LONG BASELINE ATOM INTERFEROMETRY

A DISSERTATION
SUBMITTED TO THE DEPARTMENT OF PHYSICS
AND THE COMMITTEE ON GRADUATE STUDIES
OF STANFORD UNIVERSITY
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

David Marvin Slaughter Johnson
June 2011

This dissertation is online at: http://purl.stanford.edu/pq947pb4276

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

Mark Kasevich, Primary Adviser

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

Hideo Mabuchi, Primary Adviser

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

Harindran Manoharan, Primary Adviser

Approved for the Stanford University Committee on Graduate Studies.

Patricia J. Gumport, Vice Provost Graduate Education

This signature page was generated electronically upon submission of this dissertation in electronic format. An original signed hard copy of the signature page is on file in University Archives.

Abstract

Due to its impressive sensitivity, long baseline atom interferometry is an exciting tool for tests of fundamental physics. We are currently constructing a 10-meter scale apparatus to test the Weak Equivalence Principle (WEP) using co-located Rb85 and Rb87 atom interferometers. This apparatus aims to improve the current limit on WEP violation 100-fold, which illustrates the power of this technique. This scientific goal sets stringent requirements on the kinematic preparation of the atomic test masses, the interferometer laser wavefront and stability, as well as the electromagnetic and gravitational field homogeneity of the interferometer region. The efforts to control these sources of systematic error are discussed. Additionally, applications of long baseline atom interferometry to space-based sensors for geodesy and gravitational wave detection are presented.

Contents

List of Tables

List of Figures

Chapter 1

Introduction

The observation of a falling object has provided great insights into the laws that govern our physical world. Using this technique at the leaning tower of Pisa, Galileo made the first forays into testing the equivalence principle, which would become one of the central postulates of general relativity. Similarly, Newton's apple provided a catalyst for the creation of modern physics. Modern implementations of this simple experiment have become sophisticated probes of gravity, and they have found practical applications outside the laboratory in navigation and geodesy. It is still of profound importance to be able to accurately measure the relative motion between an inertial test mass and an observer, and atom interferometry is an excellent tool to do exactly that.

Atom interferometry is a natural extension of the classic measurement of the acceleration of a test mass under acceleration [1]. Rather than using a macroscopic object, which is a large (upwards of 10^{23}!) collection of atoms that is difficult to isolate from the environment, an atom interferometer exploits de Broglie wave interference to measure the acceleration of single atoms in well-defined quantum states.

In its simplest conception, a light pulse atom interferometer uses an atom to measure the local phase of a laser beam at three places along its trajectory separated by equal times. The laser light acts as a finely corrugated ruler which is sampled by imprinting its phase on the atomic test mass at each of these three interaction points. The total phase shift imprinted on the atom is $\Delta\phi = \phi_1 - 2\phi_2 + \phi_3$ where ϕ_i is the phase of the laser at the interaction point i. The phase of the laser changes with position and time as $\phi = \mathbf{k} \cdot \mathbf{x} - \omega t$. For motion in a uniform gravitational field and interaction points separated by equal times, the total phase shift $\Delta\phi$ reduces to $\frac{1}{2}k_{\text{eff}}gT^2$ where T is the time between pulses, k_{eff} is the effective wavevector of the light (see Sec. 2 for more details), and g is the acceleration due to gravity. This phase shift increases linearly with an increase in interferometer length L by increasing the available interrogation time T, since $L = \frac{1}{2}gT^2$. For a 10 m baseline, and $k_{\text{eff}} \sim 10^7$ m^{-1}, this gives $\Delta\phi \sim 10^8$ rad per shot. Modern interferometers can resolve milli-radian phase variations per shot which gives a long baseline atom interferometer a sensitivity of $\sim 10^{-11}$ $g/\sqrt{Hz}$!

After early demonstrations of light pulse atom interferometers with short baselines L, it became immediately desirable to take advantage of the linear increase in sensitivity with length. Several atom interferometers with baseline $L \sim 1$ m have been built, and the work described below aims to push L to 10 m and eventually 30 km! These proposed, and under construction, interferometers have sensitivities that offer tantalizing opportunities to test fundamental gravitational physics in a controlled laboratory environment. The current 8.8 m apparatus under construction at Stanford will potentially increase the limit on violations of the weak equivalence principle (WEP) by 100-fold. The proposed AGIS-LEO gravitational wave detector offers a chance at detection of a predicted phenomenon of general relativity that has yet to be observed. The future of precision accelerometers promises to be quite exciting as these experiments begin to produce results.

This thesis is organized as follows. A review of light pulse atom interferometry is presented, along with a detailed systematic analysis of the Stanford ATEP (Atomic Test of the Equivalence Principle) experiment which is currently under construction. Then a proposal for AGIS-LEO (Atom Gravitational wave Interferometric Sensor - Low Earth Orbit) is discussed. Then, the experimental progress on the current effort at Stanford to construct an $L = 8.8$ m instrument is reported. Finally, two neat peripheral inventions for manipulating light using electro-optic modulators are presented in the appendix.

Chapter 2

Light Pulse Atom Interferometry

Light pulse atom interferometry has been intently studied for the past two decades. The following chapter reviews the past work and theoretical basis for this technique. This work was done in collaboration with Jason Hogan and Mark Kasevich. It has been published in *Proceedings of the International School of Physics Enrico Fermi on Atom Optics and Space Physics*. A version of it can also be found online at http://arxiv.org/abs/0806.3261. I was directly involved in producing this work.

2.1 Abstract

The light-pulse atom interferometry method is reviewed. Applications of the method to inertial navigation and tests of the Equivalence Principle are discussed.

2.2 Introduction

De Broglie wave interferometry using cold atoms is emerging as a new tool for basic science and technology. There are numerous approaches and applications which have evolved since the first demonstration experiments in the early 1990's. These notes will not attempt an exhaustive or comprehensive survey of the field. An excellent overview is provided in Ref. [2] and other lectures in this volume. These notes will focus on what has become known as light-pulse atom interferometry, which has found fruitful applications in gravitational physics and inertial sensor development.

These notes are organized as follows. We we first summarize basic theoretical concepts. We will then illustrate this formalism with a discussion of applications in inertial navigation and in a detailed design discussion of an experiment currently under development to test the Weak Equivalence Principle.

2.3 Atom interferometry overview

By analogy with their optical counterparts, atom interferometery seeks to exploit wave interference as a tool for precision metrology. In this measurement paradigm, a single particle (photon or atom) is made to coherently propagate over two paths. These paths are subsequently recombined using a beamsplitter, and their relative phase becomes manifest in the probability of detecting the particle in a given output port of the device. Hence, measuring the particle flux at the interferometer output ports enables determination of the phase shift. As this relative phase depends on physical interactions over the propagation paths, this enables the characterization of these interactions.

A key challenge for de Broglie wave interferometers is the development of techniques to coherently divide and recombine atomic wavepackets. For simplicity, consider an atom with initial momentum $\mathbf{p}$, characterized by a wavefunction $\psi(\mathbf{x}) \sim \exp\left[\frac{i}{\hbar}\mathbf{p}\cdot\mathbf{x}\right]$ [1]. The atom then is subject to a Hamiltonian interaction which is engineered to evolve the wavepacket into a momentum superposition state. One such interaction spatially modulates the amplitude of the wavefunction, so that, for example, $\exp\left[\frac{i}{\hbar}\mathbf{p}\cdot\mathbf{x}\right] \rightarrow f(\mathbf{x})\exp\left[\frac{i}{\hbar}\mathbf{p}\cdot\mathbf{x}\right]$, where $f(\mathbf{x})$ is a real periodic function with spatial frequency $\mathbf{k}$. Fourier decomposing $f(\mathbf{x})$ immediately shows that the final wavefunction is a coherent superposition of momenta $\mathbf{p}$, $\mathbf{p}\pm\hbar\mathbf{k}$, $\mathbf{p}\pm 2\hbar\mathbf{k}$, etc. In practice, such an interaction can be implemented by passing a collimated atomic beam through a microfabricated transmission grating, as demonstrated by Pritchard and co-workers [3].

Another interaction is one which spatially modulates the phase of the wavefunction. Consider, for example, $\psi(\mathbf{x}) \rightarrow \exp[i\mathbf{k}\cdot\mathbf{x}]\psi(\mathbf{x})$. This interaction results in a momentum translation $\mathbf{p}\rightarrow\mathbf{p}+\hbar\mathbf{k}$. A particularly useful implementation of this process imparts a spatial phase modulation by driving transitions between internal atomic states. For simplicity, consider a two-level atom with internal states $|1\rangle$ and $|2\rangle$ that are resonantly coupled by an applied optical traveling wave $\mathbf{E}\propto\exp[i\mathbf{k}\cdot\mathbf{x}]$ via the electric-dipole interaction $\hat{\mu}\cdot\mathbf{E}$ (where $\hat{\mu}$ is the dipole moment operator). If the atom is initially prepared in state $\psi(\mathbf{x})|1\rangle$, then following an interaction time t its state becomes $\alpha(t)\psi(\mathbf{x})|1\rangle + \beta(t)\exp[i\mathbf{k}\cdot\mathbf{x}]\psi(\mathbf{x})|2\rangle$ (see Section 2.4 for details). The interaction time can be chosen, for example, so that $|\alpha| = |\beta| = 1/\sqrt{2}$ to implement a beamsplitter (the $\pi/2$ pulse condition). In this case, the internal state of the atom becomes correlated with its external momentum. In practice, two-photon stimulated Raman transitions between groundstate hyperfine levels have proven to be particularly fruitful for implementing this class of beamsplitter. Why? Transitions are made between long lived hyperfine levels while the phase grating periodicity is twice that of a single photon optical transition (when the Raman transition is driven in a counter-propagating beam geometry).

The above mechanisms operate in free space. A new family of atom optics, based on control of atom wavepacket motion in atomic waveguides, is under development. The basic idea is that atoms are steered using microfabricated wires deposited on surfaces. These are loosely analogous to optical

[1] In reality, the wavefunction is spatially localized, so by the uncertainty principle, there must be a corresponding spread in momentum about the mean value.

fiber waveguides for light. In principle, coherent beamsplitters are implemented by the appropriate joining of waveguides. Recently, a combination of microwave and magnetic fields has allowed for creation of waveguide structures capable of coherent wavefront division. These structures have been used to demonstrate proof-of-principle interferometer topologies which have been used to study the coherence properties of the division process. The notes below will discuss free space atom optics, which have proven effective for precision measurement of inertial forces.

Exploiting the momentum exchange principles outlined above, it becomes straightforward to devise a de Broglie wave interferometer which is based on sequences of light pulses. For example, consider a three pulse sequence based on Raman transitions. An initial Raman $\pi/2$ pulse places an atom in a coherent superposition of wavepackets in states $|1\rangle$ and $|2\rangle$ whose mean momenta differ by $\hbar\mathbf{k}_{\text{eff}}$ (here $\mathbf{k}_{\text{eff}}$ is the effective wavevector of the Raman process, see below). After an interrogation time T these wavepackets separate by a distance $\hbar k_{\text{eff}}T/m$, where m is the atomic mass. A subsequent optical pulse is then applied whose duration is chosen to drive the transitions $|1\rangle \rightarrow |2\rangle$ and $|2\rangle \rightarrow |1\rangle$ with unit probability (a π pulse). This pulse has the effect of redirecting the momenta of the wavepackets so that at a time T later the wavepackets again overlap. A final $\pi/2$ pulse then serves as the exit beamsplitter.

2.4 Phase shift determination

In this section we review the method for calculating the phase difference between the two halves of the atom at the end of the light-pulse atom interferometer pulse sequence outlined above. These results are well-known [4, 5], but we are not aware of a complete, formal derivation of these rules in the literature. Other equivalent formalisms for this calculation do exist (see, for example [6, 7]). For Section 2.6.2 it is necessary to understand the formulae for the phase difference (Section 2.4.1). The proof of these formulae as well as a discussion of their range of validity is given in Section 2.4.2 but is not necessary for the rest of the paper.

2.4.1 Phase shift formulae

The main result we will show is that the total phase difference $\Delta\phi_{\text{tot}}$ between the two paths of an atom interferometer may be written as the sum of three easily calculated components:

$$\Delta\phi_{\text{tot}} = \Delta\phi_{\text{propagation}} + \Delta\phi_{\text{separation}} + \Delta\phi_{\text{laser}}. \tag{2.1}$$

For this calculation we take $\hbar = c = 1$.

The propagation phase $\Delta\phi_{\text{propagation}}$ arises from the free–fall evolution of the atom between light pulses and is given by

$$\Delta\phi_{\text{propagation}} = \sum_{\text{upper}} \left(\int_{t_I}^{t_F} (L_c - E_i) dt \right) - \sum_{\text{lower}} \left(\int_{t_I}^{t_F} (L_c - E_i) dt \right) \tag{2.2}$$

where the sums are over all the path segments of the upper and lower arms of the interferometer, and L_c is the classical Lagrangian evaluated along the classical trajectory of each path segment. In addition to the classical action, Eq. (2.2) includes a contribution from the internal atomic energy level E_i. The initial and final times t_I and t_F for each path segment, as well as L_c and E_i, all depend on the path segment.

The laser phase $\Delta\phi_{\text{laser}}$ comes from the interaction of the atom with the laser field used to manipulate the wavefunction at each of the beamsplitters and mirrors in the interferometer. At each interaction point, the component of the state that changes momentum due to the light acquires the phase of the laser $\phi_L(t_0, \mathbf{x}_c(t_0)) = \mathbf{k} \cdot \mathbf{x}_c(t_0) - \omega t_0 + \phi$ evaluated at the classical point of the interaction:

$$\Delta\phi_{\text{laser}} = \left(\sum_j \pm\phi_L(t_j, \mathbf{x}_u(t_j)) \right)_{\text{upper}} - \left(\sum_j \pm\phi_L(t_j, \mathbf{x}_l(t_j)) \right)_{\text{lower}} \tag{2.3}$$

The sums are over all the interaction points at the times t_j, and $\mathbf{x}_u(t)$ and $\mathbf{x}_l(t)$ are the classical trajectories of the upper and lower arm of the interferometer, respectively. The sign of each term depends on whether the atom gains (+) or loses (−) momentum as a result of the interaction.

The separation phase $\Delta\phi_{\text{separation}}$ arises when the classical trajectories of the two arms of the interferometer do not exactly intersect at the final beamsplitter (see Fig. 2.1). For a separation between the upper and lower arms of $\Delta\mathbf{x} = \mathbf{x}_l - \mathbf{x}_u$, the resulting phase shift is

$$\Delta\phi_{\text{separation}} = \bar{\mathbf{p}} \cdot \Delta\mathbf{x} \tag{2.4}$$

where $\bar{\mathbf{p}}$ is the average classical canonical momentum of the atom after the final beamsplitter.

2.4.2 Justification of phase shift formulae

The interferometer calculation amounts to solving the Schrodinger equation with the following Hamiltonian:

$$\hat{H}_{\text{tot}} = \hat{H}_{\text{a}} + \hat{H}_{\text{ext}} + \hat{V}_{\text{int}}(\hat{\mathbf{x}}) \tag{2.5}$$

Here $\hat{H}_{\text{a}}$ is the internal atomic structure Hamiltonian, $\hat{H}_{\text{ext}}$ is the Hamiltonian for the atom's external degrees of freedom (center of mass position and momentum), and $\hat{V}_{\text{int}}(\hat{\mathbf{x}}) = -\hat{\mu} \cdot \mathbf{E}(\hat{\mathbf{x}})$ is the atom-light interaction, which we take to be the electric dipole Hamiltonian with $\hat{\mu}$ the dipole moment operator.

The calculation is naturally divided into a series of light pulses during which $\hat{V}_{\text{int}} \neq 0$, and the segments between light pulses during which $\hat{V}_{\text{int}} = 0$ and the atom is in free-fall. When the light is off, the atom's internal and external degrees of freedom are decoupled. The internal eigenstates satisfy

$$i\partial_t \left|A_i\right\rangle = \hat{H}_{\text{a}} \left|A_i\right\rangle = E_i \left|A_i\right\rangle \tag{2.6}$$

and we write the solution as $\left|A_i\right\rangle = \left|i\right\rangle e^{-iE_i(t-t_0)}$ with time-independent eigenstate $\left|i\right\rangle$ and energy level E_i.

For the external state solution $\left|\psi\right\rangle$, we initially consider $\hat{H}_{\text{ext}} = H(\hat{\mathbf{x}}, \hat{\mathbf{p}})$ to be an arbitrary function of the external position and momentum operators:

$$i\partial_t \left|\psi\right\rangle = H(\hat{\mathbf{x}}, \hat{\mathbf{p}}) \left|\psi\right\rangle . \tag{2.7}$$

It is now useful to introduce a Galilean transformation operator

$$\hat{G}_c \equiv \hat{G}(\mathbf{x}_c, \mathbf{p}_c, L_c) = e^{i\int L_c dt} e^{-i\hat{\mathbf{p}}\cdot\mathbf{x}_c} e^{i\mathbf{p}_c\cdot\hat{\mathbf{x}}} \tag{2.8}$$

which consists of momentum boost by $\mathbf{p}_c$, a position translation by $\mathbf{x}_c$, and a phase shift. We choose to write

$$\left|\psi\right\rangle = \hat{G}_c \left|\phi_{CM}\right\rangle . \tag{2.9}$$

We will show that for a large class of relevant Hamiltonians, if $\mathbf{x}_c$, $\mathbf{p}_c$, and L_c are taken to be the classical position, momentum and Lagrangian, respectively, then $\left|\phi_{CM}\right\rangle$ is a wavepacket with $\langle\hat{\mathbf{x}}\rangle = \langle\hat{\mathbf{p}}\rangle = 0$, and the dynamics of $\left|\phi_{CM}\right\rangle$ do not affect the phase shift result (i.e., $\left|\phi_{CM}\right\rangle$ is the center of mass frame wavefunction). However, for now we maintain generality and just treat $\mathbf{x}_c$, $\mathbf{p}_c$, and L_c as arbitrary functions of time. Combining (2.7) and (2.9) results in

$$\begin{aligned} i\partial_t \left|\phi_{CM}\right\rangle &= \left\{\hat{G}_c^\dagger H(\hat{\mathbf{x}}, \hat{\mathbf{p}})\hat{G}_c - i\hat{G}_c^\dagger \partial_t \hat{G}_c\right\} \left|\phi_{CM}\right\rangle \\ &= \left\{H(\hat{\mathbf{x}} + \mathbf{x}_c, \hat{\mathbf{p}} + \mathbf{p}_c) + \dot{\mathbf{p}}_c \cdot \hat{\mathbf{x}} - (\hat{\mathbf{p}} + \mathbf{p}_c)\cdot\dot{\mathbf{x}}_c + L_c\right\} \left|\phi_{CM}\right\rangle \end{aligned} \tag{2.10}$$

where we used the following identities:

$$\begin{aligned} \hat{G}_c^\dagger \hat{\mathbf{x}} \hat{G}_c &= \hat{\mathbf{x}} + \mathbf{x}_c \\ \hat{G}_c^\dagger \hat{\mathbf{p}} \hat{G}_c &= \hat{\mathbf{p}} + \mathbf{p}_c \\ \hat{G}_c^\dagger H(\hat{\mathbf{x}}, \hat{\mathbf{p}}) \hat{G}_c &= H(\hat{\mathbf{x}} + \mathbf{x}_c, \hat{\mathbf{p}} + \mathbf{p}_c) \end{aligned} \tag{2.11}$$

Next, we Taylor expand $H(\hat{\mathbf{x}} + \mathbf{x}_c, \hat{\mathbf{p}} + \mathbf{p}_c)$ about $\mathbf{x}_c$ and $\mathbf{p}_c$,

$$H(\hat{\mathbf{x}} + \mathbf{x}_c, \hat{\mathbf{p}} + \mathbf{p}_c) = H(\mathbf{x}_c, \mathbf{p}_c) + \nabla_{\hat{\mathbf{x}}} H(\mathbf{x}_c, \mathbf{p}_c)\cdot\hat{\mathbf{x}} + \nabla_{\hat{\mathbf{p}}} H(\mathbf{x}_c, \mathbf{p}_c)\cdot\hat{\mathbf{p}} + \hat{H}_2 \tag{2.12}$$

where $\hat{H}_2$ contains all terms that are second order or higher in $\hat{\mathbf{x}}$ and $\hat{\mathbf{p}}$. (We will ultimately be allowed to neglect $\hat{H}_2$ in this calculation.) Inserting this expansion and grouping terms yields

$$i\partial_t \left|\phi_{CM}\right\rangle = \Big\{\big(H_c - \dot{\mathbf{x}}_c \cdot \mathbf{p}_c + L_c\big) + \big(\nabla_{\mathbf{x}_c} H_c + \dot{\mathbf{p}}_c\big)\cdot\hat{\mathbf{x}} + \big(\nabla_{\mathbf{p}_c} H_c - \dot{\mathbf{x}}_c\big)\cdot\hat{\mathbf{p}} + \hat{H}_2\Big\}\left|\phi_{CM}\right\rangle$$

where we have defined the classical Hamiltonian $H_c \equiv H(\mathbf{x}_c, \mathbf{p}_c)$. If we now let $\mathbf{x}_c$, $\mathbf{p}_c$, and L_c satisfy Hamilton's equations,

$$\begin{aligned}
\dot{\mathbf{x}}_c &= \nabla_{\mathbf{p}_c} H_c && (2.13)\\
\dot{\mathbf{p}}_c &= -\nabla_{\mathbf{x}_c} H_c\\
L_c &= \dot{\mathbf{x}}_c \cdot \mathbf{p}_c - H_c
\end{aligned}$$

with $\mathbf{p}_c \equiv \nabla_{\dot{\mathbf{x}}_c} L_c$ the classical canonical momentum, then $\left|\phi_{CM}\right\rangle$ must satisfy

$$i\partial_t \left|\phi_{CM}\right\rangle = \hat{H}_2 \left|\phi_{CM}\right\rangle \qquad (2.14)$$

