurement of Atomic Recoil using Atom Interferometric Techniques*. PhD thesis, York University, 2005.

[64] D. Strekalov, A. Turlapov, A. Kumarakrishnan, and T. Sleator. Periodic Structures Generated in a Cloud of Cold Atoms. *Physical Review A*, 66(2):023601, 2002.

[65] A. Kumarakrishnan, S. Cahn, U. Shim, and T. Sleator. Magnetic Grating Echoes from Laser-Cooled Atoms. *Physical Review A*, 58(5):R3387–R3390, 1998.

[66] L. Allen and J. Eberly. *Optical Resonance and Two-Level Atoms*. Dover, 1987.

[67] P. Berman and V. Malinovsky. *Principles of Laser Spectroscopy and Quantum Optics*. Princeton: Princeton University Press, 2011.

[68] G. Milburn. *X-ray Crystallography : An Introduction to the Theory and Practice of Single-Crystal Structure Analysis*. London : Butterworths, 1973.

[69] W. Bragg. The diffraction of short electromagnetic waves by a crystal. In *Proceedings of the Cambridge Philosophical Society*, volume 17, pages 43–57, 1913.

[70] I. Usón and G. Sheldrick. Advances in direct methods for protein crystallography. *Current Opinion in Structural Biology*, 9(5):643 – 648, 1999.

[71] N. Dutta, T. Cella, and J. Zilko. Integrated External Cavity Distributed Bragg Reflector Laser. *Applied Physics Letters*, 50:644, 1987.

[72] R. Bouchendira, P. Clade, S. Guellati-Khélifa, F. Nez, and F. Biraben. New Determination of the Fine Structure Constant and Test of the Quantum Electrodynamics. *Physical Review Letters*, 106(080801):7–10, 2011.

[73] I. Deutsch, R. Spreeuw, S. Rolston, and W. Phillips. Photonic Band Gaps in Optical Lattices. *Physical Review A*, 52(2):1394, 1995.

[74] B. Barrett, I. Yavin, S. Beattie, and A. Kumarakrishnan. Numerical Simulation of a Multilevel Atom Interferometer. *Physical Review A*, 82:023625, 2010.

[75] U. Janicke and M. Wilkens. Atomic Motion in a Magneto-Optical Field. *Physical Review A*, 50(4):3265, 1994.

[76] C. Adams. Atom optics. *Contemporary Physics*, 35(1):1–19, 1994.

[77] S. Walter, D. Schneble, and A. Durst. Bloch Oscillations in lLattice Potentials with Controlled Aperiodicity. *Physical Review A*, 81(3):033623, 2010.

[78] R. Zhang, R. Sapiro, R. Mhaskar, and G. Raithel. Role of the Mean Field in Bloch Oscillations of a Bose-Einstein Condensate in an Optical Lattice and Harmonic Trap. *Physical Review A*, 78:053607, 2008.

[79] G. Campbell, A. Leanhardt, J. Mun, M. Boyd, E. Streed, W. Ketterle, and D. Pritchard. Photon Recoil Momentum in Dispersive Media. *Physical Review Letters*, 94:170403, 2005.

[80] T. Mossberg, R. Kachru, E. Whittaker, and S. Hartmann. Temporally Recurrent Spatial Ordering of Atomic Population in Gases: Grating Echoes. *Physical Review Letters*, 43(12):851–855, 1979.

[81] A. Yariv. *Quantum Electronics*. New York: Wiley, 1967.

[82] C. Kittel. *Introduction to Solid State Physics*. John Wiley and Sons, Inc., 1996.

[83] S. Marksteiner, R. Walser, P. Marte, and P. Zoller. Localization of Atoms in Light Fields: Optical Molasses, Adiabatic Compression and Squeezing. *Applied Physics B*, 60(2-3):145–153, 1995.

[84] M. Weidemuller, A. Hemmerich, A. Gorlitz, T. Esslinger, and T. Hansch. Bragg Diffraction in an Atomic Lattice Bound by Light. *Physical Review Letters*, 75(25):4583–4586, 1995.

[85] M. Weidemüller, A. Görlitz, T. Hänsch, and A. Hemmerich. Local and global properties of light-bound atomic lattices investigated by bragg diffraction. *Physical Review A*, 58(6):4647, 1998.

[86] B. Barrett. Personal Communication, 2014.

[87] B. Barrett, I. Chan, and A. Kumarakrishnan. Atom-Interferometric Techniques for Measuring Uniform Magnetic Field Gradients and Gravitational Acceleration. *Physical Review A*, 84(063623), 2011.