Next we show that it is possible to choose $\left|\phi_{CM}\right\rangle$ with $\langle \hat{x}\rangle = \langle \hat{p}\rangle = 0$ for a certain class of $\hat{H}_2$, so that $\mathbf{x}_c$ and $\mathbf{p}_c$ completely describe the atom's classical center of mass trajectory. This is known as the semi-classical limit. Starting from Ehrenfest's theorem for the expectation values of $\left|\phi_{CM}\right\rangle$,

$$\begin{aligned}
\partial_t \langle \hat{x}_i\rangle &= i\Big\langle\Big[\hat{H}_2, \hat{x}_i\Big]\Big\rangle = \Big\langle \partial_{\hat{p}_i}\hat{H}_2\Big\rangle && (2.15)\\
\partial_t \langle \hat{p}_i\rangle &= i\Big\langle\Big[\hat{H}_2, \hat{p}_i\Big]\Big\rangle = -\Big\langle \partial_{\hat{x}_i}\hat{H}_2\Big\rangle && (2.16)
\end{aligned}$$

and expanding about $\langle\hat{\mathbf{x}}\rangle$ and $\langle\hat{\mathbf{p}}\rangle$,

$$\begin{aligned}
\partial_t \langle \hat{x}_i\rangle &= \Big\langle \partial_{\hat{p}_i}\hat{H}_2\Big|_{\langle\hat{\mathbf{x}}\rangle,\langle\hat{\mathbf{p}}\rangle} + \partial_{\hat{p}_j}\partial_{\hat{p}_i}\hat{H}_2\Big|_{\langle\hat{\mathbf{x}}\rangle,\langle\hat{\mathbf{p}}\rangle}(\hat{p}_j - \langle\hat{p}_j\rangle) + \partial_{\hat{x}_j}\partial_{\hat{p}_i}\hat{H}_2\Big|_{\langle\hat{\mathbf{x}}\rangle,\langle\hat{\mathbf{p}}\rangle}(\hat{x}_j - \langle\hat{x}_j\rangle)\\
&\quad + \frac{1}{2\,!}\,\partial_{\hat{p}_i}\partial_{\hat{p}_j}\partial_{\hat{p}_k}\hat{H}_2\Big|_{\langle\hat{\mathbf{x}}\rangle,\langle\hat{\mathbf{p}}\rangle}(\hat{p}_j - \langle\hat{p}_j\rangle)(\hat{p}_k - \langle\hat{p}_k\rangle) + \cdots\Big\rangle\\
\partial_t \langle \hat{p}_i\rangle &= \Big\langle \partial_{\hat{x}_i}\hat{H}_2\Big|_{\langle\hat{\mathbf{x}}\rangle,\langle\hat{\mathbf{p}}\rangle} + \partial_{\hat{x}_j}\partial_{\hat{x}_i}\hat{H}_2\Big|_{\langle\hat{\mathbf{x}}\rangle,\langle\hat{\mathbf{p}}\rangle}(\hat{x}_j - \langle\hat{x}_j\rangle) + \partial_{\hat{p}_j}\partial_{\hat{x}_i}\hat{H}_2\Big|_{\langle\hat{\mathbf{x}}\rangle,\langle\hat{\mathbf{p}}\rangle}(\hat{p}_j - \langle\hat{p}_j\rangle)\\
&\quad + \frac{1}{2\,!}\,\partial_{\hat{x}_k}\partial_{\hat{x}_j}\partial_{\hat{x}_i}\hat{H}_2\Big|_{\langle\hat{\mathbf{x}}\rangle,\langle\hat{\mathbf{p}}\rangle}(\hat{x}_j - \langle\hat{x}_j\rangle)(\hat{x}_k - \langle\hat{x}_k\rangle) + \cdots\Big\rangle
\end{aligned}$$

we find the following:

$$\begin{aligned}
\partial_t \langle \hat{x}_i\rangle &= \partial_{\hat{p}_i}\hat{H}_2\Big|_{\langle\hat{\mathbf{x}}\rangle,\langle\hat{\mathbf{p}}\rangle} + \frac{1}{2\,!}\,\partial_{\hat{p}_k}\partial_{\hat{p}_j}\partial_{\hat{p}_i}\hat{H}_2\Big|_{\langle\hat{\mathbf{x}}\rangle,\langle\hat{\mathbf{p}}\rangle}\Delta p_{jk}^2 + \cdots && (2.17)\\
\partial_t \langle \hat{p}_i\rangle &= -\,\partial_{\hat{x}_i}\hat{H}_2\Big|_{\langle\hat{\mathbf{x}}\rangle,\langle\hat{\mathbf{p}}\rangle} - \frac{1}{2\,!}\,\partial_{\hat{x}_k}\partial_{\hat{x}_j}\partial_{\hat{x}_i}\hat{H}_2\Big|_{\langle\hat{\mathbf{x}}\rangle,\langle\hat{\mathbf{p}}\rangle}\Delta x_{jk}^2 + \cdots && (2.18)
\end{aligned}$$

where $\Delta x_{jk}^2 \equiv \langle \hat{x}_j \hat{x}_k \rangle - \langle \hat{x}_j \rangle \langle \hat{x}_k \rangle$ and $\Delta p_{jk}^2 \equiv \langle \hat{p}_j \hat{p}_k \rangle - \langle \hat{p}_j \rangle \langle \hat{p}_k \rangle$ are measures of the wavepacket's width in phase space [2]. This shows that if $\hat{H}_2$ contains no terms higher than second order in $\hat{\mathbf{x}}$ and $\hat{\mathbf{p}}$, then Ehrenfest's theorem reduces to Hamilton's equations, and the expectation values follow the classical trajectories. Furthermore, this implies that we can choose $|\phi_{CM}\rangle$ to be the wavefunction in the atom's rest frame, since $\langle \hat{x} \rangle = \langle \hat{p} \rangle = 0$ is a valid solution to Eqs. (2.17) and (2.18) so long as all derivatives of $\hat{H}_2$ higher than second order vanish. In addition, even when this condition is not strictly met, it is often possible to ignore the non-classical corrections to the trajectory so long as the phase space widths Δx_{jk} and Δp_{jk} are small compared to the relevant derivatives of $\hat{H}_2$ (i.e., the semi-classical approximation). For example, such corrections are present for an atom propagating in the non-uniform gravitational field g of the Earth for which $\partial_{\hat{r}}\partial_{\hat{r}}\partial_{\hat{r}}\hat{H}_2 \sim \partial_r^2 g$. Assuming an atom wavepacket width $\Delta x \lesssim 1$ mm, the deviation from the classical trajectory is $\partial_t \langle \hat{p} \rangle \sim (\partial_r^2 g)\Delta x^2 \lesssim 10^{-20} g$, which is a negligibly small correction even in the context of the $\sim 10^{-15} g$ apparatus we describe below for testing the Equivalence Principle.

The complete solution for the external wavefunction requires a solution of Eq. (2.14) for $|\phi_{CM}\rangle$, but this is non-trivial for general $\hat{H}_2$. In the simplified case where $\hat{H}_2$ is second order in $\hat{\mathbf{x}}$ and $\hat{\mathbf{p}}$, the exact expression for the propagator is known [8] and may be used to determine the phase acquired by $|\phi_{CM}\rangle$. However, this step is not necessary for our purpose, because for second order external Hamiltonians the operator $\hat{H}_2$ does not depend on either $\mathbf{x}_c$ or $\mathbf{p}_c$. In this restricted case, the solution for the rest frame wavefunction $|\phi_{CM}\rangle$ does not depend on the atom's trajectory. Therefore, any additional phase evolution in $|\phi_{CM}\rangle$ must be the same for both arms of the interferometer and so does not contribute to the phase difference. This argument breaks down for more general $\hat{H}_2$, as does the semi-classical description of the atom's motion, but the corrections will depend on the width of $|\phi_{CM}\rangle$ in phase space as shown in Eqs. (2.17) and (2.18). We ignore all such wavepacket-structure induced phase shifts in this analysis by assuming that the relevant moments $\{\Delta x_{jk}, \Delta p_{jk}, \ldots\}$ are sufficiently small so that these corrections can be neglected. As shown above for the non-uniform $(\partial_r^2 g \neq 0)$ gravitational field of the Earth, this condition is easily met in many experimentally relevant situations.

Finally, we can write the complete solution for the free propagation between the light pulses:

$$\langle \mathbf{x} | \psi, A_i \rangle = \langle \mathbf{x} | \hat{G}_c | \phi_{CM} \rangle | A_i \rangle = e^{i \int_{t_I}^{t_F} L_c dt} e^{i \mathbf{p}_c \cdot (\mathbf{x} - \mathbf{x}_c)} \phi_{CM}(\mathbf{x} - \mathbf{x}_c) | i \rangle e^{-i E_i (t_F - t_I)} \tag{2.19}$$

We see that this result takes the form of a traveling wave with de Broglie wavelength set by $\mathbf{p}_c$ multiplied by an envelope function $\phi_{CM}(\mathbf{x})$, both of which move along the classical path $\mathbf{x}_c$. Also, the wavepacket accumulates a propagation phase shift given by the classical action along this path,

[2] In general, there will also be cross terms with phase space width such as $\langle \hat{x}_j \hat{p}_k \rangle - \langle \hat{x}_j \rangle \langle \hat{p}_k \rangle$.

as well as an additional phase shift arising from the internal atomic energy:

$$\Delta\phi_{\text{propagation}} = \sum_{\text{upper}} \left(\int_{t_I}^{t_F} (L_c - E_i) dt \right) - \sum_{\text{lower}} \left(\int_{t_I}^{t_F} (L_c - E_i) dt \right) \tag{2.20}$$

where the sums are over all the path segments of the upper and lower arms of the interferometer, and t_I, t_F, L_c, and E_i all depend on the path.

Next, we consider the time evolution while the light is on and $\hat{V}_{\text{int}} \neq 0$. In this case, the atom's internal and external degrees of freedom are coupled by the electric dipole interaction, so we work in the interaction picture using the following state ansatz:

$$|\Psi\rangle = \int d\mathbf{p} \sum_i c_i(\mathbf{p}) \, |\psi_{\mathbf{p}}\rangle \, |A_i\rangle \tag{2.21}$$

where we have used the momentum space representation of $|\phi_{CM}\rangle$ and so $|\psi_{\mathbf{p}}\rangle \equiv \hat{G}_c e^{-i\hat{H}_2(t-t_0)} \, |\mathbf{p}\rangle$. Inserting this state into the Schrodinger equation gives the interaction picture equations,

$$\begin{aligned} i\partial_t \, |\Psi\rangle \quad &= \quad i \int d\mathbf{p} \sum_i \frac{\partial c_i(\mathbf{p})}{\partial t} \, |\psi_{\mathbf{p}}\rangle \, |A_i\rangle + \hat{H}_{\text{a}} \, |\Psi\rangle + \hat{H}_{\text{ext}} \, |\Psi\rangle = \hat{H}_{\text{tot}} \, |\Psi\rangle & (2.22) \\ \Rightarrow \dot{c}_i(\mathbf{p}) \quad &\equiv \quad \frac{\partial c_i(\mathbf{p})}{\partial t} = \frac{1}{i} \int d\mathbf{p}' \sum_j c_j(\mathbf{p}') \, \langle A_i| \, \langle \psi_{\mathbf{p}}| \, \hat{V}_{\text{int}}(\hat{\mathbf{x}}) \, |\psi_{\mathbf{p}'}\rangle \, |A_j\rangle & (2.23) \end{aligned}$$

where we used (2.6) and (2.7) as well as the orthonormality of $|A_i\rangle$ and $|\psi_{\mathbf{p}}\rangle$. The interaction matrix element can be further simplified by substituting in $|\psi_{\mathbf{p}}\rangle$ and using identity (2.11):

$$\begin{aligned} \langle \psi_{\mathbf{p}}| \, \hat{V}_{\text{int}}(\hat{\mathbf{x}}) \, |\psi_{\mathbf{p}'}\rangle \quad &= \quad \langle \mathbf{p}| \, e^{i\hat{H}_2(t-t_0)} \hat{V}_{\text{int}}(\hat{\mathbf{x}} + \mathbf{x}_c) e^{-i\hat{H}_2(t-t_0)} \, |\mathbf{p}'\rangle & (2.24) \\ &= \quad \langle \mathbf{p}| \, \hat{V}_{\text{int}}(\hat{\mathbf{x}} + \mathbf{x}_c) \, |\mathbf{p}'\rangle \, e^{i\left(\frac{\mathbf{p}^2}{2m} - \frac{\mathbf{p}'^2}{2m}\right)(t-t_0)} & \end{aligned}$$

where we have made the simplifying approximation that $\hat{H}_2 \approx \frac{\hat{\mathbf{p}}^2}{2m}$. This approximation works well as long as the light pulse time $\tau \equiv t - t_0$ is short compared to the time scale associated with the terms dropped from $\hat{H}_2$. For example, for an atom in the gravitational field of Earth, this approximation ignores the contribution $m(\partial_r g)\hat{x}^2$ from the gravity gradient, which for an atom of size $\Delta x \approx 1$ mm leads to a frequency shift $\sim m(\partial_r g)\Delta x^2 \sim 1$ mHz. For a typical pulse time $\tau < 1$ ms, the resulting errors are $\lesssim 1$ μrad and can usually be neglected. Generally, in this analysis we will assume the short pulse (small τ) limit and ignore all effects that depend on the finite length of the light pulse. These systematic effects can sometimes be important, but they are calculated elsewhere[9][10]. In the case of the ^{87}Rb–^{85}Rb Equivalence Principle experiment we discuss below, such errors are common-mode suppressed in the differential signal because we use the same laser pulse to manipulate both atoms simultaneously.

As mentioned before, we typically use a two photon process for the atom optics (i.e., Raman or Bragg) in order to avoid transferring population to the short-lived excited state. However, from the point of view of the current analysis, these three-level systems can typically be reduced to effective two-level systems[11][12]. Since the resulting phase shift rules are identical, we will assume a two-level atom coupled to a single laser frequency to simplify the analysis. Assuming a single traveling wave excitation $\mathbf{E}(\hat{\mathbf{x}}) = \mathbf{E}_0 \cos(\mathbf{k}\cdot\hat{\mathbf{x}} - \omega t + \phi)$, Eq. (2.23) becomes

$$\dot{c}_i(\mathbf{p}) = \frac{1}{2i}\int d\mathbf{p}' \sum_j \Omega_{ij}\, c_j(\mathbf{p}') \langle \mathbf{p}| \left(e^{i(\mathbf{k}\cdot(\hat{\mathbf{x}}+\mathbf{x}_c)-\omega t+\phi)} + h.c.\right) |\mathbf{p}'\rangle\, e^{i\int_{t_0}^{t} \omega_{ij} + \frac{\mathbf{p}^2}{2m} - \frac{\mathbf{p}'^2}{2m} dt} \tag{2.25}$$

where the Rabi frequency is defined as $\Omega_{ij} \equiv \langle i|(-\hat{\mu}\cdot\mathbf{E}_0)|j\rangle$ and $\omega_{ij} \equiv E_i - E_j$. Now we insert the identity

$$\mathbf{k}\cdot(\hat{\mathbf{x}}+\mathbf{x}_c) - \omega t + \phi = \underbrace{\mathbf{k}\cdot\hat{\mathbf{x}}}_{\text{boost}} + \underbrace{\Big(\mathbf{k}\cdot\mathbf{x}_c(t_0) - \omega t_0 + \phi\Big)}_{\text{laser phase}} + \underbrace{\int_{t_0}^{t}(\mathbf{k}\cdot\dot{\mathbf{x}}_c - \omega)dt}_{\text{Doppler shift}} \tag{2.26}$$

into Eq. (2.25) and perform the integration over $\mathbf{p}'$ using $\langle \mathbf{p}| e^{\pm i\mathbf{k}\cdot\hat{\mathbf{x}}} |\mathbf{p}'\rangle = \langle \mathbf{p}|\mathbf{p}' \pm \mathbf{k}\rangle$:

$$\begin{aligned}\dot{c}_i(\mathbf{p}) = \frac{1}{2i}\sum_j \Omega_{ij} \Big\{ & c_j(\mathbf{p}-\mathbf{k}) e^{i\phi_L} e^{i\int_{t_0}^{t}(\omega_{ij}-\omega+\mathbf{k}\cdot\dot{\mathbf{x}}_c+\frac{\mathbf{k}\cdot\mathbf{p}}{m}-\frac{\mathbf{k}^2}{2m})dt} + \\ & c_j(\mathbf{p}+\mathbf{k}) e^{-i\phi_L} e^{-i\int_{t_0}^{t}(-\omega_{ij}-\omega+\mathbf{k}\cdot\dot{\mathbf{x}}_c+\frac{\mathbf{k}\cdot\mathbf{p}}{m}+\frac{\mathbf{k}^2}{2m})dt} \Big\}\end{aligned} \tag{2.27}$$

where we define the laser phase at point $\{t_0, \mathbf{x}_c(t_0)\}$ as $\phi_L \equiv \mathbf{k}\cdot\mathbf{x}_c(t_0) - \omega t_0 + \phi$. Finally, we impose the two-level constraint ($i = 1, 2$) and consider the coupling between $c_1(\mathbf{p})$ and $c_2(\mathbf{p}+\mathbf{k})$:

$$\begin{aligned}\dot{c}_1(\mathbf{p}) &= \frac{1}{2i}\Omega\, c_2(\mathbf{p}+\mathbf{k}) e^{-i\phi_L} e^{-i\int_{t_0}^{t}\Delta(\mathbf{p})dt} \\ \dot{c}_2(\mathbf{p}+\mathbf{k}) &= \frac{1}{2i}\Omega^* c_1(\mathbf{p}) e^{i\phi_L} e^{i\int_{t_0}^{t}\Delta(\mathbf{p})dt}\end{aligned} \tag{2.28}$$

Here the detuning is $\Delta(\mathbf{p}) \equiv \omega_0 - \omega + \mathbf{k}\cdot(\dot{\mathbf{x}}_c + \frac{\mathbf{p}}{m}) + \frac{\mathbf{k}^2}{2m}$, the Rabi frequency is $\Omega \equiv \Omega_{12} = (\Omega_{21})^*$, and $\omega_0 \equiv \omega_{21} > 0$. In arriving at Eqs. (2.28) we made the rotating wave approximation[13], dropping terms that oscillate at $(\omega_0 + \omega)$ compared to those oscillating at $(\omega_0 - \omega)$. Also, $\Omega_{ii} = 0$ since the $|A_i\rangle$ are eigenstates of parity and $\hat{\mu}$ is odd.

The general solution to (2.28) is

$$\begin{pmatrix} c_1(\mathbf{p},t) \\ c_2(\mathbf{p}+\mathbf{k},t) \end{pmatrix} = \begin{pmatrix} \Lambda_c(\mathbf{p}) e^{-\frac{i}{2}\Delta(\mathbf{p})\tau} & -i\Lambda_s(\mathbf{p}) e^{-\frac{i}{2}\Delta(\mathbf{p})\tau} e^{-i\phi_L} \\ -i\Lambda_s^*(\mathbf{p}) e^{\frac{i}{2}\Delta(\mathbf{p})\tau} e^{i\phi_L} & \Lambda_c^*(\mathbf{p}) e^{\frac{i}{2}\Delta(\mathbf{p})\tau} \end{pmatrix} \begin{pmatrix} c_1(\mathbf{p},t_0) \\ c_2(\mathbf{p}+\mathbf{k},t_0) \end{pmatrix} \tag{2.29}$$

$$\Lambda_c(\mathbf{p}) = \cos\left(\frac{1}{2}\sqrt{\Delta(\mathbf{p})^2 + |\Omega|^2}\,\tau\right) + i\frac{\Delta(\mathbf{p})}{\sqrt{\Delta(\mathbf{p})^2 + |\Omega|^2}}\sin\left(\frac{1}{2}\sqrt{\Delta(\mathbf{p})^2 + |\Omega|^2}\,\tau\right) \tag{2.30}$$

$$\Lambda_s(\mathbf{p}) = \frac{\Omega}{\sqrt{\Delta(\mathbf{p})^2 + |\Omega|^2}}\sin\left(\frac{1}{2}\sqrt{\Delta(\mathbf{p})^2 + |\Omega|^2}\,\tau\right) \tag{2.31}$$

In integrating (2.28) we applied the short pulse limit in the sense of $\mathbf{k}\cdot\ddot{\mathbf{x}}_c\tau^2 \ll 1$, ignoring changes of the atom's velocity during the pulse. For an atom falling in the gravitational field of the Earth, even for pulse times $\tau \sim 10\ \mu$s this term is $\sim kg\tau^2 \sim 10^{-2}$ rad which is non-negligible at our level of required precision. However, for pedagogical reasons we ignore this error here. Corrections due to the finite pulse time are suppressed in the proposed differential measurement between Rb isotopes since we use the same laser to simultaneously manipulate both species (see Section 2.6.1).

For simplicity, from now on we assume the light pulses are on resonance: $\Delta(0) = 0$. We also take the short pulse limit in the sense of $|\Delta(\mathbf{p}) - \Delta(0)|\,\tau \ll 1$ so that we can ignore all detuning systematics. This condition is automatically satisfied experimentally, since only the momentum states that fall within the Doppler width $\sim \tau^{-1}$ of the pulse will interact efficiently with the light.

$$\begin{pmatrix} c_1(\mathbf{p},t) \\ c_2(\mathbf{p}+\mathbf{k},t)\end{pmatrix} = \begin{pmatrix} \Lambda_c & -i\Lambda_s e^{-i\phi_L} \\ -i\Lambda_s^* e^{i\phi_L} & \Lambda_c \end{pmatrix}\begin{pmatrix} c_1(\mathbf{p},t_0) \\ c_2(\mathbf{p}+\mathbf{k},t_0)\end{pmatrix} \qquad \begin{array}{c}\Lambda_c = \cos\frac{|\Omega|\tau}{2} \\ \Lambda_s = \frac{\Omega}{|\Omega|}\sin\frac{|\Omega|\tau}{2}\end{array} \tag{2.32}$$

In the case of a beamsplitter ($\frac{\pi}{2}$ pulse), we choose $|\Omega|\,\tau = \frac{\pi}{2}$, whereas for a mirror ($\pi$ pulse) we set $|\Omega|\,\tau = \pi$:

$$\Lambda_{\pi/2} = \begin{pmatrix} \frac{1}{\sqrt{2}} & \frac{-i}{\sqrt{2}}\,e^{-i\phi_L} \\ \frac{-i}{\sqrt{2}}\,e^{i\phi_L} & \frac{1}{\sqrt{2}} \end{pmatrix} \qquad \Lambda_\pi = \begin{pmatrix} 0 & -i\,e^{-i\phi_L} \\ -i\,e^{i\phi_L} & 0 \end{pmatrix} \tag{2.33}$$

These matrices encode the rules for the imprinting of the laser's phase on the atom: the component of the atom $c_1(\mathbf{p},t_0)$ that gains momentum from the light (absorbs a photon) picks up a phase $+\phi_L$, and the component of the atom $c_2(\mathbf{p}+\mathbf{k},t_0)$ that loses momentum to the light (emits a photon) picks up a phase $-\phi_L$. Symbolically,

$$|\mathbf{p}\rangle \quad\rightarrow\quad |\mathbf{p}+\mathbf{k}\rangle\, e^{i\phi_L} \tag{2.34}$$

$$|\mathbf{p}+\mathbf{k}\rangle \quad\rightarrow\quad |\mathbf{p}\rangle\, e^{-i\phi_L} \tag{2.35}$$

As a result, the total laser phase shift is

$$\Delta\phi_{\text{laser}} = \left(\sum_j \pm\phi_L(t_j, \mathbf{x}_u(t_j))\right)_{\text{upper}} - \left(\sum_j \pm\phi_L(t_j, \mathbf{x}_l(t_j))\right)_{\text{lower}} \tag{2.36}$$

where the sums are over all of the atom–laser interaction points $\{t_j, \mathbf{x}_u(t_j)\}$ and $\{t_j, \mathbf{x}_l(t_j)\}$ along the upper and lower arms, respectively, and the sign is determined by Eqs. (2.34) and (2.35).

The final contribution to $\Delta\phi_{\text{tot}}$ is the separation phase, $\Delta\phi_{\text{separation}}$. As shown in Fig. 2.1, this shift arises because the endpoints of the two arms of the interferometer need not coincide at the time of the final beamsplitter. To derive the expression for separation phase, we write the state of the atom at time $t = t_0 + \tau$ just after the final beamsplitter pulse as

$$|\Psi(t)\rangle = |\Psi_u(t)\rangle + |\Psi_l(t)\rangle \tag{2.37}$$

where $|\Psi_u(t)\rangle$ and $|\Psi_l(t)\rangle$ are the components of the final state that originate from the upper and lower arms, respectively. Just before the final beamsplitter pulse is applied, we write the state of each arm as

$$\begin{aligned} |\Psi_u(t_0)\rangle &= \int d\mathbf{p}\, c_1(\mathbf{p}, t_0)\hat{G}_u(t_0)\, |\mathbf{p}\rangle\, |A_1\rangle\, e^{i\theta_u} && (2.38) \\ |\Psi_l(t_0)\rangle &= \int d\mathbf{p}\, c_2(\mathbf{p}, t_0)\hat{G}_l(t_0)\, |\mathbf{p}\rangle\, |A_2\rangle\, e^{i\theta_l} && (2.39) \end{aligned}$$

where $\hat{G}_u \equiv \hat{G}(\mathbf{x}_u, \mathbf{p}_u, L_u)$ and $\hat{G}_l \equiv \hat{G}(\mathbf{x}_l, \mathbf{p}_l, L_l)$ are the Galilean transformation operators for the upper and lower arm, respectively. These operators translate each wavepacket in phase space to the appropriate position ($\mathbf{x}_u$ or $\mathbf{x}_l$) and momentum ($\mathbf{p}_u$ or $\mathbf{p}_l$). Here we have assumed for clarity that prior to the final beamsplitter the upper and lower arms are in internal states $|A_1\rangle$ and $|A_2\rangle$ with amplitudes $c_1(\mathbf{p}, t_0)$ and $c_2(\mathbf{p}, t_0)$, respectively; identical results are obtained in the reversed case. We have also explicitly factored out the dynamical phases θ_u and θ_l accumulated along the upper and lower arms, respectively, which contain by definition all contributions to laser phase and propagation phase acquired prior to the final beamsplitter.

We write the wavefunction components after the beamsplitter in the form of Eq. (2.21):

$$|\Psi_u(t)\rangle = \int d\mathbf{p} \sum_i c_i^{(u)}(\mathbf{p}, t)\hat{G}_u\, |\mathbf{p}\rangle\, |A_i\rangle \tag{2.40}$$

$$|\Psi_l(t)\rangle = \int d\mathbf{p} \sum_i c_i^{(l)}(\mathbf{p}, t)\hat{G}_l\, |\mathbf{p}\rangle\, |A_i\rangle \tag{2.41}$$

where we invoked the short pulse limit so that $e^{-i\hat{H}_2\tau} \approx 1$. Next we time evolve the states using Eq. (2.32) assuming a perfect $\frac{\pi}{2}$ pulse and using the initial conditions given in Eqs. (2.38) and (2.39): namely, $c_1^{(u)}(\mathbf{p}, t_0) = c_1(\mathbf{p}, t_0)e^{i\theta_u}$ and $c_2^{(u)}(\mathbf{p}, t_0) = 0$ for the upper arm and $c_1^{(l)}(\mathbf{p}, t_0) = 0$ and $c_2^{(l)}(\mathbf{p}, t_0) = c_2(\mathbf{p}, t_0)e^{i\theta_l}$ for the lower arm.

$$|\Psi_u(t)\rangle = \int d\mathbf{p}\, c_1(\mathbf{p}, t_0) \left\{ \frac{1}{\sqrt{2}}\hat{G}_u\, |\mathbf{p}\rangle\, |A_1\rangle + \frac{-i}{\sqrt{2}} e^{i\phi_L(\mathbf{x}_u)}\hat{G}_u\, |\mathbf{p} + \mathbf{k}\rangle\, |A_2\rangle \right\} e^{i\theta_u} \tag{2.42}$$

$$|\Psi_l(t)\rangle = \int d\mathbf{p}\, c_2(\mathbf{p} + \mathbf{k}, t_0) \left\{ \frac{-i}{\sqrt{2}} e^{-i\phi_L(\mathbf{x}_l)}\hat{G}_l\, |\mathbf{p}\rangle\, |A_1\rangle + \frac{1}{\sqrt{2}}\hat{G}_l\, |\mathbf{p} + \mathbf{k}\rangle\, |A_2\rangle \right\} e^{i\theta_l} \tag{2.43}$$

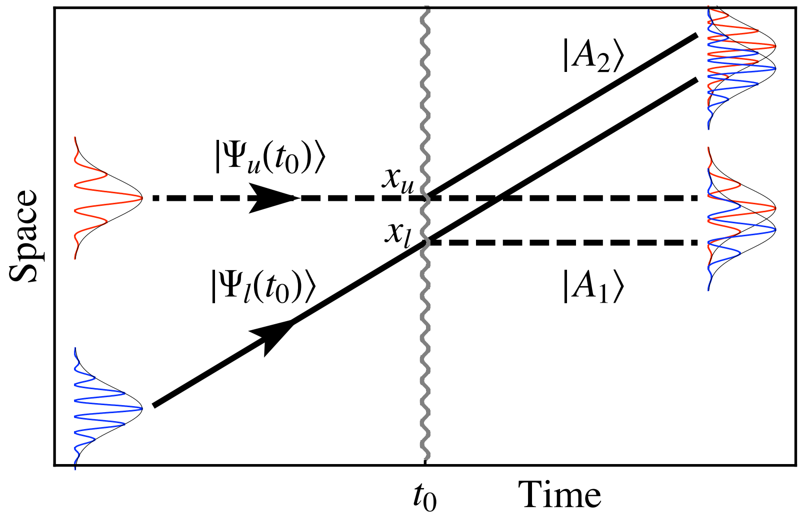

Figure 2.1: Separation Phase. This is a magnified view of the end of the interferometer which shows the upper and lower arms converging at the final beamsplitter at time t_0, and the resulting interference. The dashed and solid lines designate the components of the wavefunction in internal states $|A_1\rangle$ and $|A_2\rangle$, respectively. After the beamsplitter, each output port consists of a superposition of wavepackets from the upper and lower arm. Any offset $\Delta\mathbf{x} = \mathbf{x}_l - \mathbf{x}_u$ between the centers of the wavepacket contributions to a given output port results in a separation phase shift.

We now project into position space and perform the $\mathbf{p}$ integrals,

$$\langle \mathbf{x}|\Psi_u(t)\rangle = \frac{c_1(\mathbf{x}-\mathbf{x}_u,t_0)}{\sqrt{2}}\left\{e^{i\mathbf{p}_u\cdot(\mathbf{x}-\mathbf{x}_u)}\,|A_1\rangle - ie^{i\phi_L(\mathbf{x}_u)}e^{i(\mathbf{p}_u+\mathbf{k})\cdot(\mathbf{x}-\mathbf{x}_u)}\,|A_2\rangle\right\}e^{i\theta_u} \quad (2.44)$$

$$\langle \mathbf{x}|\Psi_l(t)\rangle = \frac{c_2(\mathbf{x}-\mathbf{x}_l,t_0)}{\sqrt{2}}\left\{-ie^{-i\phi_L(\mathbf{x}_l)}e^{i(\mathbf{p}_l-\mathbf{k})\cdot(\mathbf{x}-\mathbf{x}_l)}\,|A_1\rangle + e^{i\mathbf{p}_l\cdot(\mathbf{x}-\mathbf{x}_l)}\,|A_2\rangle\right\}e^{i\theta_l} \quad (2.45)$$

where we identified the Fourier transformed amplitudes using $c_i(\mathbf{x}-\mathbf{x}_c,t_0) = \int d\mathbf{p}\,\langle \mathbf{x}-\mathbf{x}_c|\mathbf{p}\rangle\,c_i(\mathbf{p},t_0)$. The resulting interference pattern in position space is therefore

$$\begin{aligned}
\langle \mathbf{x}|\Psi(t)\rangle &= \langle \mathbf{x}|\Psi_u(t)\rangle + \langle \mathbf{x}|\Psi_l(t)\rangle \\
&= \frac{1}{\sqrt{2}}\,|A_1\rangle\left\{c_1(\mathbf{x}-\mathbf{x}_u,t_0)e^{i\theta_u}e^{i\mathbf{p}_u\cdot(\mathbf{x}-\mathbf{x}_u)} - i\,c_2(\mathbf{x}-\mathbf{x}_l,t_0)e^{i\theta_l}e^{-i\phi_L(\mathbf{x}_l)}e^{i(\mathbf{p}_l-\mathbf{k})\cdot(\mathbf{x}-\mathbf{x}_l)}\right\} \\
&+ \frac{1}{\sqrt{2}}\,|A_2\rangle\left\{c_2(\mathbf{x}-\mathbf{x}_l,t_0)e^{i\theta_l}e^{i\mathbf{p}_l\cdot(\mathbf{x}-\mathbf{x}_l)} - i\,c_1(\mathbf{x}-\mathbf{x}_u,t_0)e^{i\theta_u}e^{i\phi_L(\mathbf{x}_u)}e^{i(\mathbf{p}_u+\mathbf{k})\cdot(\mathbf{x}-\mathbf{x}_u)}\right\}
\end{aligned}$$

The probability of finding the atom in either output port $|A_1\rangle$ or $|A_2\rangle$ is

$$|\langle A_1|\,\langle \mathbf{x}|\Psi(t)\rangle|^2 \quad = \quad \frac{|c_1|^2+|c_2|^2}{2}+\frac{1}{2}\left(i\,c_1\,c_2^*\,e^{i\Delta\phi_1}+c.c.\right) \tag{2.46}$$

$$|\langle A_2|\,\langle \mathbf{x}|\Psi(t)\rangle|^2 \quad = \quad \frac{|c_1|^2+|c_2|^2}{2}-\frac{1}{2}\left(i\,c_1\,c_2^*\,e^{i\Delta\phi_2}+c.c.\right) \tag{2.47}$$

with $c_1 \equiv c_1(\mathbf{x}-\mathbf{x}_u, t_0)$ and $c_2 \equiv c_2(\mathbf{x}-\mathbf{x}_l, t_0)$. For the total phase shift we find

$$\Delta\phi_1 \quad \equiv \quad \Big\{\theta_u + \mathbf{p}_u\cdot(\mathbf{x}-\mathbf{x}_u)\Big\} - \Big\{\theta_l - \phi_L(\mathbf{x}_l) + (\mathbf{p}_l-\mathbf{k})\cdot(\mathbf{x}-\mathbf{x}_l)\Big\} \tag{2.48}$$

$$= \quad \underbrace{\theta_u - \Big(\theta_l - \phi_L(\mathbf{x}_l)\Big)}_{\Delta\phi_{\text{propagation},1}\;+\;\Delta\phi_{\text{laser},1}} \;+\; \underbrace{\bar{\mathbf{p}}_1\cdot\Delta\mathbf{x}}_{\Delta\phi_{\text{separation},1}} \;+\Delta\mathbf{p}\cdot(\mathbf{x}-\bar{\mathbf{x}}) \tag{2.49}$$

and

$$\Delta\phi_2 \quad \equiv \quad \Big\{\theta_u + \phi_L(\mathbf{x}_u) + (\mathbf{p}_u+\mathbf{k})\cdot(\mathbf{x}-\mathbf{x}_u)\Big\} - \Big\{\theta_l + \mathbf{p}_l\cdot(\mathbf{x}-\mathbf{x}_l)\Big\} \tag{2.50}$$

$$= \quad \underbrace{\Big(\theta_u + \phi_L(\mathbf{x}_u)\Big) - \theta_l}_{\Delta\phi_{\text{propagation},2}\;+\;\Delta\phi_{\text{laser},2}} \;+\; \underbrace{\bar{\mathbf{p}}_2\cdot\Delta\mathbf{x}}_{\Delta\phi_{\text{separation},2}} \;+\Delta\mathbf{p}\cdot(\mathbf{x}-\bar{\mathbf{x}}) \tag{2.51}$$

where $\bar{\mathbf{p}}_1 = \frac{\mathbf{p}_u+(\mathbf{p}_l-\mathbf{k})}{2}$ and $\bar{\mathbf{p}}_2 = \frac{(\mathbf{p}_u+\mathbf{k})+\mathbf{p}_l}{2}$ are the average momenta in the $|A_1\rangle$ (slow) and $|A_2\rangle$ (fast) output ports, respectively. In general, the separation phase is

$$\Delta\phi_{\text{separation}} = \bar{\mathbf{p}}\cdot\Delta\mathbf{x} \tag{2.52}$$

which depends on the separation $\Delta\mathbf{x} \equiv \mathbf{x}_l - \mathbf{x}_u$ between the centers of the wavepackets from each arm as well as the average canonical momentum $\bar{\mathbf{p}}$ in the output port. We point out that even though the definitions (2.48) and (2.50) use the same sign convention as our previous expressions for laser (2.36) and propagation (2.20) phase in the sense of $(\;)_{\text{upper}} - (\;)_{\text{lower}}$, the separation vector $\Delta\mathbf{x}$ is defined as $(\mathbf{x})_{\text{lower}} - (\mathbf{x})_{\text{upper}}$.

Notice that the phase shift expressions (2.49) and (2.51) contain a position dependent piece $\Delta\mathbf{p}\cdot(\mathbf{x}-\bar{\mathbf{x}})$, where $\bar{\mathbf{x}} \equiv \frac{\mathbf{x}_u+\mathbf{x}_l}{2}$ and $\Delta\mathbf{p} = (\mathbf{p}_u+\mathbf{k})-\mathbf{p}_l = \mathbf{p}_u-(\mathbf{p}_l-\mathbf{k})$, owing to the fact that the contributions from each arm may have different momenta after the last beamsplitter. Typically this momentum difference is very small, so the resulting phase variation has a wavelength that is large compared to the spatial extent of the wavefunction. Furthermore, this effect vanishes completely in the case of spatially averaged detection over a symmetric wavefunction.

Finally, we show that the total phase shifts $\Delta\phi_1$ and $\Delta\phi_2$ for the two output ports are actually equal, as required by conservation of probability. According to Eqs. (2.49) and (2.51), the contributions to the total phase differ in the following ways:

$$\begin{aligned}\Big(\Delta\phi_{\text{propagation},1} + \Delta\phi_{\text{laser},1}\Big) - \Big(\Delta\phi_{\text{propagation},2} + \Delta\phi_{\text{laser},2}\Big) &= \phi_L(\mathbf{x}_l) - \phi_L(\mathbf{x}_u) \\ &= \mathbf{k}\cdot(\mathbf{x}_l - \mathbf{x}_u) = \mathbf{k}\cdot\Delta\mathbf{x}\end{aligned}$$

$$\Delta\phi_{\text{separation},1} - \Delta\phi_{\text{separation},2} = \bar{\mathbf{p}}_1\cdot\Delta\mathbf{x} - \bar{\mathbf{p}}_2\cdot\Delta\mathbf{x} = -\mathbf{k}\cdot\Delta\mathbf{x}$$

Together these results imply that $\Delta\phi_1 = \Delta\phi_2$ and prove that the total interferometer phase shift $\Delta\phi_{\text{tot}}$ is independent of the output port.

The accuracy of the above formalism is dependent on the applicability of the aforementioned stationary phase approximation as well as the short pulse limit. The stationary phase approximation breaks down when the external Hamiltonian varies rapidly compared to the phase space width of the atom wavepacket. The short pulse limit requires that the atom's velocity not change appreciably during the duration of the atom-light interaction. Both approximations are justified to a large degree for a typical light pulse atom interferometer, but in the most extreme high precision applications such as we consider here, important corrections are present. However, we emphasize that these errors due to finite pulse duration and wavepacket size are well-known, previously established backgrounds.

2.5 Applications in inertial navigation

The navigation problem is easily stated: How do we determine a platform's trajectory as a function of time? In the 20th century, solutions to the problem have led to the development of exquisitely refined hardware systems and navigation algorithms. Today we take for granted that a hand-held GPS reciever can be used to obtain meter level position determination. When GPS is unavailable (for example, when satellites are not in view), position determination becomes much less accurate. In this case, stand-alone "black-box" inertial navigation systems, comprised of a suite of gyroscope and acclerometers, are used to infer position changes by integrating the outputs of these inertial force sensors. State-of-the-art commercial grade navigation systems have position drift errors of kilometers per hour of navigation time, significantly worse than the GPS solution. How can we close the performance gap between GPS and inertial systems? The way forward is improved instrumentation: better gyroscopes and accelerometers.

The light-pulse interferometry method is well suited to inertial applications. As shown in Section 2.4.1, the phase shift for the light pulse interferometer consists of contributions from path phases, optical interactions, and separation phases. However, for the sensitivity range of interest, the optical phase shifts dominate, and therefore $\Delta\phi_{tot} \approx \Delta\phi_{laser}$. In this case there is a straightforward interpretation for the operation of the sensor: the sensor registers the time evolution of the relative position of the atomic wavepackets with respect to the sensor case (defined by the opto-mechanical hardware for the laser beams) using optical interferometry. Since distances are measured in terms of the wavelength of light, and the atom is in a benign environment (spurious forces, such as from magnetic field gradients, are extremely small – below $10^{-10}g$), the sensors are characterized by superbly stable, low-noise operation.

The phase shift between the two paths is inferred by measuring the probability of finding the atoms in a given output port. Since we are only concerned with the optical phase shift, we can calculate the sensor output using Eq. (2.3). We see that for a standard $\pi/2 - \pi - \pi/2$ excitation sequence [1], $\Delta\phi_{laser} = \mathbf{k}_1 \cdot (\mathbf{x}_1^u - \mathbf{x}_1^0) - \mathbf{k}_2 \cdot (\mathbf{x}_2^u - \mathbf{x}_2^0) - \mathbf{k}_2 \cdot (\mathbf{x}_2^l - \mathbf{x}_2^0) + \mathbf{k}_3 \cdot (\mathbf{x}_3^l - \mathbf{x}_3^0)$. Here the subscript indexes each of the three successive optical interactions, $\mathbf{x}_i^u$ and $\mathbf{x}_i^l$ are the semi-classical positions of the atom along the upper and lower paths, respectively, at the time of each interaction in a non-rotating, inertial coordinate system, $\mathbf{x}_i^0$ is the phase reference for the optical fields, and $\mathbf{k}_i$ is the propagation vector of the laser field associated with each pulse.

2.5.1 Gyroscope

Assuming the atoms have initial velocity $\mathbf{v}_0$ and the effective Raman propagation vectors initially have common orientations $\mathbf{k}_0$ which rotate with angular rate $\mathbf{\Omega}$, it is straightforward to show that $\Delta\phi_{laser} \approx \mathbf{k}_0 \cdot (2\mathbf{v}_0 \times \mathbf{\Omega})T^2$. This configuration is well-suited to precision measurements of platform rotations. This expression can be put in a form analogous to the Sagnac shift for optical interferometers by noting that for small rotation rates Ω, $\mathbf{A} = \frac{\hbar}{m}(\mathbf{k_0} \times \mathbf{v_0})T^2$ is the area enclosed by the interfering paths. Thus this shift can also be written in the Sagnac form $\phi = \frac{2m}{\hbar}\mathbf{\Omega} \cdot \mathbf{A}$ – proportional to the product of the enclosed area and rotation rate.

Gyroscopes built on this principle have achieved performance levels in the laboratory which compare favorably with state-of-the art gyroscopes, as shown in Figs. 2.2, 2.3, 2.4 and 2.5. Key figures of merit for gyroscope performance include gyroscope noise (often referred to as angle random walk), bias stability (stability of output for a null input) and scale factor stability (stability of the multiplier between the input rotation rate and output phase shift). The laboratory gyroscope illustrated below has achieved a demonstrated angle random walk of 3 μdeg/hr$^{1/2}$, bias stability of 60 μdeg/hr (upper limit) and scale factor stability of 5 ppm (upper limit). Key drivers in the stability of the gyroscope outputs are the stability of the intensities used to drive the Raman transitions, and alignment stability of the Raman beam optical paths. Non-inertial phase shifts associated with magnetic field inhomogeneities and spurious AC Stark shifts are nulled using a case reversal technical where the propagation directions of the Raman beams are periodically reversed using electo-optic methods.

2.5.2 Accelerometer

If the platform containing the laser beams accelerates, or if the atoms are subject to a gravitational acceleration, the laser phase shift then contains acceleration terms $\Delta\phi = \mathbf{k}_0 \cdot \mathbf{a}T^2$. For a stationary interferometer, with the laser beams vertically directed, this phase shift measures the acceleration

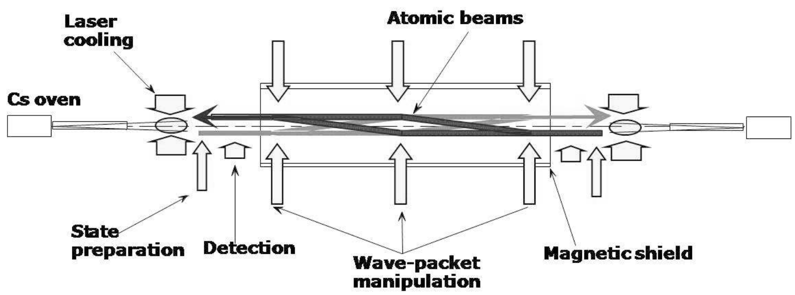

Figure 2.2: Schematic illustration of an atomic beam gyroscope.