[88] D. Steck. Rubidium 85 D line data, version 2.1.6. https://steck.us/alkalidata/. Accessed: 29/03/2016.

[89] D. Steck. Rubidium 87 D line data, version 2.1.5. https://steck.us/alkalidata/. Accessed: 29/03/2016.

[90] J Dalibard and C. Cohen-Tannoudji. Laser Cooling Below the Doppler Limit by Polarization Gradients: Simple Theoretical Models. *JOSA B*, 6(11):2023–2045, 1989.

[91] M. Littman and H. Metcalf. Spectrally narrow pulsed dye laser without beam expander. *Appl. Opt.*, 17(14):2224–2227, 1978.

[92] K. Liu and M. Littman. Novel geometry for single-mode scanning of tunable lasers. *Opt. Lett.*, 6(3):117–118, 1981.

[93] K. Macadam, A. Steinbach, and C. Wieman. A Narrow-Band Tunable Diode Laser System with Grating Feedback, and a Saturated Absorption Spectrometer for Cs and Rb. *American Journal of Physics*, 60(12):1098–1111, 1992.

[94] K. Libbrecht and J. Hall. A Low-Noise High-Speed Diode Laser Current Controller. *Review of Scientific Instruments*, 64(8):2133–2135, 1993.

[95] L. Ricci, M. Weidemüller, T. Esslinger, A. Hemmerich, C. Zimmermann, V. Vuletic, W. Konig, and T. Hansch. A Compact Grating-Stabilized Diode Laser System for Atomic Physics. *Optics communications*, 117(95):541–549, 1995.

[96] E. Cook, P. Martin, T. Brown-heft, J. Garman, and D. Steck. High Passive-Stability Diode-Laser Design for use in Atomic-Physics Experiments. *Review of Scientific Instruments*, 83(043101):1–9, 2012.

[97] Toptica Photonics AG. Tunable diode lasers. http://www.toptica.com/fileadmin/Editors_English/11_brochures_datasheets/toptica_BR_Scientific_Lasers.pdf. Accessed: 11/08/2017.

[98] R. Berthiaume, A. Vorozcovs, and A. Kumarakrishnan. An auto-lock laser system for long term frequency stabilization. *APS DAMOP*, May 2010. Conference presentation.

[99] W. Riley. *Handbook of Frequency Stability Analysis*. NIST Special Publication 1065, 2008.

[100] D. Vukobratovich and R. Richard. Flexure Mounts for High-Resolution Optical Elements. In R. Bieringer and K. Harding, editors, *Optomechanical and Electro-Optical Design of Industrial Systems*, volume 0959, pages 18–36. SPIE, 1988.

[101] A. Hemmerich and T. Hansch. Two-Dimensional Atomic Crystal Bound by Light. *Physical Review Letters*, 70(4):410–413, 1993.

[102] P. Verkerk, B. Lounis, C. Salomon, and C. Cohen-Tannoudji. Dynamics and Spatial Order of Cold Cesium Atoms in a Periodic Optical Potential. *Physical Review Letters*, 68(26):3861–3864, 1992.

[103] D. Allan. Statistics of Atomic Frequency Standards. *Proceedings of the IEEE*, 54(2):221–230, 1966.

[104] D. Allan. Time and Frequency (Time-Domain) Characterization, Estimation, and Prediction of Precision Clocks and Oscillators. *IEEE transactions on ultrasonics, ferroelectrics, and frequency control*, 34(6):647–54, 1987.

[105] Stanford Research Systems, Inc. Prs10 rubidium oscillator specifications. http://thinksrs.com/products/PRS10.htm. Accessed: 16/06/2015.

[106] P. Williams, A. Rose, G. Day, T. Milner, and M. Deeter. Temperature dependence of the verdet constant in several diamagnetic glasses. *Appl. Opt.*, 30(10):1176–1178, 1991.

[107] S. Winter, C. Mok, and A. Kumarakrishnan. Tools for Laser Spectroscopy: The Design and Construction of a Faraday Isolator. *Canadian Journal of Physics*, 84(9):845–855, 2006.

[108] D. Steck. Rubidium 87 D line data, 2013.

[109] M. Artoni, G. La Rocca, and F. Bassani. Resonantly absorbing one-dimensional photonic crystals. *Physical Review E*, 72(4):046604, 2005.

[110] A. Schilke, C. Zimmermann, and W. Guerin. Photonic properties of one-dimensionally-ordered cold atomic vapors under conditions of electromagnetically induced transparency. *Physical Review A*, 86(2):023809, 2012.