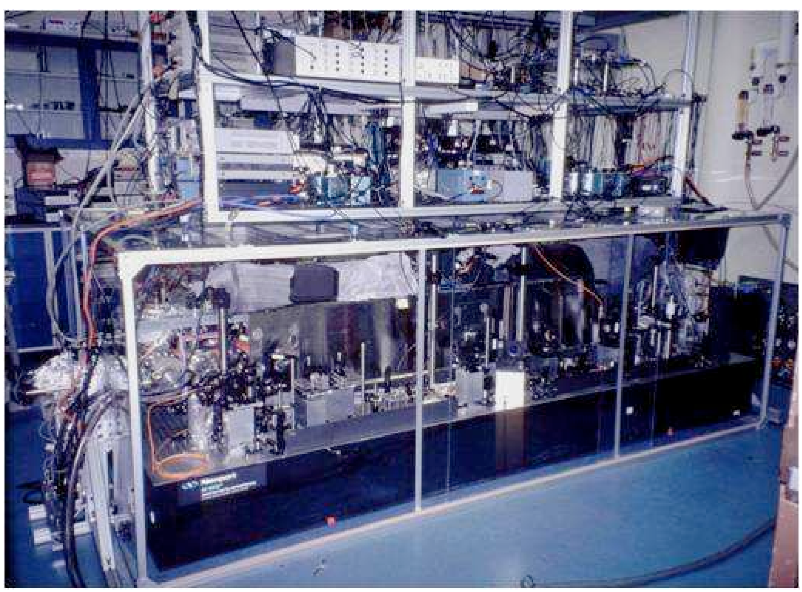

Figure 2.3: Photograph of gyroscope.

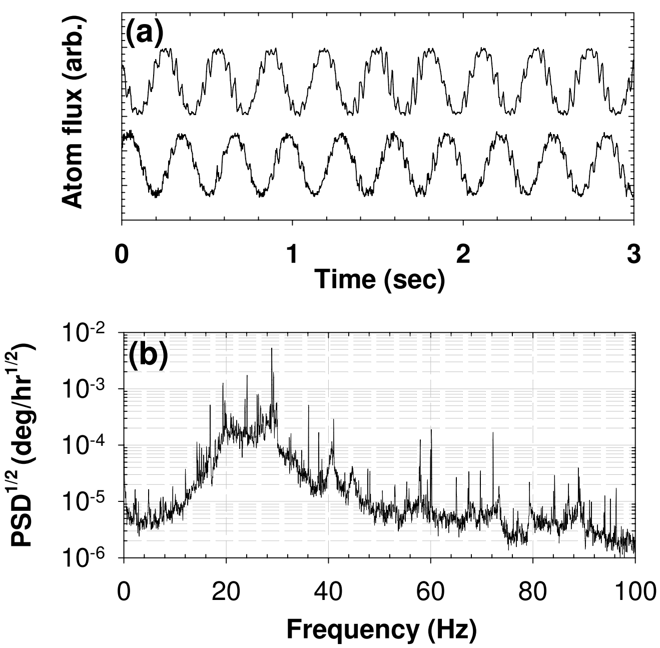

Figure 2.4: Gyroscope interference fringes (a) and power spectral density of gyroscope noise (b). The increase in noise in the 10-50 Hz band is due to technical rotation noise sources in the building (the gyroscope was mounted directly to the building floor). See Ref. [14].

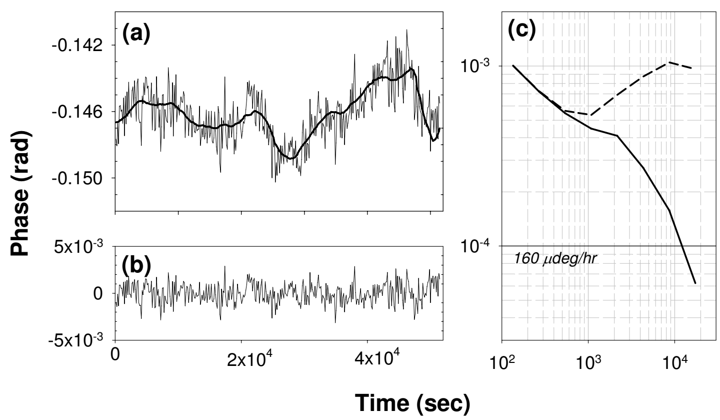

Figure 2.5: (a) Drift in case reversed gyroscope output with time and its correlation with instrument temperature (solid black). (b) Gyro noise output residual after subtracting temperature model. (c) Allan deviation of gyro output before and after temperature compensation (see Ref. [14]).

due to gravity. Remarkably, part per billion level agreement has been demonstrated between the output of an atom interferometer gravimeter and a conventional, "falling-corner-cube" gravimeter. [15]. In the future, compact, geophysical (10^{-8} g accuracy) grade instruments should enable low cost gravity field surveys (see Fig. 2.6). For this type of instrument, laser cooling methods are used to initially prepare ensembles of roughly 10^7 atoms at kinetic temperatures approaching 1 μK. At these low temperatures, the rms velocity spread of the atomic ensembles is a few cm/sec. Cold atom ensembles are then launched on ballistic trajectories. In this configuration, the time between laser pulses exceeds 100 msec, which means wavepackets separate by roughly 1 mm over the course of the interferometer sequence. The phase shift is read-out by detecting the number of atoms in each final state using resonance fluorescence and normalized detection methods.

In general, both rotation and acceleration terms are present in the sensor outputs. For navigation applications, the rotation response needs to be isolated from the acceleration response. In practice, this is accomplished by using multiple atom sources and laser beam propagation axes. For example, for the gyroscope illustrated in Fig. 2.2, counter-propagating atom beams are used to isolate rotation induced phase shifts from acceleration induced shifts. It is interesting to note that the same apparatus is capable of simultaneous rotation and acceleration outputs – a significant benefit for navigation applications which require simultaneous output of rotation rate and acceleration for three mutually orthogonal axes. Since gyroscope and accelerometer operation rest on common principles and common hardware implementations, integration of sensors into a full inertial base is straightforward. Of course, particular hardware implementations depend on the navigation platform

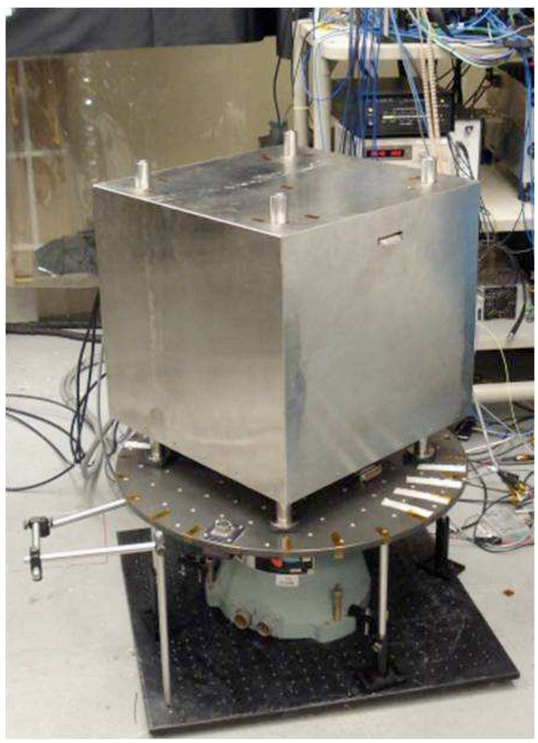

Figure 2.6: Compact accelerometer. This instrument supports accelerometer, gyroscope and gravity gradiometer operation modes. The compact accelerometer has demonstrated microGal sensitivity.

(e.g. ship, plane, land vehicle) and trajectory dynamics.

2.5.3 Gravity gradiometer

There is an additional complication in navigation system architecture for high accuracy navigation applications: the so-called "problem of the vertical." Terrestrial navigation requires determining platform position in the gravity field of the Earth. Due to the Equivalence Principle, navigation system accelerometers do not distinguish between the acceleration due to gravity and platform acceleration. So in order to determine platform trajectory in an Earth-fixed coordinate system, the local acceleration due to gravity needs to be subtracted from accelerometer output in order to determine the acceleration of the vehicle with respect to the Earth. This means that the local acceleration due to gravity needs to be independently known. For example, existing navigation systems use a gravity map to make this compensation. However, in present systems, this map does not have enough resolution or accuracy for meter-level position determination. To give a feeling for orders of magnitude, a 10^{-7} error in knowledge of the local acceleration due to gravity integrates to meter-level position errors in 1 hour.

There are at least two paths forward: 1) better maps or 2) on-the-fly gravity field determination. Improved maps can be obtained with more precise surveys. On-the-fly determination seems impossible, due to the Equivalence Principle (since platform accelerations cannot be discriminated from the acceleration due to gravity). However, the outputs from a gravity gradiometer – an instrument which measures changes in the acceleration due to gravity over fixed baselines – can be used for this purpose. The idea is to integrate the gravity gradient over the inferred trajectory to determine gravity as a function of position. In principle, such an instrument can function on a moving platform, since platform accelerations cancel as a common mode when the output from spatially separated accelerometers are differenced to obtain the gradient. In practice, such a strategy places hard requirements on the stability of the component accelerometers: their responses need to be matched to an exceptional degree in order to discriminate gravity gradient induced accelerations (typically below 10^{-9} g) from other sensor error sources.

Due to the stability of their acceleration outputs, a pair of light pulse accelerometers is well-suited to gravity gradient instrumentation. The basic idea is to simultaneously create two spatially separated interferometers using a common laser beam. In this way, technical acceleration noise of the measurement platform is a common-mode noise source which leads to near identical phase shifts in each accelerometer. On the other hand, a gravity gradient along the measurement axis results in a residual differential phase shift. This configuration has been used to measure the gravity gradient of the Earth, as well as the gravity gradient associated with nearby mass distributions, as illustrated in Fig. 2.7. Laboratory gravity gradiometers have achieved resolutions below 1 E (where 1 E = 10^{-9} sec^{-2}). This configuration has also been used to measure the Newtonian constant of gravity G [16, 17, 18]. In future navigation systems, an ensemble of accelerometers, configured along

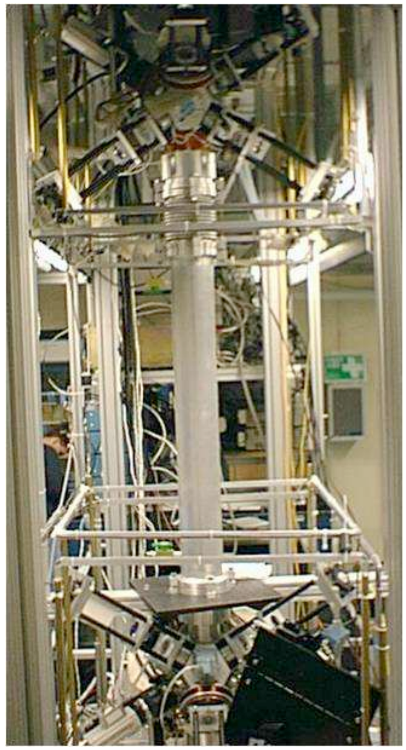

Figure 2.7: Vertical axis gravity gradiometer. Two atom interferometer accelerometers separated along a vertical measurement axis are simultaneously interrogated by a common set of Raman laser beams. This apparatus was used to measure G, Newton's constant [16].

independent measurement axes, could acquire the full gravity gradient tensor.

In addition to their role as navigation aids, gravity gradiometers have applications in geodesy and oil/mineral exploration. The idea here is that mass/density anomalies associated with interesting geophysical features (such as kimberlite pipes – in the case of diamond exploration – or salt domes – in the case of oil exploration) manifest as gravity anomalies. In some cases, these anomalies can be pronounced enough to be detected by a gravity gradiometer from an airborne platform. Atom-based gravity gradiometers appear to have competitive performance figures of merit for these applications as compared with existing technologies. For these applications, the central design challenge is realization of an instrument which has very good noise performance, but also is capable of sustained operation on a moving platform. Figs. 2.8, 2.9 illustrate a system currently under development for this purpose.

Figure 2.8: Gravity gradiometer with horizontal measurement axis. Raman beams propagate along the axis defined by the tube connecting each accelerometer housing.

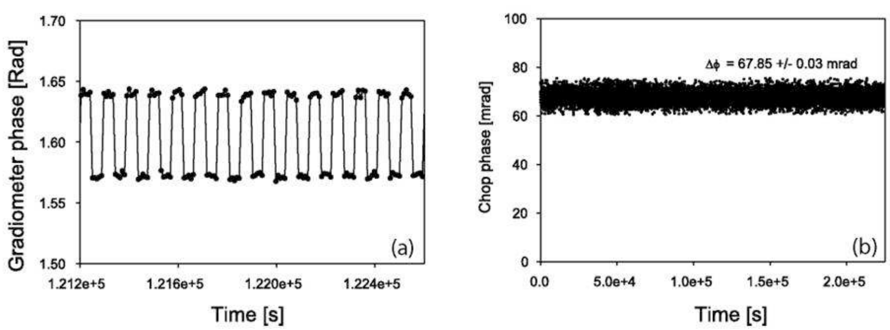

Figure 2.9: Gravity gradiometer response to a proof mass which is periodically brought in close proximity to one of the accelerometer regions. For the proof masses used for this demonstration, the apparatus is capable of resolving changes in G at the 3×10^{-4} level [19].

2.6 Application to tests of the Equivalence Principle

Precision tests of the Equivalence Principle (EP) promise to provide insight into fundamental physics. Since the EP is one of the central axioms of general relativity (GR), these experiments are powerful checks of gravity and can tightly constrain new theories. Furthermore, EP experiments test for hypothetical fifth forces since many examples of new forces are EP–violating[20].

The Equivalence Principle has several forms, with varying degrees of universality. Here we consider tests of the Weak Equivalence Principle, which can be stated as follows: the motion of a body in a gravitational field in any local region of space-time is indistinguishable from its motion in a uniformly accelerated frame. This implies that the body's inertial mass is equal to its gravitational mass, and that all bodies locally fall at the same rate under gravity, independent of their mass or composition.

The results of EP experiments are typically expressed in terms of the Eötvös parameter $\eta = \Delta a/\bar{a}$, where Δa is the EP violating differential acceleration between the two test bodies and $\bar{a}$ is their average acceleration [21]. Currently, two conceptually different experiments set the best limits on the Equivalence Principle. Lunar Laser Ranging (LLR), which tests the EP by comparing the acceleration of the Earth and Moon as they fall toward the Sun, limits EP violation at $\eta = (-1.0 \pm 1.4) \times 10^{-13}$ [22]. Recently, the Eöt-Wash group has set a limit of $\eta = (0.3 \pm 1.8) \times 10^{-13}$ using an Earth-based torsion pendulum apparatus[23]. Several proposed satellite missions aim to improve on these limits by observing the motion of macroscopic test bodies in orbit around the Earth[24, 25]. Here we discuss our effort to perform a ground-based EP test using individual atoms with a goal of measuring $\eta \sim 10^{-15}$. Instead of macroscopic test masses, we compare the simultaneous acceleration under gravity of freely-falling cold atom clouds of ^{87}Rb and ^{85}Rb using light-pulse atom interferometry [1].

Light-pulse atom interferometers have already been used to make extremely accurate inertial force measurements in a variety of configurations, including gyroscopes, gradiometers, and gravimeters. For example, the local gravitational acceleration g of freely-falling Cs atoms was measured with an accuracy $\sim 10^{-9}g$ [15]. Gravity gradiometers have been used to suppress noise as well as many systematic errors that are present in absolute g measurements by comparing the acceleration of two displaced samples of atoms. A differential measurement of this kind was used to measure the Newtonian constant of gravity G with an accuracy of $\sim 3 \times 10^{-3}G$ [16]. The EP measurement we describe here benefits from an analogous differential measurements strategy, where in this case the common-mode noise suppression arises from a comparison between two co-located isotopes of different mass, rather than between spatially separated atoms as in a traditional gradiometer.

2.6.1 Proposed experiment overview

The proposed experiment ideally consists of simultaneously observing the free-fall motion of the two Rb isotopes in the absence of all non-gravitational forces. To this end, the measurement is performed inside a 10 cm diameter by 8.8 m long cylindrical ultra high vacuum chamber. To maximize their free-fall time, the atoms are launched in a vertical fountain geometry from the bottom of the chamber. Light-pulse atom interferometry is performed while the atoms are in free-fall, and the resulting phase shift is sensitive to the atoms' acceleration. Figure 2.10 is a space-time diagram depicting the trajectories that each atom follows during the free-fall interferometry sequence.

To maximize the cancellation of spurious effects, both the ^{87}Rb and ^{85}Rb atom clouds are launched at the same time and are made to follow the same trajectories as closely as possible. We launch both isotopes from the same magnetic trap in order to minimize any differences between their initial positions and velocities. As a result of the small isotope shift between ^{87}Rb and ^{85}Rb, we are able to use the same laser pulses to simultaneously manipulate them during the interferometer sequence (see Fig. 2.10). Using the same laser makes the apparatus insensitive to pulse timing jitter and dramatically reduces the phase noise stability requirements of the lasers.

A single measurement of acceleration in our atom interferometer consists of three steps: atom cloud preparation, interferometer pulse sequence, and detection. In the first step a sub-microkelvin cloud of $\sim 10^7$ atoms is formed using laser cooling and evaporative cooling in a TOP trap [26]. This dilute ensemble of cold atoms is then launched vertically with velocity $v_z \sim 13$ m/s by transferring momentum from laser light using an accelerated optical lattice potential [27]. This technique allows for precise control of the launch velocity and, because it is a coherent process, it avoids heating the cloud via spontaneous emission.

In the second phase of the measurement, the atoms follow free-fall trajectories and the interferometry is performed. A sequence of laser pulses serve as beamsplitters and mirrors that coherently divide each atom's wavepacket and then later recombine it to produce interference (see Fig. 2.10). The atom beamsplitter is typically implemented using a stimulated two-photon process (Raman or Bragg transitions), resulting in a net momentum transfer of $\mathbf{k}_{\text{eff}} = \mathbf{k_2} - \mathbf{k_1} \approx 2\mathbf{k_2}$ at each interaction. Since the acceleration sensitivity of the interferometer is proportional to the effective momentum $\hbar k_{\text{eff}}$ transferred to the atom during interactions with the laser, we intend to take advantage of more sophisticated atom optics. Large momentum transfer (LMT) beamsplitters with $\hbar k_{\text{eff}} = 24\hbar k$ have been demonstrated [28], and up to $100\hbar k$ may be possible. Promising LMT beamsplitter candidates include optical lattice manipulations [27], sequences of Raman pulses [29] and adiabatic passage methods [30].

The third and final step of each acceleration measurement is atom detection. At the end of the interferometer sequence, each atom is in a superposition of the two output velocity states, as shown by the diverging paths on the right in Fig. 2.10. These two final velocity states are directly analogous to the two output ports of a Mach-Zehnder light interferometer after the final

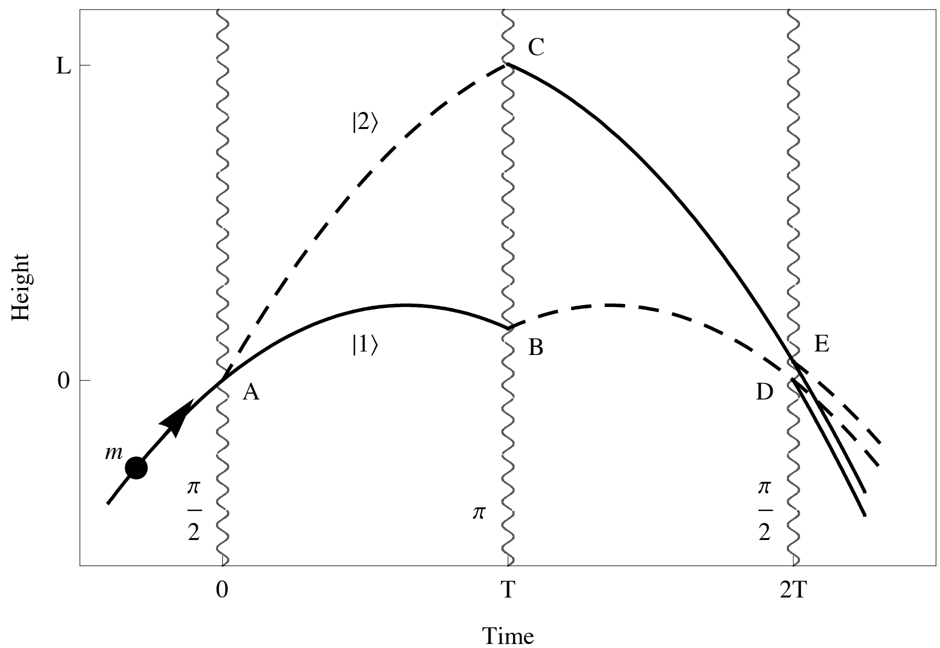

Figure 2.10: Space–time diagram for a single atom of mass m during the interferometer pulse sequence. The atom is launched with velocity v_z from the bottom of the vacuum system. At time $t = 0$, a $\frac{\pi}{2}$ (beamsplitter) pulse is applied to coherently divide the atom wavefunction. After a time T, a π (mirror) pulse is applied that reverses the relative velocity between the wavefunction components. A final $\frac{\pi}{2}$ (beamsplitter) pulse at time $2T$ results in interference between the two space–time paths. The interferometer phase shift is inferred by measuring the probability of detecting the atom in either state $|1\rangle$ (solid line) or state $|2\rangle$ (dashed line). Note that points D and E are in general spatially separated in the presence of non-uniform forces, leading to a separation phase shift.

recombining beamsplitter. As with a light interferometer, the probability that an atom will be found in a particular output port depends on the relative phase acquired along the two paths of the atom interferometer. Since the output states differ in velocity by $\sim \hbar k_{\text{eff}}/m$, they spatially separate over time. After an appropriate drift time, the two velocity groups can be separately resolved, and the populations can be then measured by fluorescence imaging.

We now consider the expected sensitivity of our differential ^{87}Rb–^{85}Rb accelerometer. Recent atom interferometers have demonstrated sensor noise levels limited only by the quantum projection noise of the atoms (atom shot noise) [31]. Assuming a time–average atom flux of $n = 10^6$ atoms/s, the resulting shot noise–limited phase sensitivity is $\sim \frac{1}{\sqrt{n}} = 10^{-3}$ rad/$\sqrt{\text{Hz}}$. The phase shift in an atom interferometer due to a constant acceleration g is $\Delta\phi = k_{\text{eff}} g T^2$ [5]. Taking advantage of the $L \approx 8.8$ m vacuum system allows for a long interrogation time of up to $T = 1.34$ s. Finally, using $\hbar k_{\text{eff}} = 10\hbar k$ LMT beamsplitters results in an acceleration sensitivity of $\sim 7 \times 10^{-13}$ $g/\sqrt{\text{Hz}}$ and a precision of $< 10^{-15} g$ after ~ 1 day of integration. In the most conservative case, constraining ourselves to conventional $2\hbar k$ atom optics leads to a precision of $< 10^{-15} g$ after ~ 1 month of integration. This estimate is based on realistic extrapolations from current performance levels, which are at $10^{-10} g$ [16].

2.6.2 Error model

An accurate test of the EP requires a thorough understanding of potential backgrounds. To reach the goal sensitivity, we must control spurious accelerations to $< 10^{-15} g$. Systematic errors at this level can arise from many sources, including gravity gradients, Earth's rotation, and electromagnetic forces. To calculate these contributions to the phase shift, we follow the prescription outlined in Section 2.4. We take the atom's Lagrangian in the lab frame to be

$$L = \frac{1}{2}m(\dot{\mathbf{r}} + \mathbf{\Omega} \times (\mathbf{r} + \mathbf{R}_e))^2 - m\phi(\mathbf{r} + \mathbf{R}_e) - \frac{1}{2}\alpha\mathbf{B}(\mathbf{r})^2 \tag{2.53}$$

where $\mathbf{r}$ is the position of the atom in the lab frame, $\mathbf{R}_e = (0, 0, R_e)$ is the radius of the Earth, $\mathbf{\Omega} = (0, \Omega_y, \Omega_z)$ is the Earth's rotation rate, and $\phi(\mathbf{r})$ is the gravitational potential. In the chosen coordinate system, z is the vertical direction in the lab and $\mathbf{\Omega}$ lies in the y-z plane. We then expand ϕ in a Taylor series about $\mathbf{R}_e$,

$$\phi(\mathbf{r} + \mathbf{R}_e) = -\left(\mathbf{g} \cdot \mathbf{r} + \frac{1}{2!}(T_{ij})r_i r_j + \frac{1}{3!}(Q_{ijk})r_i r_j r_k + \frac{1}{4!}(S_{ijkl})r_i r_j r_k r_l\right) \tag{2.54}$$

where Earth's gravity field is $\mathbf{g} \equiv -\nabla\phi(\mathbf{R}_e)$, the gravity gradient tensor is $T_{ij} \equiv \partial_j g_i$, the second gradient tensor is $Q_{ijk} \equiv \partial_k \partial_j g_i$, the third gradient tensor is $S_{ijkl} \equiv \partial_l \partial_k \partial_j g_i$, and repeated indices are summed over. Since $\hat{\mathbf{z}}$ is in the vertical direction in the lab we have that $\mathbf{g} = (0, 0, -g)$ and $g = 9.8$ m/s^2. The interferometer follows a fountain geometry which is approximately one-dimensional along the z-direction, so we only include Q_{zzz} and S_{zzzz} and safely ignore the other second and third gradient tensor terms. Likewise, in this analysis we assume that off-diagonal gradient tensor terms T_{ij} with $i \neq j$ are small and can be ignored (this is exactly true for a perfectly spherical Earth). The effects of higher-order moments of the gravitational field are treated separately using a perturbative calculation as described in Section 2.6.2.

Because magnetic fields can cause significant systematic errors, the atoms are prepared in one of the magnetic field insensitive clock states ($|m_F = 0\rangle$ states). The residual energy shift in a magnetic field $\mathbf{B}$ is then $U_B = \frac{1}{2}\alpha\mathbf{B}^2$, where α is the second order Zeeman shift coefficient. We consider magnetic fields of the form

$$\mathbf{B}(\mathbf{r}) = \left(B_0 + \frac{\partial B}{\partial z} z\right)\hat{\mathbf{z}} \tag{2.55}$$

where B_0 is a constant bias magnetic field, and $\frac{\partial B}{\partial z}$ is the gradient of the background magnetic field. While this linear model is sufficient for slowly varying fields, in Section 2.6.2 we describe a perturbative calculation that can account for more complicated magnetic field spatial profiles.

We do not include additional electromagnetic forces in the Lagrangian as their accelerations are well below our systematic threshold. For neutral atoms, electric fields are generally not a concern since the atom's response is second order. Furthermore, electric fields are easily screened by the metallic vacuum chamber, leading to negligibly small phase shifts. Short range effects due to the Casimir [32] force or local patch potentials [33] are also negligible since the atoms are kept far (> 1 cm) from all surfaces throughout the experiment.

As explained in Section 2.4.2, we point out that the phase shift derived from Eq. 2.1 is only exactly correct for Lagrangians that are second order in position and velocity. When this is not true, as is the case in Eq. 2.53 when $Q_{ijk} \neq 0$ and $S_{ijkl} \neq 0$, the semiclassical formalism breaks down and there are quantum corrections to the phase shift. However, these corrections depend on the size of the atom wavepacket compared to the length scale of variation of the potential, and are typically negligible for wavepackets ~ 1 mm in size.

To analytically determine the trajectories $\mathbf{r}(t)$, we solve the Euler-Lagrange equations using a power series expansion in t:

$$r_i(t) = \sum_{n=0}^{N} a_{in}(t - t_0)^n \qquad (i = 1, 2, 3) \tag{2.56}$$

The coefficients a_{in} are determined recursively after substitution into the equations of motion. This expansion converges quickly as long as $\Omega T \ll 1$ and $\left|\frac{r^n}{g}\frac{\partial^n g}{\partial r^n}\right| \ll 1$. For our apparatus with characteristic length $r \sim 10$ m and time $T \sim 1$ s these conditions are easily met, since $\Omega T \sim 10^{-4}$ rad and $\left|\frac{r^n}{g}\frac{\partial^n g}{\partial r^n}\right| \sim \left(\frac{r}{R_e}\right)^n \lesssim 10^{-6}$ assuming a spherical Earth. With these trajectories and the interferometer geometry shown in Fig. 2.10 we obtain the following expressions for the phase shift in the slow (state $|1\rangle$) output port:

$$\Delta\phi_{\text{propagation}} = \frac{1}{\hbar}((S_{AC}+S_{CE})-(S_{AB}+S_{BD})) \tag{2.57}$$

$$\Delta\phi_{\text{laser}} = \phi_L(\mathbf{r}_A,0)-\phi_L(\mathbf{r}_C,T)-\phi_L(\mathbf{r}_B,T)+\phi_L(\mathbf{r}_D,2T) \tag{2.58}$$

$$\Delta\phi_{\text{separation}} = \frac{1}{2\hbar}(\mathbf{p}_D+\mathbf{p}_E)\cdot(\mathbf{r}_D-\mathbf{r}_E) \tag{2.59}$$

where S_{ij} is the classical action along the path segment between points $\mathbf{r}_i$ and $\mathbf{r}_j$, and $\mathbf{p}_i = \partial_{\dot{\mathbf{r}}} L(\mathbf{r}_i)$ is the classical canonical momentum at point $\mathbf{r}_i$ after the final beamsplitter. The laser phase shift at each interaction point is

$$\phi_L(\mathbf{r},t) = \mathbf{k}_{\text{eff}}\cdot\mathbf{r} - \omega_{\text{eff}}t + \phi_0 \tag{2.60}$$

where $\mathbf{k}_{\text{eff}}$ and ω_{eff} are the effective propagation vector and frequency, respectively, for whatever atom–laser interaction is used to implement the atom optics. In the case of the stimulated two–photon processes mentioned earlier, $\mathbf{k}_{\text{eff}} = \mathbf{k_2} - \mathbf{k_1}$ and $\omega_{\text{eff}} = \omega_2 - \omega_1 = (k_2 - k_1)/c$.

Using the above method, we computed the phase shift response for a single atom interferometer, and the results are shown in Table 2.1. The values of the experimental parameters used to generate this list are representative of the 8.8 m apparatus described previously. Many of these terms are common to both species, and in order obtain our $< 10^{-15}g$ sensitivity, we rely on their common mode cancellation. In Table 2.2 we compute the differential phase shift between a ^{87}Rb and a ^{85}Rb interferometer. The two species have different masses m and second order Zeeman coefficients α, as well as potentially different launch kinematics $\mathbf{r}(0)$ and $\dot{\mathbf{r}}(0)$. To create Table 2.2, we parameterized the launch kinematics with a differential velocity $\delta\mathbf{v} = (\delta v_x, \delta v_y, \delta v_z)$ and initial position $\delta\mathbf{r} = (\delta x, \delta y, \delta z)$ between the centroids of the two isotope clouds. Residual systematic phase errors are the result of differential accelerations that arise from gravity gradients, second gravity gradients, coriolis and centrifugal forces, and magnetic forces on the atoms.

As justified below, we expect to achieve experimental parameters that reduce the majority of the systematic errors below our experimental threshold. However, the first several terms in Table 2.2 are still too large. In order to further reduce these backgrounds, we can employ propagation reversal to suppress all terms $\propto k_{\text{eff}}^2$. This well–known technique entails reversing the laser propagation vector $\mathbf{k}_{\text{eff}} \longrightarrow -\mathbf{k}_{\text{eff}}$ on subsequent trials and then subtracting the two results [16]. This suppresses terms 1, 4, 9, and 10 by $\Delta k_{\text{eff}}/k_{\text{eff}}$, where Δk_{eff} is the error in k_{eff} made as a result of the reversal. Reducing these terms below our systematic threshold requires $\Delta k_{\text{eff}}/k_{\text{eff}} < 10^{-5}$. The main acceleration signal and all other terms linear in k_{eff} are not suppressed by this subtraction.

After propagation vector reversal, the last important background phase shifts arise from the differential coriolis and centrifugal acceleration between the isotopes (Table 2.2 terms 2, 5, 7, and 8), and from the Earth's gravity gradient (Table 2.2 terms 3 and 6). We discuss the techniques used to control these remaining systematics in Section 2.6.3.

	Phase shift	Size (rad)	Fractional size
1	$-k_{\text{eff}}gT^2$	-2.85×10^{8}	1.00
2	$k_{\text{eff}}R_e\Omega_y^2T^2$	6.18×10^{5}	2.17×10^{-3}
3	$-k_{\text{eff}}T_{zz}v_zT^3$	1.58×10^{3}	5.54×10^{-6}
4	$\frac{7}{12}k_{\text{eff}}gT_{zz}T^4$	-9.21×10^{2}	3.23×10^{-6}
5	$-3k_{\text{eff}}v_z\Omega_y^2T^3$	-5.14	1.80×10^{-8}
6	$2k_{\text{eff}}v_x\Omega_yT^2$	3.35	1.18×10^{-8}
7	$\frac{7}{4}k_{\text{eff}}g\Omega_y^2T^4$	3.00	1.05×10^{-8}
8	$-\frac{7}{12}k_{\text{eff}}R_eT_{zz}\Omega_y^2T^4$	2.00	7.01×10^{-9}
9	$-\frac{\hbar k_{\text{eff}}^2}{2m}T_{zz}T^3$	7.05×10^{-1}	2.48×10^{-9}
10	$\frac{3}{4}k_{\text{eff}}gQ_{zzz}v_zT^5$	9.84×10^{-3}	3.46×10^{-11}
11	$-\frac{7}{12}k_{\text{eff}}Q_{zzz}v_z^2T^4$	-7.66×10^{-3}	2.69×10^{-11}
12	$-\frac{7}{4}k_{\text{eff}}R_e\Omega_y^4T^4$	-6.50×10^{-3}	2.28×10^{-11}
13	$-\frac{7}{4}k_{\text{eff}}R_e\Omega_y^2\Omega_z^2T^4$	-3.81×10^{-3}	1.34×10^{-11}
14	$-\frac{31}{120}k_{\text{eff}}g^2Q_{zzz}T^6$	-3.39×10^{-3}	1.19×10^{-11}
15	$-\frac{3\hbar k_{\text{eff}}^2}{2m}\Omega_y^2T^3$	-2.30×10^{-3}	8.06×10^{-12}
16	$\frac{1}{4}k_{\text{eff}}T_{zz}^2v_zT^5$	2.19×10^{-3}	7.68×10^{-12}
17	$-\frac{31}{360}k_{\text{eff}}gT_{zz}^2T^6$	-7.53×10^{-4}	2.65×10^{-12}
18	$3k_{\text{eff}}v_y\Omega_y\Omega_zT^3$	2.98×10^{-4}	1.05×10^{-12}
19	$-k_{\text{eff}}\Omega_y\Omega_zy_0T^2$	-7.41×10^{-5}	2.60×10^{-13}
20	$-\frac{3}{4}k_{\text{eff}}R_eQ_{zzz}v_z\Omega_y^2T^5$	-2.14×10^{-5}	7.50×10^{-14}
21	$\frac{31}{60}k_{\text{eff}}gR_eQ_{zzz}\Omega_y^2T^6$	1.47×10^{-5}	5.17×10^{-14}
22	$\frac{3}{2}k_{\text{eff}}T_{zz}v_z\Omega_y^2T^5$	-1.42×10^{-5}	5.00×10^{-14}
23	$-\frac{7}{6}k_{\text{eff}}T_{zz}v_x\Omega_yT^4$	1.08×10^{-5}	3.81×10^{-14}
24	$-2k_{\text{eff}}T_{xx}\Omega_yx_0T^3$	-6.92×10^{-6}	2.43×10^{-14}
25	$-\frac{7\hbar k_{\text{eff}}^2}{12m}Q_{zzz}v_zT^4$	-6.84×10^{-6}	2.40×10^{-14}
26	$-\frac{7}{6}k_{\text{eff}}T_{xx}v_x\Omega_yT^4$	-5.42×10^{-6}	1.90×10^{-14}
27	$-\frac{31}{60}k_{\text{eff}}gT_{zz}\Omega_y^2T^6$	4.90×10^{-6}	1.72×10^{-14}
28	$k_{\text{eff}}T_{xx}v_z\Omega_y^2T^5$	4.75×10^{-6}	1.67×10^{-14}
29	$\frac{3\hbar k_{\text{eff}}^2}{8m}gQ_{zzz}T^5$	4.40×10^{-6}	1.55×10^{-14}
30	$\frac{31}{360}k_{\text{eff}}R_eT_{zz}^2\Omega_y^2T^6$	1.63×10^{-6}	5.74×10^{-15}
31	$-\frac{31}{90}k_{\text{eff}}gT_{xx}\Omega_y^2T^6$	-1.63×10^{-6}	5.74×10^{-15}
32	$\frac{\hbar k_{\text{eff}}^2}{8m}T_{zz}^2T^5$	9.78×10^{-7}	3.43×10^{-15}
33	$-\frac{\hbar k_{\text{eff}}\alpha B_0(\partial_zB)T^2}{m}$	-7.67×10^{-8}	2.69×10^{-16}
34	$\frac{31}{60}k_{\text{eff}}gS_{zzzz}v_z^2T^6$	-7.52×10^{-8}	2.64×10^{-16}
35	$-\frac{1}{4}k_{\text{eff}}S_{zzzz}v_z^3T^5$	3.64×10^{-8}	1.28×10^{-16}
36	$\frac{31}{72}k_{\text{eff}}T_{zz}Q_{zzz}v_z^2T^6$	-3.13×10^{-8}	1.10×10^{-16}

Table 2.1: Phase shift response for a single atom interferometer $\frac{\pi}{2}-\pi-\frac{\pi}{2}$ sequence given the Lagrangian in Eq. 2.53. Column 3 shows the fractional size of each term compared to the acceleration signal $k_{\text{eff}}gT^2$. All terms with fractional phase shift $>10^{-16}$ are included. The numbers are for a ^{87}Rb interferometer with the following parameters: $k_{\text{eff}}=2k=2\cdot\frac{2\pi}{780\text{ nm}}$, $T_{zz}=-2g/R_e$, $T_{xx}=T_{yy}=g/R_e$, $Q_{zzz}=6g/R_e^2$, $S_{zzzz}=-24g/R_e^3$, $R_e=6.72\times10^6$m, $B_0=100$ nT, and $\partial_zB=0.1$ nT/m. The Earth's rotation rate is given by $\Omega_y=\Omega\cos\theta_{\text{Lat}}$ and $\Omega_z=\Omega\sin\theta_{\text{Lat}}$ with $\Omega=7.27\times10^{-5}$ rad/s and $\theta_{\text{Lat}}=37.4$ degrees North latitude. The initial position of the atom in the lab is taken as $\mathbf{r}(0)=(x_0,y_0,0)$, with $x_0=1$ mm and $y_0=1$ mm. The initial velocity is $\dot{\mathbf{r}}(0)=(v_x,v_y,v_z)$, with $v_x=1$ mm/s, $v_y=1$ mm/s, and $v_z=13.2$ m/s.

	Phase shift	Size (rad)	Fractional size
1	$-\frac{1}{2}\left(\frac{1}{\mathrm{m85}}-\frac{1}{m_{87}}\right)\hbar k_{\mathrm{eff}}^2 T_{zz} T^3$	1.66×10^{-2}	5.83×10^{-11}
2	$2k_{\mathrm{eff}}\delta v_x \Omega_y T^2$	3.35×10^{-3}	1.18×10^{-11}
3	$-k_{\mathrm{eff}} T_{zz} \delta v_z T^3$	1.44×10^{-4}	5.05×10^{-12}
4	$-\frac{3}{2}\left(\frac{1}{m_{85}}-\frac{1}{m_{87}}\right)\hbar k_{\mathrm{eff}}^2 \Omega_y^2 T^3$	-5.40×10^{-5}	1.90×10^{-13}
5	$-3k_{\mathrm{eff}}\Omega_y^2 \delta v_z T^3$	-4.68×10^{-6}	1.64×10^{-14}
6	$-k_{\mathrm{eff}} T_{zz} \delta z T^2$	8.93×10^{-7}	3.14×10^{-15}
7	$-k_{\mathrm{eff}}\delta y \Omega_y \Omega_z T^2$	-7.41×10^{-7}	2.60×10^{-15}
8	$3k_{\mathrm{eff}}\delta v_y \Omega_y \Omega_z T^3$	2.98×10^{-7}	1.05×10^{-15}
9	$-\frac{7}{12}\left(\frac{1}{m_{85}}-\frac{1}{m_{87}}\right)\hbar k_{\mathrm{eff}}^2 Q_{zzz} v_z T^4$	-1.61×10^{-7}	5.65×10^{-16}
10	$\frac{3}{8}\left(\frac{1}{m_{85}}-\frac{1}{m_{87}}\right)\hbar k_{\mathrm{eff}}^2 g Q_{zzz} T^5$	1.03×10^{-7}	3.63×10^{-16}
11	$-\left(\frac{\alpha_{85}}{m_{85}}-\frac{\alpha_{87}}{m_{87}}\right)\hbar k_{\mathrm{eff}} B_0 (\partial_z B) T^2$	-9.94×10^{-8}	3.49×10^{-16}
12	$-2k_{\mathrm{eff}} T_{xx} \delta x \Omega_y T^3$	-6.92×10^{-8}	2.43×10^{-16}

Table 2.2: Differential phase shift between ^{87}Rb and ^{85}Rb. To create the differential phase shift list we parameterized the launch kinematics with a differential velocity ($\delta v_x = 1\ \mu$m/s, $\delta v_y = 1\ \mu$m/s, $\delta v_z = 12\ \mu$m/s) and position ($\delta x = 1\ \mu$m, $\delta y = 1\ \mu$m, $\delta z = 10$ nm) between the centroids of the two isotope clouds. All other parameters are the same as in Table 2.1. Column 3 shows the fractional size of each term compared to the acceleration signal $k_{\mathrm{eff}} g T^2$. We include all terms with a fractional phase shift $> 10^{-16}$.

Gravity Inhomogeneities

The Taylor series expansion of the gravitational potential (see Eq. 2.54) is a good approximation of the coarse structure of Earth's gravity on length scales of $\sim R_e$, the radius of the Earth. However, local gravity can also vary on much shorter length scales in a way that depends on the specific mass distribution surrounding the experiment, and these gravity inhomogeneities can result in spurious phase shifts. Since these inhomogeneities can be rapidly spatially varying, the Taylor series expansion is not well-suited for their description. Instead, we leverage the fact that these inhomogeneities are typically small in magnitude and solve for the induced phase shift using first–order perturbation theory [4]. This linearization allows us to make a Fourier decomposition of the phase shift response in terms of the spatial wavelengths of the local g–field.

First, we assume a one-dimensional gravitational potential perturbation of the form $\delta\phi(z)$. The gravity field perturbation along the vertical (z) direction is defined as $\delta g_z(z) \equiv -\partial_z \delta\phi$ and may be written as

$$\delta g_z(z) = \int \widetilde{\delta g}_z(\lambda) e^{\frac{i2\pi z}{\lambda}} d\lambda \tag{2.61}$$

where $\widetilde{\delta g}_z(\lambda)$ is the Fourier component of a gravity perturbation with wavelength λ. The total phase shift due to gravity inhomogeneities summed over all wavelengths is

$$\Delta\phi_g = \int T_{gz}(\lambda)\widetilde{\delta g}_z(\lambda) d\lambda \tag{2.62}$$

where $T_{gz}(\lambda)$ is the interferometer's gravity perturbation response function. Qualitatively, the response to short wavelengths is suppressed since the interferometer averages over variations that are smaller than its length [34]. The response is flat for wavelengths longer than the scale of the interferometer, and in the limit where $\lambda \sim R_e$ this analysis smoothly approaches the results of our Taylor series calculation described above.

For the ^{87}Rb–^{85}Rb EP measurement, we are interested in the differential phase response between the isotopes. Figure 2.11 shows the differential response function $\Delta T_{gz}(\lambda) \equiv |(T_{gz})_{87} - (T_{gz})_{85}|$ for gravity inhomogeneities. Once again, short wavelength variations are suppressed since each interferometer spatially averages over a ~ 10 m region. The peak response occurs at a length scale set by the spatial separation of the arms of a single interferometer $\Delta z = \frac{\hbar k_{\mathrm{eff}}}{m} T \sim 16$ mm. Perfect differential cancellation between isotopes is not achieved because the spatial separation of the arms is mass dependent. Additionally, the long wavelength differential response is suppressed because the differences between the isotope trajectories are negligible when compared to variations with length scales much longer than Δz.

The differential response curve allows us to compute systematic errors arising from the specific gravity environment of our interferometer. Quantitative estimates of these effects requires knowledge of the local $\delta g_z(z)$, which may be obtained through a combination of modelling and characterization.

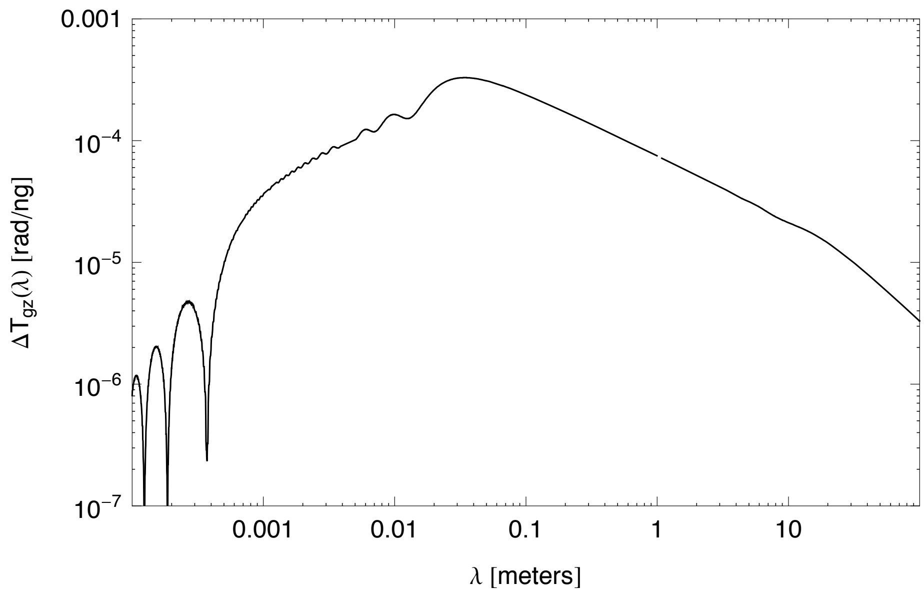

Figure 2.11: Differential gravity response function versus spatial wavelength λ between simultaneous ^{87}Rb and ^{85}Rb interferometers. Short wavelengths are averaged over by each individual interferometer, while long wavelength inhomogeneities cancel as a common-mode between the two species. This response curve assumes identical launch kinematics for the two isotopes.

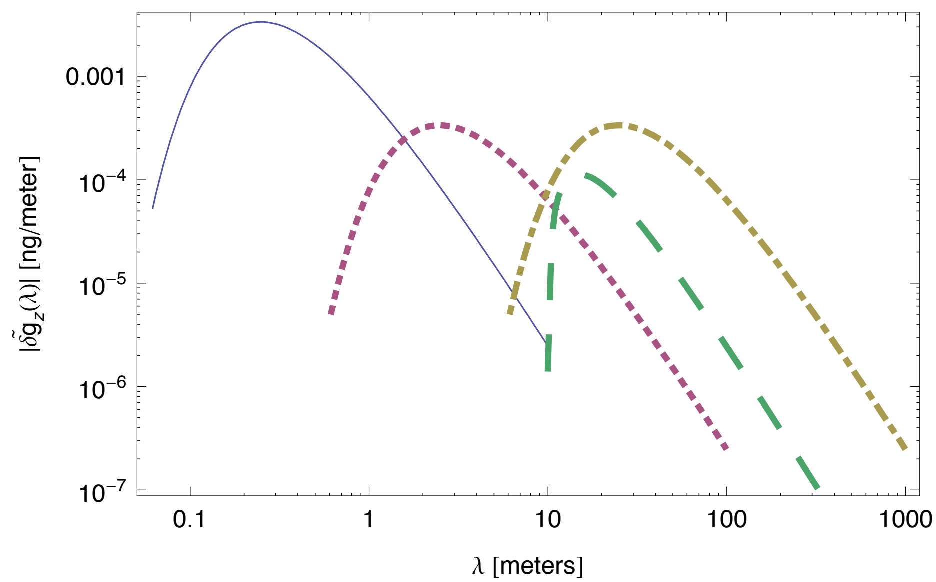

Figure 2.12: The magnitude power spectra of the local gravitational field, $\widetilde{\delta g}_z$, for several example sources. The solid (blue) curve is a 10^{-2} kg point source, 10 cm from the center of the interferometer. Similarly, the dotted (purple) curve is a 1 kg source at 1 m and the dash-dotted (yellow) curve is 1000 kg at 10 m. The long-dashed (green) curve is a thin 10 m long rod of mass 10 kg, parallel to the interferometer, whose center is 1 m from the interferometer.

The atom interferometer itself can be used as a precision gravimeter for mapping $\delta g_z(z)$ in situ. By varying the launch velocity, initial vertical position, and interrogation time T, the position of each gravity measurement can be controlled.

Figure 2.11 shows that the differential ^{87}Rb–^{85}Rb interferometer is maximally sensitive to short wavelength ($\lambda \sim 1-10$ cm) gravitational inhomogeneities. To investigate the impact of local uneven mass distributions on the experiment, we compute the spectrum $\widetilde{\delta g}_z(\lambda)$ of various sources at different distances from the interferometer. These results are shown in Fig. 2.12. When combined with our response function (Fig. 2.11), we see that for typical mass inhomogeneities, only those that are within a few centimeters of the interferometer can cause potentially significant systematic phase shifts. These nearby inhomogeneities result in phase errors of $\sim 10^{-6}$ rad, which is slightly above our target sensitivity. It will therefore be especially important for the EP measurement that we characterize the local g–field at the centimeter scale.

Magnetic field inhomogeneities

The linear expansion of $\mathbf{B}$ in Eq. 2.55 approximates large scale variation of the magnetic field. However, local field inhomogeneities may exist on short length scales due to the presence of nearby

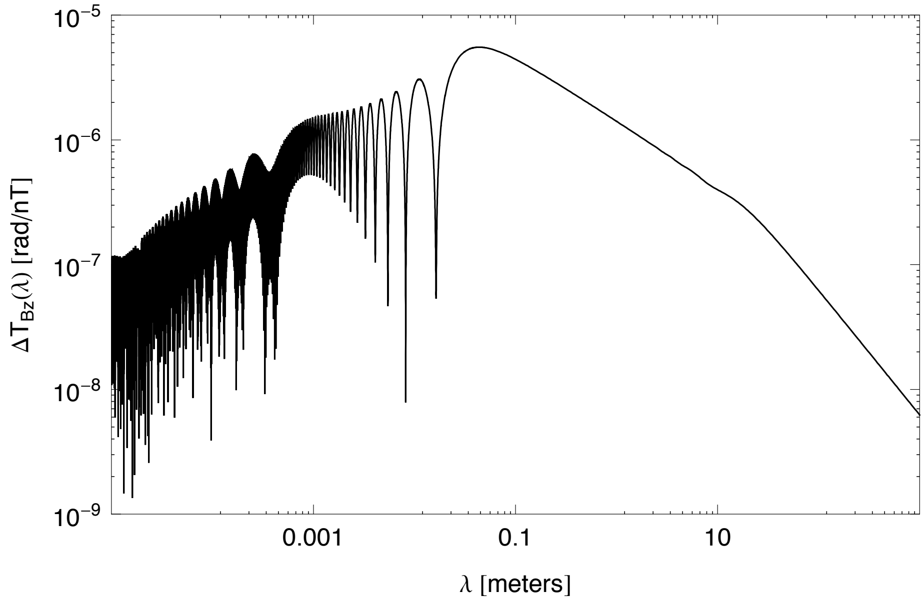

Figure 2.13: Differential magnetic field response function between simultaneous ^{87}Rb and ^{85}Rb interferometers. Short wavelengths are averaged over by each individual interferometer, while long wavelength inhomogeneities cancel as a common-mode between the two species. This response curve assumes identical launch kinematics for the two isotopes.

magnetic materials. These variations are not well approximated by a Taylor series expansion. Using the same procedure described above for gravity inhomogeneities, we write the local magnetic field as

$$B(z) = \int \widetilde{B}_z(\lambda) e^{\frac{i2\pi z}{\lambda}} d\lambda \tag{2.63}$$

where $\widetilde{B}_z(\lambda)$ is the Fourier component of a field perturbation with wavelength λ. The total phase shift from magnetic field inhomogeneities is

$$\Delta\phi_B = \int T_{Bz}(\lambda) \widetilde{B}_z(\lambda) d\lambda \tag{2.64}$$

Here $T_{Bz}(\lambda)$ is the interferometer's magnetic inhomogeneity response function.

As with gravity above, we compute the differential response function $\Delta T_{Bz}(\lambda) \equiv |(T_{Bz})_{87} - (T_{Bz})_{85}|$ between ^{87}Rb and ^{85}Rb (see Fig. 2.13). The differential response arises because the isotopes have different second order Zeeman coefficients α, as well as different masses. This sensitivity curve drives our magnetic shield design requirements, as discussed in Section 2.6.3.

2.6.3 Controlling potential systematic errors

Rotation of the Earth

The largest systematic term in the phase shift expansion for a dual species differential interferometer after propagation reversal is due to the rotation of the Earth. Specifically, a differential acceleration due to the coriolis force occurs if the isotopes have different transverse velocities δv_x (Table 2.2, term 2). Reducing this phase shift below our systematic threshold would require $\delta v_x < 10^{-11}$ m/s, which is challenging. However, this specification can be relaxed by artificially making the rotation rate zero.

To a good approximation[3], the atoms are only affected by the Earth's rotation through their coupling to the laser, so the coriolis acceleration can be eliminated by rotating the laser in the opposite direction of the Earth's rotation. In order to calculate the effect of this rotation compensation, we performed the phase shift calculation using a rotating $\mathbf{k}_{\text{eff}}$. Following the work of [15], we use a retro-reflection configuration to deliver the laser beams $\mathbf{k}_1$ and $\mathbf{k}_2$ to the atoms. We rotate $\mathbf{k}_{\text{eff}}$ by actuating the retro-reflection mirror. As a result, the incoming beam remains pointing along the z-direction and only the reflected beam rotates. With this configuration, $\mathbf{k}_{\text{eff}}$ is given by

$$\mathbf{k}_{\text{eff}} = -2k\hat{\mathbf{n}}(\hat{\mathbf{n}} \cdot \hat{\mathbf{k}}) \tag{2.65}$$

where $\hat{\mathbf{n}}$ is the time-dependent unit normal vector of the retro-reflection mirror, and $\hat{\mathbf{k}}$ is a unit vector in the direction of the fixed delivery beam. Notice that the direction of $\mathbf{k}_{\text{eff}}$ rotates as desired, but its length now depends on angle[4].

The resulting phase shift list appears in Table 2.3, with $\delta\Omega_y$ and $\delta\Omega_z$ the errors in the applied counter-rotation rate. Assuming a transverse velocity difference of $\delta v_x \sim 1$ μm/s, these rotation compensation errors must be kept below $10^{-5}\Omega_{\text{Earth}} \approx 1$ nrad/s. Methods for measuring angles with nanoradian precision have already been demonstrated [35]. In order to actuate the mirror at this level of precision we can use commercially available sub-nm accurate piezo-electric actuators along with active feedback.

Notice that not all rotation-related phase errors are removed by rotation compensation. Terms that arise from the differential centrifugal acceleration between the isotopes (e.g., Table 2.3 terms 3 and 4) are not suppressed. Physically, this is a consequence of the fact that the retro-reflection mirror that we use to change the laser's angle is displaced from the center of rotation of the Earth by R_e. Therefore, although we can compensate for the angle of the laser by counter-rotating, the retro-reflection mirror remains attached to the rotating Earth, leading to a centrifugal acceleration

[3]The atoms are also weakly coupled electromagnetically and gravitationally to the local environment, which is fixed to the rotating Earth. These cross-couplings to rotation are generally not important because the dominant gravitational interaction with the Earth is spherically symmetric, and all electromagnetic interactions with the atom (e.g. with the applied magnetic bias field) are naturally small.

[4]This small change in magnitude of k_{eff} does not lead to any problematic phase errors in the interferometer since the total angle through which the laser rotates is only $\Delta\theta = 2\Omega_{\text{Earth}}T \sim 10^{-4}$ rad, and the effect is $\mathcal{O}(\Delta\theta^2)$.

	Phase shift	Size (rad)	Fractional size
1	$-\frac{1}{2}\left(\frac{1}{m_{85}}-\frac{1}{m_{87}}\right)\hbar k_{\text{eff}}^2 T_{zz} T^3$	1.66×10^{-2}	5.83×10^{-11}
2	$-k_{\text{eff}} T_{zz} \delta v_z T^3$	1.44×10^{-4}	5.05×10^{-12}
3	$-\frac{3}{2}\left(\frac{1}{m_{85}}-\frac{1}{m_{87}}\right)\hbar k_{\text{eff}}^2 \Omega_y^2 T^3$	-5.40×10^{-5}	1.90×10^{-13}
4	$-3k_{\text{eff}} \Omega_y^2 \delta v_z T^3$	-4.68×10^{-6}	1.64×10^{-14}
5	$-k_{\text{eff}} T_{zz} \delta z T^2$	8.93×10^{-7}	3.14×10^{-15}
6	$-\frac{7}{12}\left(\frac{1}{m_{85}}-\frac{1}{m_{87}}\right)\hbar k_{\text{eff}}^2 Q_{zzz} v_z T^4$	-1.61×10^{-7}	5.65×10^{-16}
7	$\frac{3}{8}\left(\frac{1}{m_{85}}-\frac{1}{m_{87}}\right)\hbar k_{\text{eff}}^2 Q_{zzz} g T^5$	1.03×10^{-7}	3.63×10^{-16}
8	$-\left(\frac{\alpha_{85}}{m_{85}}-\frac{\alpha_{87}}{m_{87}}\right)\hbar k_{\text{eff}} B_0 (\partial_z B) T^2$	-9.94×10^{-8}	3.49×10^{-16}
9	$k_{\text{eff}} T_{xx} \delta x \Omega_y T^3$	3.46×10^{-8}	1.22×10^{-16}
10	$-2k_{\text{eff}} \delta v_x \delta \Omega_y T^2$	-3.35×10^{-8}	1.18×10^{-16}

Table 2.3: Differential phase shift list with rotation compensation. Terms 1, 3, 6, and 7 will be suppressed by the propagation reversal technique described in Section 2.6.2.

of the phase fronts. After propagation reversal, the only term of this type that is significant is $\sim k_{\text{eff}} \Omega_y^2 \delta v_z T^3$ (Table 2.3 term 4). However, this term is smaller than and has the same scaling with experimental control parameters as the gravity gradient phase shift (Table 2.3 term 2), so the constraints described in Section 2.6.3 to suppress the gravity gradient terms are sufficient to control this centrifugal term as well.

One potential obstacle in achieving the required transverse velocity constraint of $\delta v_x \sim 1$ μm/s is the expected micro-motion the atoms experience in the TOP magnetic trap prior to launch [36]. Micro-motion orbital velocities in a tight TOP trap such as ours can approach ~ 1 cm/s in the transverse plane. Although the differential orbital velocities are suppressed by the ^{87}Rb–^{85}Rb mass ratio, the resulting $\delta v_x \sim 100$ μm/s is still too large. This problem can potentially be solved by adiabatically reducing the magnetic field gradient and increasing the rotating field frequency prior to launch.

Gravity gradients

The largest systematic background after rotation compensation is due to the gravity gradient along the vertical direction of the apparatus ($T_{zz} = \partial_z g_z$). Since gravity is not uniform, the two isotopes experience a different average acceleration if their trajectories are not identical. This effect causes a differential phase shift proportional to the initial spatial separation and initial velocity difference between the isotopes (see Table 2.3, terms 2 and 5). Assuming a spherical Earth model, the gravity gradient felt by the atoms is $T_{zz} \sim 3 \times 10^{-16} g/\text{nm}$, which means that the initial vertical position difference between the isotopes δz must be < 1 nm and the initial vertical velocity difference δv_z

must be < 1 nm/s in order to reduce the systematic phase shift beneath our threshold.

The experiment is designed to initially co-locate the two isotope clouds at the nm level by evaporative cooling both species in the same magnetic trap. For trapping, we use the state $|F = 2, m_F = 2\rangle$ for ^{87}Rb and $|F = 3, m_F = 3\rangle$ for ^{85}Rb since these states have the same magnetic moment [37]. The mass difference between the isotopes leads to a differential trap offset in the combined magnetic and gravitational potential given by

$$\Delta z_{\text{trap}} = \frac{g\Delta m}{\mu_B B''} \tag{2.66}$$

where Δm is the ^{87}Rb–^{85}Rb mass difference, B'' is the magnetic field curvature of the trap, and μ_B is the Bohr magneton. Our TOP magnetic trap is designed to provide a field curvature $B'' \sim 4 \times 10^5$ Gauss/cm^2 which reduces the trap offset to $\Delta z_{\text{trap}} \approx 10$ nm. The resulting systematic error is $\sim 10^{-14} g$, but it can be subtracted during the analysis given a knowledge of Δz_{trap} at the $\sim 10\%$ level. This offset can be inferred from a measurement of the field curvature B'' of the trap (e.g., by measuring the trap oscillation frequency). The gravity gradient must also be known, but this can be characterized in situ by using the interferometer as a gradiometer [16].

Control of the gravity gradient phase shift (Table 2.3 term 2) requires that the differential launch velocity be $\delta v_z \leq 1$ nm/s. Therefore we cannot employ standard launch techniques (e.g., moving molasses) since the velocity uncertainty is fundamentally limited by the photon recoil velocity $v_R \sim 6$ mm/s due to spontaneous emission. Instead, the atoms are launched using an accelerated optical lattice potential [27]. We launch the two species using the same far-detuned ($\sim$ 200GHz) optical lattice, coherently transferring $\sim 2200\hbar k$ of momentum to each cloud. Because the two species have different masses, they have different Bloch oscillation times $\tau_B = \frac{\hbar k_{\text{eff}}}{ma}$, where a is the lattice acceleration. As a result, after adiabatically ramping down the lattice potential, the two species are in different momentum eigenstates since they have absorbed a different number of photons. The differential velocity after launch is then

$$\delta v_L = \hbar k_{\text{eff}} \left(\frac{N_{85}}{m_{85}} - \frac{N_{87}}{m_{87}} \right) \tag{2.67}$$

where N_{85} and N_{87} are the number of photons transferred to ^{85}Rb and ^{87}Rb, respectively. We choose the integers $N_{85} = 2168$ and $N_{87} = 2219$ such that their ratio is as close to the isotope mass ratio as possible, resulting in $\delta v_L \sim 12$ μm/s. After launch, we can perform a velocity selective transition to pick out a common class from the overlapping distributions of the two isotopes, which at the expense of atom number could conceivably allow us to achieve our differential velocity constraint.

There are several additional ways to reduce the gravity gradient systematics beyond precise control of the launch kinematics. We can implement a 4-pulse sequence ($\frac{\pi}{2} - \pi - \pi - \frac{\pi}{2}$) which suppresses all phase shift terms $\propto T^3$ at the cost of an order one loss in acceleration sensitivity [6]. This eliminates the velocity dependent gravity gradient phase shifts but would still require that we maintain tight control over the initial differential vertical position between the isotope clouds.

Secondly, we can potentially reduce the local gravity gradient T_{zz} by applying appropriate trim masses around the apparatus. It has been shown [34] that in principle a local mass distribution can effectively cancel the gravity gradient of the Earth for a 10 m-scale apparatus. Reducing T_{zz} by an order of magnitude would relax our initial position constraint to the level provided by the expected value of Δz_{trap}, thereby removing the requirement for subtraction during data analysis.

Magnetic fields

The magnetic field phase shift appearing in Table 2.3 (term 8) constrains the maximum allowed linear field gradient to $\partial_z B < 0.1$ nT/meter. In the interferometer region, the measured gradient of the Earth's field is $\sim 3\ \mu$T/m, and therefore we require a shielding ratio of at least $\sim 5 \times 10^4$. In addition to suppressing the field gradient, the magnetic response function (see Fig. 2.13) indicates that the field must be uniform on length scales ~ 1 cm. Large magnetic shields with similar performance have been demonstrated [38]. The magnetic shielding for our interferometer region is provided by a three–layer concentric cylindrical shield made of high permeability material. To maintain a pristine magnetic environment, we use an aluminum vacuum chamber and non-magnetic materials inside the shielded region.

In order to verify the performance of the magnetic shield, we must characterize the field. As with gravity inhomogeneities, the atom interferometer can be used to map the local magnetic field in situ, in this case by using a magnetic field sensitive ($m_F \neq 0$) state [39].

2.7 Conclusion

In these notes we have given an overview of the light-pulse method, and discussed applications to inertial navigation and a test of the Equivalence Principle. As the field continues to progress, we see two trends evolving. First, a steady evolution in the technology associated with laser manipulation of atoms will lead to progressive development of compact, field-ready inertial sensors. Such developments include integrated photonics packages for the laser and optics paths, and integrated optics bench and vacuum systems. On the other hand, continued evolution of high performance science instruments will extend the physics reach of this technology. For example, next generation light-pulse atom interference systems appear capable of placing superb limits on atom charge neutrality [40], making terrestrial tests of General Relativity [34, 41], and possibly detecting low frequency gravitational waves [42, 43].

Chapter 3

AGIS-LEO

Long baseline atom interferometry is potentially the most sensitive technique to measure the motion of inertial test masses. Below is a proposal for a space-based, atom interferometric gravitational wave detector. The work presented below was done in concert with the AGIS-LEO collaboration, which consists of Jason M. Hogan, David M. S. Johnson, Susannah Dickerson, Tim Kovachy, Alex Sugarbaker, Sheng-wey Chiow, Peter W. Graham, Mark A. Kasevich, Babak Saif, Surjeet Rajendran, Philippe Bouyer, Bernard D. Seery, Lee Feinberg, and Ritva Keski-Kuha. The content below is tentatively accepted to be published in the journal General Relativity and Gravitation. A version of it can be found online at http://arxiv.org/abs/1009.2702. I was directly involved in producing this work.

3.1 Introduction and Science Goals

Gravitational waves (GWs) provide an unexplored window into the Universe. The vast majority of current astrophysical observations rely on electromagnetic waves. Unlike electromagnetic waves, which originate from moving charges, gravitational waves are sourced by energy (*i.e.*, the stress energy tensor) and can thus reveal entirely new types of sources that are difficult or impossible to observe electromagnetically, including black hole, white dwarf, and neutron star binaries. Further, gravitational waves are not screened even by dense concentrations of matter and can thus be used to probe environments, including the very early universe, that are inaccessible to any other telescope. Unlike electromagnetic waves, gravitational waves permit observations beyond the surface of last scattering. Such observations could change our understanding of the fundamental laws of physics, shedding light on the birth of our universe in reheating after inflation or revealing evidence that our universe underwent a phase transition at extremely high temperatures. Historically, the ability to observe new sections of the electromagnetic spectrum has led to many new discoveries. Similarly, the gravitational spectrum seems likely to be filled with sources that will provide unique information about astrophysics, cosmology, and high energy physics [44, 45, 46]. However, it may be the unpredicted sources of gravitational radiation that have the greatest effect on our understanding of the universe.

We propose a low Earth orbit gravitational wave observatory based on atom interferometry, AGIS-LEO, which will be sensitive to signals in the frequency band 50 mHz - 10 Hz, complementary to the reach of the LIGO, VIRGO and LISA instruments. This proposal extends the AGIS measurement concept presented in recent work [43, 47]. The frequency band probed by AGIS-LEO is home to a variety of interesting astrophysical phenomena. Textbook sources of gravitational radiation are compact object binaries, which give a coherent, oscillatory gravitational wave signal. These binaries consist of systems where both components are compact objects such as white dwarfs, neutron stars, or black holes [45]. They strongly radiate gravitational waves because they contain large-mass stars orbiting in close proximity, and the frequency of the emitted waves traverses through the AGIS-LEO frequency band as the binaries merge.

The remaining lifetime of a binary source before coalescence is strongly correlated with the frequency of the gravitational wave that it emits. Lifetimes in the AGIS-LEO band range from several thousand years for a source at 10 mHz to about a week for a source at 1 Hz. These observation times are significantly longer than the lifetime of a comparable source in the frequency band of modern ground-based laser interferometer detectors (e.g., LIGO or VIRGO). This long lifetime enhances the number of such sources in this frequency band relative to that of LIGO or VIRGO and increases their detectability owing to the increased observation time. Additionally, due to their large physical size, the mergers of white dwarfs end before they enter the frequency band accessible to LIGO and VIRGO but lie well within the frequency band of AGIS-LEO. Numerous such binaries are expected to lie in the Milky Way, and AGIS-LEO may be able to detect the brightest of these sources, such as a compact solar mass binary within ~ 100 kiloparsec.

Other major sources of gravitational radiation in the AGIS-LEO band include the inspirals of solar mass black holes into intermediate mass ($10^2 M_\odot \lessapprox M \lessapprox 10^4 M_\odot$) black holes out to distances of ~ 1 megaparsec and into massive ($10^4 M_\odot \lessapprox M \lessapprox 10^7 M_\odot$) black holes out to distances of ~ 10 megaparsecs. There are significant uncertainties in the expected merger rate of such objects, particularly because their detection is a challenge through electromagnetic methods. In fact, the detection of the gravitational waves emitted by these objects is one of the cleanest ways to observe them, and hence there is great scientific importance in achieving this goal.

The detection and parameter estimation of these mergers could play an important role in enhancing our understanding of the growth of structure in our Universe. For example, they can help uncover the nature of seed black holes that lead to the formation of massive black holes at the centers of galaxies. These binary mergers will also test general relativity in the strong field regime - a regime which is currently inaccessible to terrestrial and solar system tests of gravity. Furthermore, the detection of a large number of such binary mergers may permit their use as standard gravitational wave sirens (analogous to an electromagnetic standard candle), thereby providing a new, precise cosmological probe [48].

There are also several possible cosmological sources of gravitational radiation, described below, which are more easily observable in the low frequency band below 10 Hz. Although these sources are by no means certain to exist, the discovery of any one would be incredibly important for cosmology and high energy physics.

We believe our universe began in a period of inflation; the subsequent reheating produced all the matter we see today. This period of reheating after inflation produces gravitational waves which are potentially strong enough to be detectable [49, 50, 51]. The gravitational wave spectrum produced by reheating is expected to be peaked; the frequency of the peak is model-dependent and varies with the scale of reheating. After redshifting, this peak can lie anywhere between roughly 10 mHz and 10^9 Hz, so it may be possible for AGIS-LEO to observe gravitational waves from reheating.

A first-order phase transition in the early universe can produce gravitational waves through bubble nucleation and turbulence [52, 53]. A phase transition at the well-motivated TeV scale is likely to produce gravitational waves with a frequency today in a range near 100 mHz. In some models with new physics at the weak scale [54], including some supersymmetric [55] and warped extra-dimensional [56] models, the electroweak phase transition can produce gravitational waves well above the threshold for detection by atom interferometers.

A network of cosmic strings produces a stochastic background of gravitational waves from cusps, kinks, and vibrations of the strings. If such a network exists, calculations [57, 58] indicate that the gravitational waves produced may well be strong enough to be observable in the low frequency band from $\sim$ 100 mHz to 10 Hz for cosmic string tensions as low as $G\mu \sim 10^{-8}$ (or lower for a more advanced detector that could be realized in the future).

There are many other possible sources of gravitational waves from fundamental physics in the early universe, including Goldstone modes of scalar fields [59], or radion modes and fluctuations of our brane in an extra dimensional scenario [60, 61]. Pre-big bang [62] or extended [63] inflation could also lead to a strong gravitational wave signal in the relevant frequency range. Other astrophysical sources may lead to an interesting stochastic gravitational wave background (for a review see [46]).

3.2 Mission Overview

The promising scientific and technical applications of atomic physics techniques in space have motivated a number of ambitious proposals in recent years [64, 65, 66], and significant work has already been done toward developing space-qualified cold atom-based instruments [67, 68]. Here we propose to search for gravitational waves using light-pulse atom interferometry [2, 69]. (An alternative approach using confined lattice-hold interferometry will also be explored in Sec. 3.6.1.) In this section, we describe atom interferometry and how it can be used to make a sensitive space-based gravitational wave detector, AGIS-LEO. We will also introduce the multiple-satellite configurations that arise naturally from the application of atom interferometry to the detection of gravitational waves.

3.2.1 Atom Interferometry

In a light-pulse atom interferometer (AI), an atom is forced to follow a superposition of two spatially separated free-fall paths [1]. This is accomplished by coherently splitting the atom wavefunction with a pulse of light that transfers momentum $\hbar\mathbf{k}_{\text{eff}}$ to a part of the atom [2], where $\mathbf{k}_{\text{eff}}$ is the effective wavevector of the light. As a result of the momentum recoil from the interaction with the light, the two halves of the atom wavefunction spatially separate along the light propagation direction, each following its own classical path due to the velocity difference. Note that the interferometer paths considered here are all nominally one-dimensional along the direction of $\mathbf{k}_{\text{eff}}$; unwanted displacements perpendicular to $\mathbf{k}_{\text{eff}}$ arise only because of Coriolis forces present in low Earth orbit and will be discussed below (see Figs. 3.2 and 3.3 for examples). When the atom wavefunction is later recombined, the resulting interference pattern depends on the relative phase accumulated along the two paths. This phase shift results from both the free-fall evolution of the quantum state along each path and from the local phase of the laser, which is imprinted on the atom during each of the light pulses [70]. Since it precisely compares the motion of the atom to the reference frame defined by the laser phase fronts, the phase shift is sensitive to inertial forces present during the interferometer sequence. To be concrete, the lasers are ultimately anchored to the frame of the AI sensor, so the AI phase shift is a record of the motion of the inertial atom with respect to the sensor frame as measured by the laser phase fronts, which act as a sub-micron periodicity ruler. Equivalently, the atom interferometer phase shift can be viewed as a clock comparison between the time kept by the laser's phase evolution and the atom's own internal clock. Sensitivity to gravitational waves may be understood as arising from this time comparison, since the presence of space-time strain changes the light travel time between the atom and the laser [42].

Gravitational waves can be detected by recording the phase shift they induce in the atom interferometer. Since this phase shift is also directly sensitive to unavoidable phase noise and vibration noise present on the laser light pulses that act as the atom beamsplitters and mirrors, a practical gravitational wave detector must be differential. The concept of the AGIS-LEO instrument is to compare the phase shifts of two separate atom interferometers that are manipulated by the same laser, and that are therefore subject to the same laser noise. The differential phase shift is still sensitive to the gravitational wave while the laser noise is suppressed as a common mode. [42]

A number of proposals for gravitational wave detection using atom interferometry have been discussed based on a variety of implementations [71, 72, 73, 74, 75, 76, 77, 78, 79], including the light pulse technique. The principle difference between these schemes and the implementation discussed

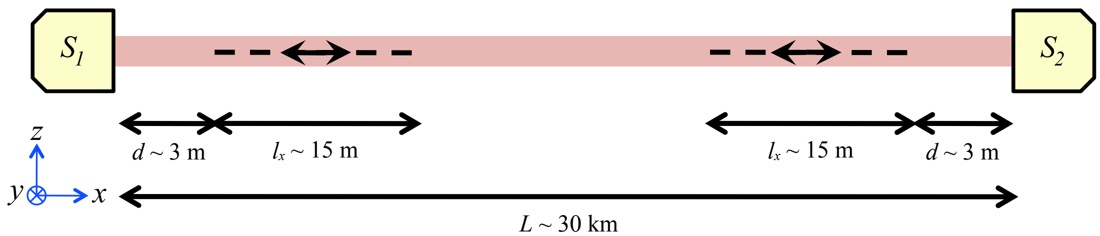

Figure 3.1: Schematic of one beam line of the proposed setup. Two satellites S_1 and S_2 house the lasers and atom sources. The dashed lines represent the path of length l_x traveled by the atoms during the interferometer sequence. At their closest approach, the atoms are a distance d from the satellites. The red band represents the paths of the lasers along the axis between the satellites.

here is that the AGIS-LEO configuration is based on a differential measurement comparing the phases of two widely separated atom interferometers. Using the AGIS protocol, the atoms effectively function as test masses that record the effects of the gravitational wave on the propagating laser field as it travels between the two interferometers [42]. This offers the significant practical advantage that the instrument sensitivity can easily be increased by using longer baseline separations between the interferometers without having to increase the size of the interferometers themselves.

Figure 3.1 illustrates the AGIS-LEO differential interferometer concept. Two satellites S_1 and S_2 separated by baseline L provide a pair of counter-propagating laser beams required to implement the necessary two-photon transitions commonly used for atom optics in alkali atoms [1, 80, 81]. The atom interferometer regions (dashed lines) are each near one of the laser sources and have size l_x along the laser direction (the x-direction). As mentioned above, the atom wavepackets are separated along the laser propagation direction so that the paths of the interferometer arms are nominally one dimensional along x. Both satellites are in orbit around the Earth, and for simplicity we take the orbital angular velocity vector to be $\overrightarrow{\Omega} = \Omega_{\mathrm{or}}\hat{\mathbf{y}}$ (for a leader-follower orbit, as seen in Fig. 3.6).

The specific sequence of light pulses used to implement the atom interferometer is an important consideration since it affects both the gravitational wave response and the sensitivity to noise. The standard three-pulse accelerometer pulse sequence ($\pi/2 - \pi - \pi/2$) consists of an initial beamsplitter pulse, a mirror pulse a time T later, and a final interfering beamsplitter pulse after an additional time T. Such an interferometer is sensitive to gravitational waves as well as uniform acceleration; this interferometer geometry is discussed in detail in [43]. However, in a leader-follower low Earth orbit (see Fig. 3.6) the three-pulse sequence cannot be used by the AGIS-LEO instrument because the large rotation bias prevents the overlap of the interferometer arms. In order to observe interference and thereby measure a phase shift, the arms of the atom interferometer must overlap at the end of the pulse sequence within the size of the atomic wavepacket (i.e., within the spatial coherence length of the atom). For ^{87}Rb atoms, the coherence length of the atom's external degrees of freedom is $\Delta x \simeq \hbar/\Delta p \sim 10\ \mu\mathrm{m}\sqrt{\frac{100\ \mathrm{pK}}{\tau}}$ for atom ensemble temperature τ. The large Coriolis forces experienced

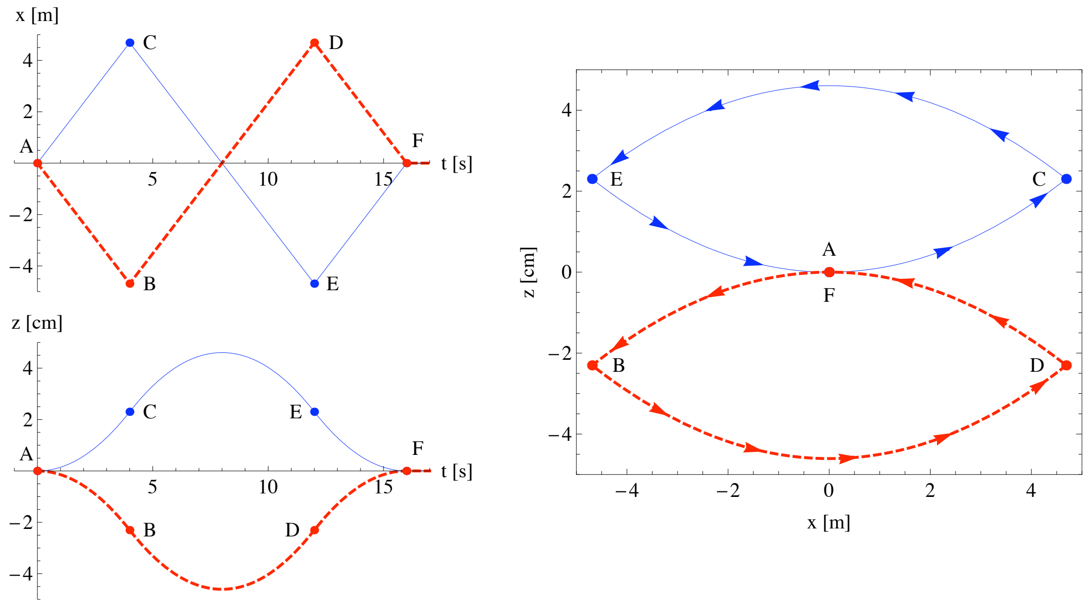

Figure 3.2: Four-pulse atom interferometer sequence using double-diffraction LMT beamsplitters under the influence of a leader-follower low Earth orbit rotation bias. The upper arm of the interferometer is the solid (blue) path and the lower arm the dashed (red) path. Interactions with the light pulses are shown as points labeled by letters A through F. The interferometer laser pulses are directed along the x-direction with $\mathbf{k}_i = \kappa_i k_{\text{eff}} \hat{\mathbf{x}}$ for $1 \leq i \leq 4$. The normalized wavevectors have values $\kappa_1 = 1$, $\kappa_2 = 2$, $\kappa_3 = 2$ and $\kappa_4 = 1$. The pulses occur at times $t_1 = 0$, $t_2 = T$, $t_3 = 3T$ and $t_4 = 4T$ with interrogation time $T = 4$ s. See Fig. 3.1 for the definition of coordinate axes.

in low Earth orbit can cause deviations to the nominal paths of the atom interferometer's arms that exceed this coherence length, causing the interferometer not to close. To address this problem, AGIS-LEO uses multiple-pulse sequences with appropriate symmetry to facilitate interferometer closure.

An example of a pulse sequence that closes even in the presence of large Coriolis forces is the symmetric four-pulse sequence ($\pi/2{-}\pi{-}\pi{-}\pi/2$) shown in Fig. 3.2, where the four laser pulses propagate along the x direction with effective wavevectors $\mathbf{k}_i = k_i \hat{\mathbf{x}}$ for $1 \leq i \leq 4$. This interferometer uses two mirror pulses, resulting in a two-loop interferometer geometry along the x direction. Although the interferometer is nominally one dimensional along x, the Coriolis force causes transverse deflections so that at point C the upper arm z-position is $z_C = \frac{1}{2}\left|2\overrightarrow{\Omega} \times \mathbf{v}\right| T^2 = \Omega_{\text{or}}(\hbar k_{\text{eff}}/m)T^2 \sim 2$ cm, where once again Ω_{or} is the orbital angular velocity of the AGIS satellites. For increased symmetry between the arms, the interferometer in Fig. 3.2 uses double-diffraction beamsplitters [82] in which the upper and lower arms of the interferometer are given a positive and negative recoil kick, respectively. In contrast to the typical single-diffraction beamsplitter where only one arm changes velocity, the double-diffraction beamsplitter results in symmetric Coriolis deflections which are needed to close the interferometer.

Multiple-pulse sequences such as the one in Fig. 3.2 also help reduce systematic errors in the instrument. For example, by adjusting the time spacings between the pulses, the four-pulse sequence can be tuned to be insensitive to either uniform accelerations or gravity gradients and Coriolis effects [6, 43]. This is useful for suppressing spurious noise from these sources. The sequence shown in Fig. 3.2 uses a symmetric pulse spacing with pulses occurring at times $t_1 = 0$, $t_2 = T$, $t_3 = 3T$, and $t_4 = 4T$. As a result, this interferometer is insensitive to acceleration noise, but it remains sensitive to gravity gradients.

A five-pulse sequence ($\pi/2 - \pi - \pi - \pi - \pi/2$) offers a number of advantages that make it a promising candidate for the AGIS-LEO mission. Figure 3.3 shows an example of a symmetric five-pulse sequence with pulses at times $t_1 = 0$, $t_2 = T$, $t_3 = 3T$, $t_4 = 5T$, and $t_5 = 6T$ and with variable pulse wavevectors k_i along the x-direction assuming large momentum transfer (LMT) atom optics with magnitudes $k_1 = k_{\text{eff}}$, $k_2 = \frac{9}{4}k_{\text{eff}}$, $k_3 = \frac{5}{2}k_{\text{eff}}$, $k_4 = \frac{9}{4}k_{\text{eff}}$, and $k_5 = 2k_{\text{eff}}$, where $\hbar k_{\text{eff}} = 200\hbar k$ (see Sec. 3.3.2). The three mirror pulses in this sequence result in a three-loop interferometer geometry along the x-direction, and these particular t_i and k_i values are chosen to simultaneously meet the constraints of interferometer closure in x and z and symmetry in time (see the detailed discussion in Sec. 3.5.1). For four- and five-pulse sequences with the same low-frequency detection limit, the z-deflection of the five-pulse interferometer is intrinsically smaller than that of the four-pulse interferometer since it goes through zero at the midpoint. Smaller deflection is desirable since the atoms must stay within the beam waist of the interferometer lasers (see Sec. 3.4). In addition, the five-pulse sequence is insensitive to both acceleration noise and gravity gradient induced background noise which reduces the temperature requirements (kinematic constraints) of the atom source (see Sec. 3.3.1).

3.2.2 The Gravitational Wave Phase Shift

The gravitational wave-induced phase difference between two five-pulse ($\pi/2-\pi-\pi-\pi-\pi/2$) atom interferometers separated by baseline distance L along the direction of propagation of the laser light is

$$\Delta\phi_{\text{GW}} = 8k_{\text{eff}}hL\sin^4(\omega T/2)\left(\frac{7+8\cos\omega T}{2}\right)\sin\theta_{\text{GW}} \tag{3.1}$$

where h is the gravitational wave strain for a gravitational wave of frequency ω, T is the interrogation time of the interferometer, and k_{eff} is the effective wavevector of the AI beamsplitter. Here $\theta_{\text{GW}} = \omega t$ is the phase of the gravitational wave at the time t of the measurement. This result follows from a fully relativistic phase shift calculation discussed in [41, 34] applied to the five-pulse sequence described in Sec. 3.2.1.

To maximize the strain sensitivity of the detector, the baseline L should be made as large as

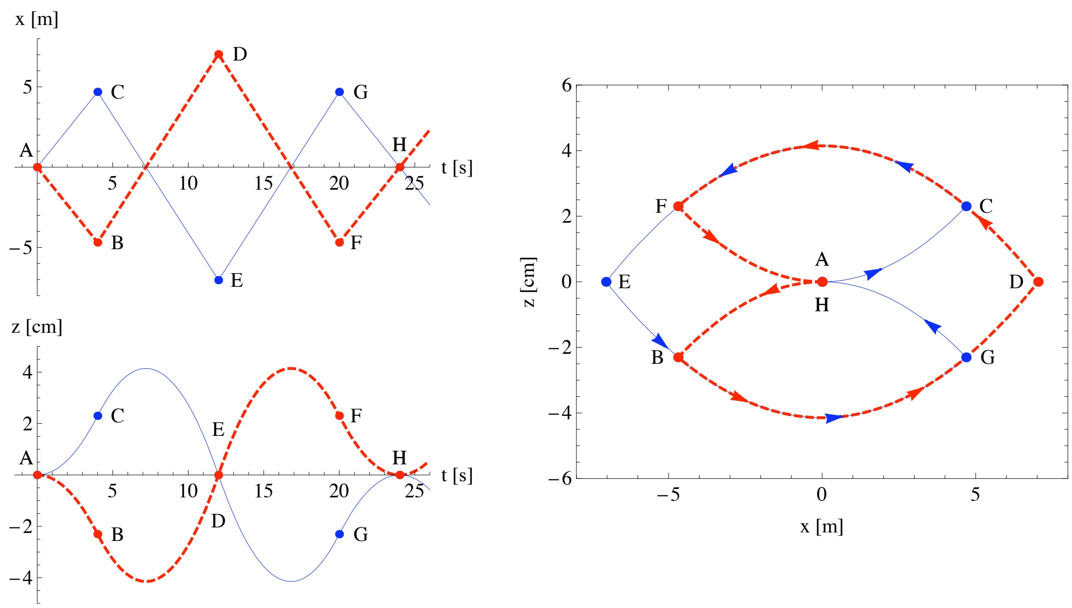

Figure 3.3: Five-pulse atom interferometer sequence using double-diffraction LMT beamsplitters under the influence of a leader-follower low Earth orbit rotation bias. The upper arm of the interferometer is the solid (blue) path and the lower arm is the dashed (red) path. Interactions with the light pulses are shown as points labeled by letters A through H. The interferometer laser pulses are directed along the x-direction with $\mathbf{k}_i = \kappa_i k_{\text{eff}} \hat{\mathbf{x}}$ for $1 \leq i \leq 5$. The normalized wavevectors have values $\kappa_1 = 1$, $\kappa_2 = \frac{9}{4}$, $\kappa_3 = \frac{5}{2}$, $\kappa_4 = \frac{9}{4}$ and $\kappa_5 = 2$. The pulses occur at times $t_1 = 0$, $t_2 = T$, $t_3 = 3T$, $t_4 = 5T$ and $t_5 = 6T$ with interrogation time $T = 4$ s. See Fig. 3.1 for the definition of coordinate axes.

possible. This is a straightforward scaling to implement since only the laser light needs to travel the distance L between the atom interferometers; each AI can remain relatively small, situated at the ends of the long baseline (see Fig. 3.1). From Eq. (3.1), the interferometer is maximally sensitive to GW frequencies at and above the corner frequency $\omega_c = \frac{2}{T}\cos^{-1}\sqrt{\frac{3}{8}} \approx 2\pi \cdot 0.29/T$ at which point $\Delta\phi_{\mathrm{GW}} = (125/16)k_{\mathrm{eff}}hL\sin\theta_{\mathrm{GW}}$. For lower frequencies $\omega < \omega_c$, the 3 dB point occurs at $\omega_{3\mathrm{dB}} = \frac{2}{T}\csc^{-1}[4/\sqrt{5}] \approx 2\pi \cdot 0.19/T$ and sensitivity falls off as ω^4. At higher frequencies $\omega > \omega_c$, the envelope of the sensitivity curve is constant and the periodic anti-resonant frequencies that are present in (Eq. 1) can easily be avoided by varying T [43].

In order to observe a gravitational wave, the phase shift $\Delta\phi_{\mathrm{GW}}$ must be sampled repeatedly as it oscillates in time due to the evolution of θ_{GW}. To avoid aliasing, the sampling rate f_r must be at least twice the frequency of the gravitational wave. The sampling rate can be increased (without decreasing T) by operating multiple concurrent interferometers using the same hardware [43]. In these multiplexed interferometers, atom clouds with different velocities would be addressed via Doppler shifts of the laser light.

3.2.3 Gravitational Wave Sensitivity

The phase sensitivity of an atom interferometer is limited by quantum projection noise $\delta\phi = 1/\mathrm{SNR} = 1/\sqrt{N_a}$, where N_a is the number of detected atoms per second. For shot-noise limited atom detection, the resulting strain sensitivity to gravitational waves becomes $h = \frac{\delta\phi}{8k_{\mathrm{eff}}L}\csc^4(\omega T/2)\left(\frac{2}{7+8\cos\omega T}\right)$. Figure 3.4 shows the AGIS-LEO mission strain sensitivity (enveloped) for a single pair of satellites using the five-pulse interferometer sequence from Fig. 3.3 with $\hbar k_{\mathrm{eff}} = 200\hbar k$ beamsplitters, an $L = 30$ km baseline, a $T = 4$ s interrogation time, and a phase sensitivity of $\delta\phi = 10^{-4}$ rad/$\sqrt{\mathrm{Hz}}$. Although the sensitivity spectrum contains periodic nulls corresponding to frequencies at which $\Delta\phi_{\mathrm{GW}} = 0$, in practice the entire strain sensitivity envelope shown in Fig. 3.4 can be accessed by scanning T over a small range on repeated interferometer sequences.

Note that the AGIS-LEO experimental parameters that define the sensitivity curve in Fig. 3.4 exceed the current state-of-the-art. Specifically, the LMT beamsplitter momenta (see Sec. 3.3.2) and cold atom source flux (see Sec. 3.3.1) both represent specifications that are likely to be feasible in the near future. In fact, many of the parameters that set the instrument sensitivity can eventually be improved further as AI technology continues to evolve. Figure 3.5 shows the gravitational wave sensitivity that would be possible with a more advanced version of the AGIS instrument. This future sensitivity curve assumes the same five-pulse interferometer sequence and $\hbar k_{\mathrm{eff}} = 200\hbar k$ atom optics, but with an $L = 10^4$ km baseline, a $T = 40$ s interrogation time, and an improved phase sensitivity of $\delta\phi = 10^{-5}$ rad/$\sqrt{\mathrm{Hz}}$. We note that this sensitivity curve does not include limits that might be imposed by the systematic effects relevant to such a future apparatus, since its purpose is simply to illustrate the potential for atomic gravitational wave detectors to improve as technology advances. The large baseline of this future instrument would require a solar orbit and is not addressed here.

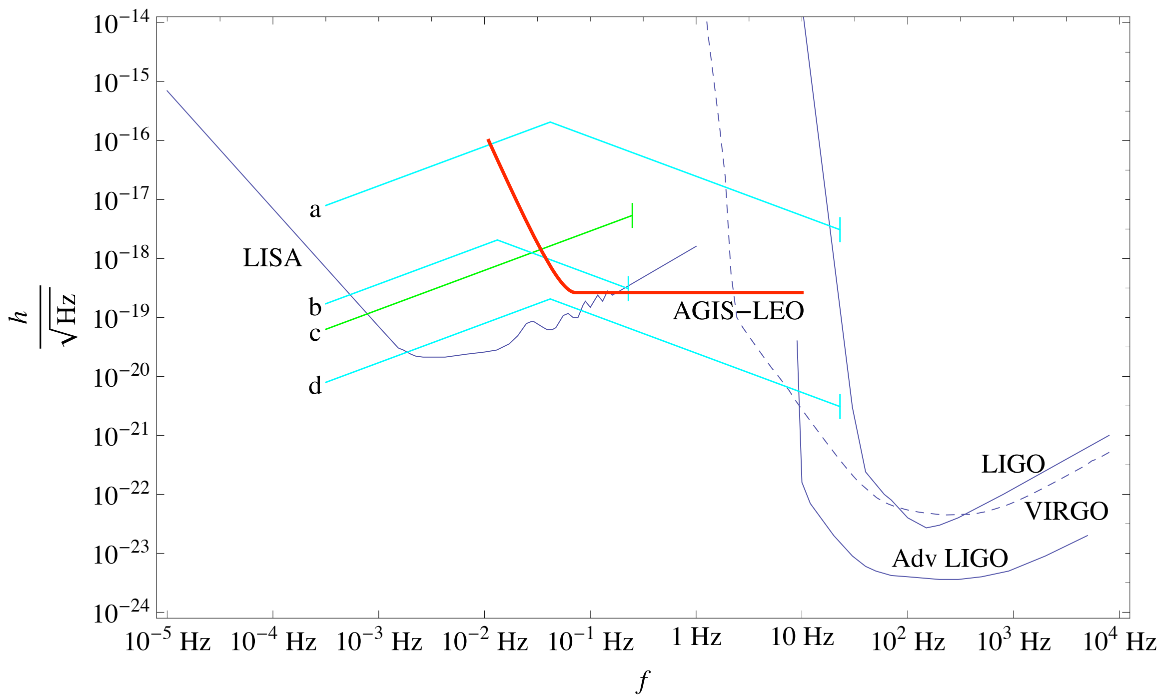

Figure 3.4: The thick (red) curve shows the enveloped strain sensitivity versus gravitational wave frequency of a five-pulse interferometer sequence for a single arm of the proposed AGIS-LEO instrument. Analogous plots for LIGO [83], Advanced LIGO [84], VIRGO [84] and LISA [85] are included for comparison. Curves a-d show gravitational wave source strengths after integrating over the lifetime of the source or one year, whichever is shorter: (a) represents inspirals of $10^3 M_\odot$, $1 M_\odot$ intermediate mass black hole binaries at 10 kpc, (b) inspirals of $10^5 M_\odot$, $1 M_\odot$ massive black hole binaries at 10 Mpc, (c) white dwarf binaries at 10 kpc, and (d) inspirals of $10^3 M_\odot$, $1 M_\odot$ intermediate mass black hole binaries at 10 Mpc.

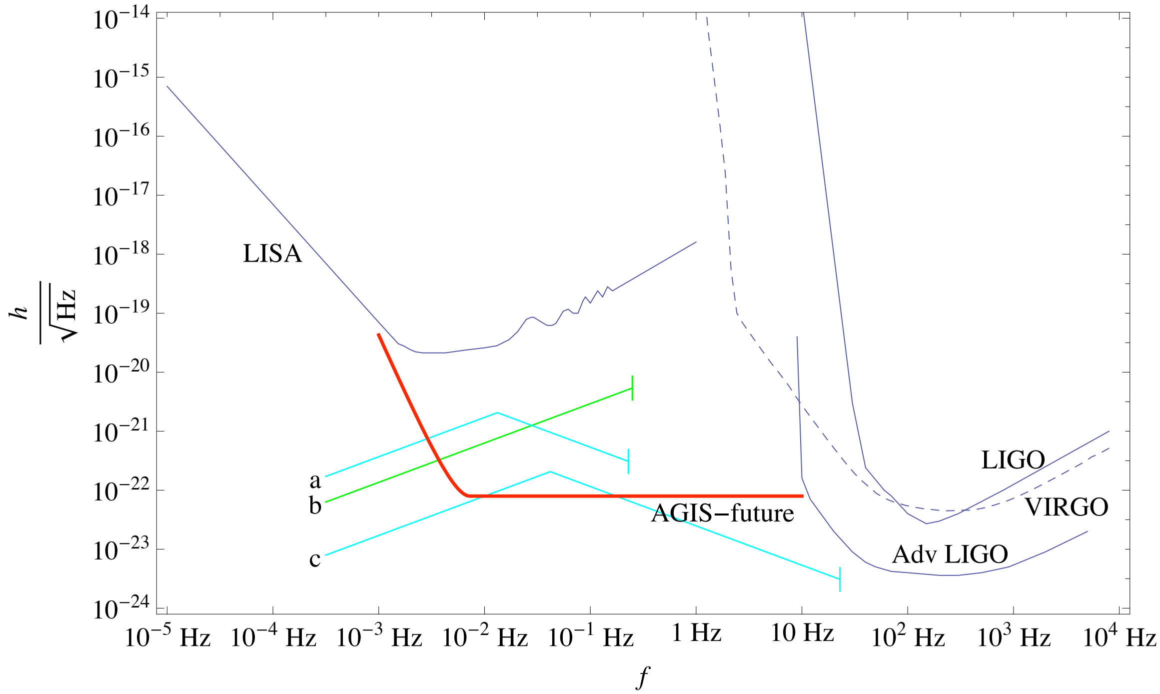

Figure 3.5: The thick (red) curve shows the enveloped strain sensitivity versus gravitational wave frequency of a five-pulse interferometer sequence for a single arm of a sensor that could be realized in the future. Analogous plots for LIGO [83], Advanced LIGO [84], VIRGO [84] and LISA[85] are included for comparison. Curves a-c show gravitational wave source strengths after integrating over the lifetime of the source or one year, whichever is shorter: (a) represents inspirals of $10^5 M_\odot$, $1 M_\odot$ massive black hole binaries at 10 Gpc, (b) white dwarf binaries at 10 Mpc, and (c) inspirals of $10^3 M_\odot$, $1 M_\odot$ intermediate mass black hole binaries at 10 Gpc. It should be noted that the stochastic background of white dwarf binaries will likely overwhelm all but the strongest of these sources at frequencies less than ~ 0.5 Hz (see Fig. 3.9). Note also that this sensitivity curve assumes atom shot noise limited detection and does not include limits that could be imposed by systematic effects present in the future apparatus.

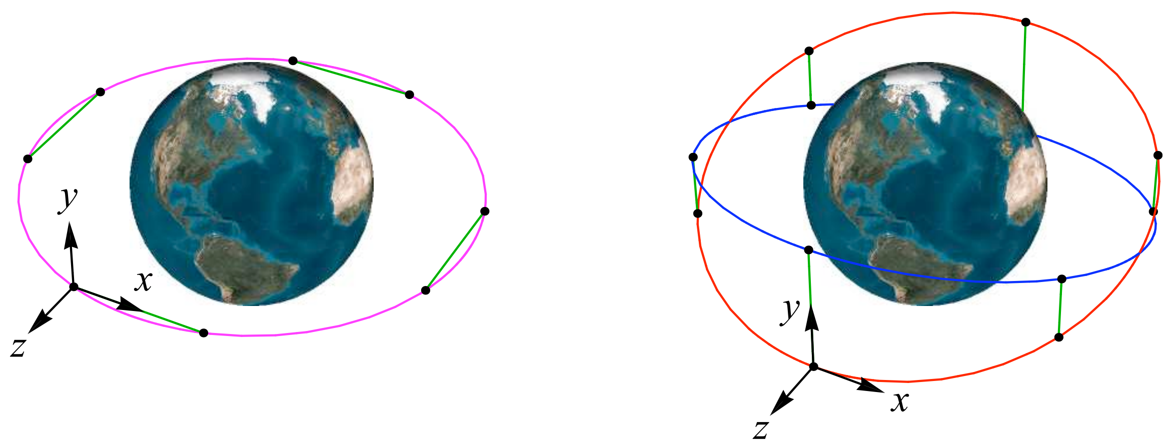

Figure 3.6: Conceptual orbit diagrams for two possible AGIS-LEO two-satellite configurations, not to scale. In each figure, the locations of the satellites (black dots) are shown at quarter-intervals (left) or sixth-intervals (right) of the orbital period, connected by a green line corresponding to the interferometer baseline axis of the gravitational wave detector at that instant. *Left*: The leader-follower (LF) configuration assumed in this paper unless otherwise specified. *Right*: The inclined-great-circles (IGC) configuration that can be used to mitigate rotation effects. Note the difference in definitions of the coordinate axes in the satellite frames between the LF and IGC configurations. In the LF configuration (as in Fig 3.1), the x-direction is fixed to the rotating baseline and the y-direction is normal to the satellites' mutual orbital plane. In the IGC configuration, the baseline lies on the y-axis and the x-axis is chosen to be tangential to the satellite's orbital curve.

3.2.4 Satellite Configuration

Two-Satellite Configurations

As shown in the preceding sections, gravitational waves can be detected with a pair of atom interferometers manipulated by two satellites separated by a baseline of length L. Ideally, the satellites would be inertial and the baseline length would be constant. Instead, many practical issues must be addressed for the non-inertial low Earth orbits we discuss here. The simplest two-satellite configuration is the circular leader-follower (LF) configuration (see Fig. 3.6), in which both satellites follow circular orbits with the same altitude and angle of inclination, but with different angles of true anomaly (i.e. phases). This yields baselines of constant length that also rotate constantly at the rate Ω_{or} as the satellites orbit the Earth. While this rotation bias yields off-axis Coriolis deflections (as discussed in Sec. 3.2.1 and Figs. 3.2 and 3.3), the LF configuration provides a simple starting point for our analysis. Thus, an equatorial, circular, two-satellite, LF orbit will be assumed throughout this paper unless stated otherwise.

However, it should be noted that other orbital configurations exist for two satellites that have some advantages over the simple LF configuration. For instance, the inclined-great-circles (IGC) configuration shown on the right in Fig. 3.6 reduces the impact of rotation bias at the cost of a variable baseline length. An IGC configuration can be achieved by placing two satellites in circular orbits of the same radius, but with different angles of inclination. If the true anomaly difference between the satellites were zero, then the baseline axis would not rotate, but it would flip twice per period and its length would oscillate. Practically, however, it is necessary to add some anomaly to ensure that the satellites do not collide. In this case, one satellite appears to move in an ellipse centered on the other. The eccentricity of this ellipse is determined by the ratio between the anomaly and the inclination, allowing continuous tuning from a configuration that favors nearly constant orientations at the cost of length variations to a configuration with invariant lengths and constant rotation rates. Thus, while any orbit pair at equal altitude requires the satellites to circle each other once per period, the IGC configuration can be used to concentrate most of this rotation at the orbit crossing points. During the rest of the orbit, the baseline deviates little from one of two antiparallel orientations. For instance, at an orbital altitude of 1000 km (105 minute period), an IGC orbit with a maximum separation of 30 km and a minimum separation of 100 m exhibits sub-degree angular deviations during two 45-minute windows centered at the points of maximum separation.

The potential benefits of the IGC configuration are substantial. Most important, the small angular deviations decrease rotation bias and Coriolis deflections (see Fig. 3.7), relaxing many of the constraints based on $\Omega_{\rm or}$ in the LF configuration. Further improvements are then possible with star-tracker-based satellite pointing stabilization and active servoing of the beam delivery optics to compensate for residual rotation [70]. Finally, if the baseline axis points toward the Sun, this configuration makes a sunshield more practical, since it would only have to protect the interferometer region along its transverse extent rather than its full length. (As will be discussed in Sec. 3.5.2, the Earth's shadow could also be used for sun shielding, yielding about 35 minutes of uptime per period in an equatorial orbit with an altitude of 1000 km. This corresponds well with the 45-minute window of small angular deviations discussed previously.)

The variable baseline length in an IGC configuration implies that the atoms from the two satellites have a non-zero relative velocity along the interferometer axis. This requires that the lasers be phase-chirped to account for the Doppler shift caused by this changing relative velocity over the duration of the interferometer. Additionally, the atoms released from the two satellites will have different respective Doppler shifts. Thus, two different sets of frequencies will be needed to address a pair of interferometers in a differential configuration (with one interferometer released from each satellite). The two frequency sets can be realized by electro-optically generating multiple sidebands from a common source laser.

The relative velocity between the two interferometers is advantageous for Doppler-based multiplexing of multiple interferometers during the same time window. The relative velocity between the atoms and the satellite depends on ΔT, the time since the start of each interferometer. Thus, each cloud is naturally tagged with a different frequency which allows for easy individual addressability.

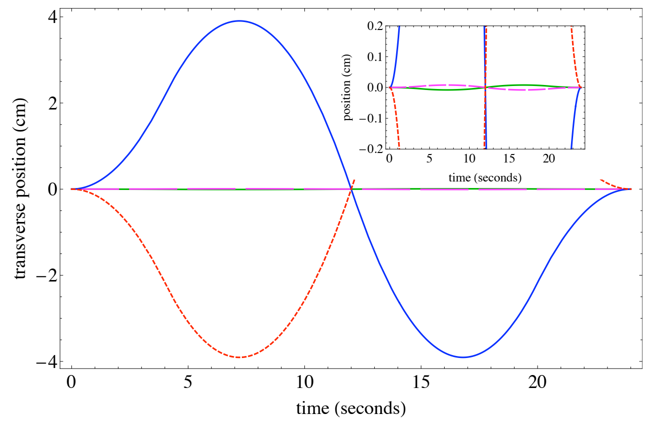

Figure 3.7: Coriolis deflection from the interferometer axis for inclined-great-circles and leader-follower orbits. Leader-follower deflection is shown in solid (blue) and dotted (red) lines with a maximum transverse separation of 8 cm between the two arms of the interferometer, while the inclined-great-circles deflection is shown in solid (green) and dashed (purple) lines with a maximum separation of 160 μm. The inset shows a close-up view in order to detail the inclined-great-circles deflection.

Three-Satellite Configurations

While a single pair of atom interferometers along a common baseline is sufficient for detecting gravitational waves, a multiple-arm configuration with several AI pairs offers many potential advantages, which we will discuss shortly. Our three-arm proposal consists of three satellites maintaining constant separations of $\sim$ 30 km in formation flight (Fig. 3.8). Such a formation can be pictured as a chief point on a circular Earth-orbit, with deputy satellites in circular orbits in a Hill frame around the chief point. It has been shown that such orbits exist with low station-keeping costs even when one accounts for the J_2 oblateness of the Earth [86]. Lasers located inside each satellite send light to both of the other satellites, forming a set of three arms consisting of counter-propagating light beams. Each arm has an atom interferometer at both ends - near, but outside of, the satellites (Fig. 3.1). The pair of atom interferometers along a given arm is manipulated by the counter-propagating light beams of that arm, and the phase difference between these AIs due to a gravitational wave is given by Eq. (3.1). As a result, each arm of the instrument acts as an independent gravitational wave detector.

The desire for multiple, independent GW detectors in the AGIS instrument is motivated by one of our primary science goals: the detection of stochastic GWs. While a coherent GW signal, such as one from a binary merger, can be detected by a single arm of AGIS, stochastic GWs can be detected only by studying correlations in the signals from at least two independent GW detectors. In Fig. 3.9, the sensitivities of the AGIS-LEO and future instruments are plotted against expected stochastic backgrounds. Following the analysis in [43], we display the sensitivities as 95% confidence limits on

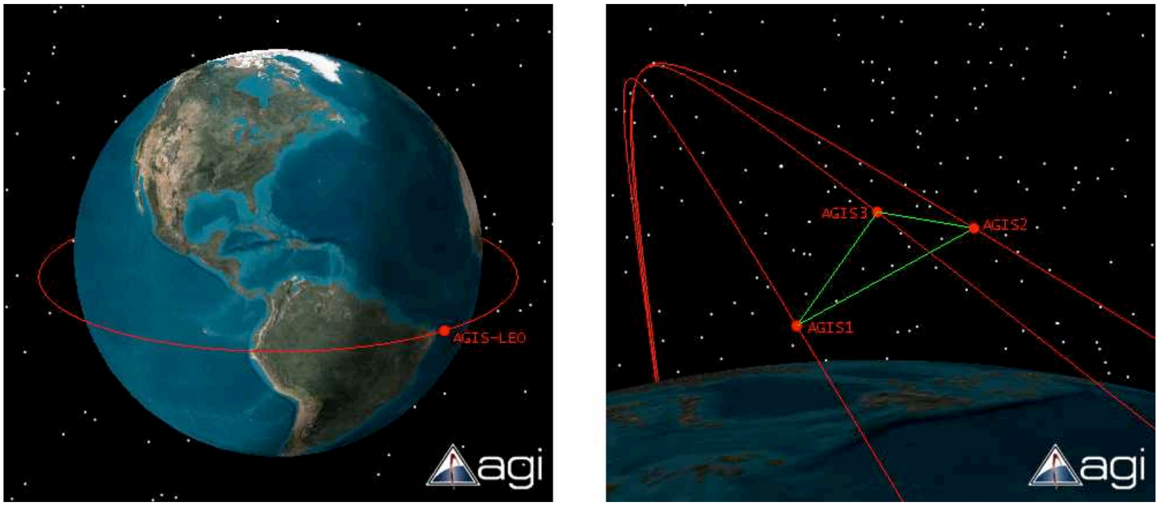

Figure 3.8: AGIS-LEO three-satellite orbit simulations. *Left*: Equatorial orbit of the entire AGIS-LEO constellation at an altitude of 1000 km. *Right*: Closer view of three AGIS satellites in an equilateral triangular formation. The simulations were done using the STK software package from AGI.

the gravitational wave energy density, Ω_{GW}. To account for the effect of the satellite configuration on the instrument's sensitivity to stochastic sources, the calculation includes the appropriate overlap function value $\gamma = \pi$ for an equilateral triangle geometry [87]. As indicated by the sensitivity of the future-technology curve, sensors using AIs have the potential to detect stochastic GWs from cosmic strings and, depending on the exact frequency cutoff of the white dwarf background, possibly even inflation and TeV scale phase transitions.

There are also technical advantages to the AGIS-LEO three satellite configuration over a single arm design. Additional suppression of laser phase noise and vibration noise is possible by combining the signals from multiple arms. Noise that originates from a particular satellite can have a significant common component for the two arms serviced by that satellite. Comparison of the signals from multiple arms can be used to subtract this common noise, allowing for additional noise cancellation in the event that the single arm differential suppression is not sufficient. Finally, the use of three independent arms provides redundancy and improves uptime. A final determination of the AGIS-LEO orbital configuration will require a more thorough analysis of the costs and benefits outlined here.

3.3 Instrument Overview

A single phase measurement in an atom interferometer consists of three steps: atom cloud preparation, interferometer pulse sequence, and detection. In order to attain the 10^{-4} rad/$\sqrt{\text{Hz}}$ phase

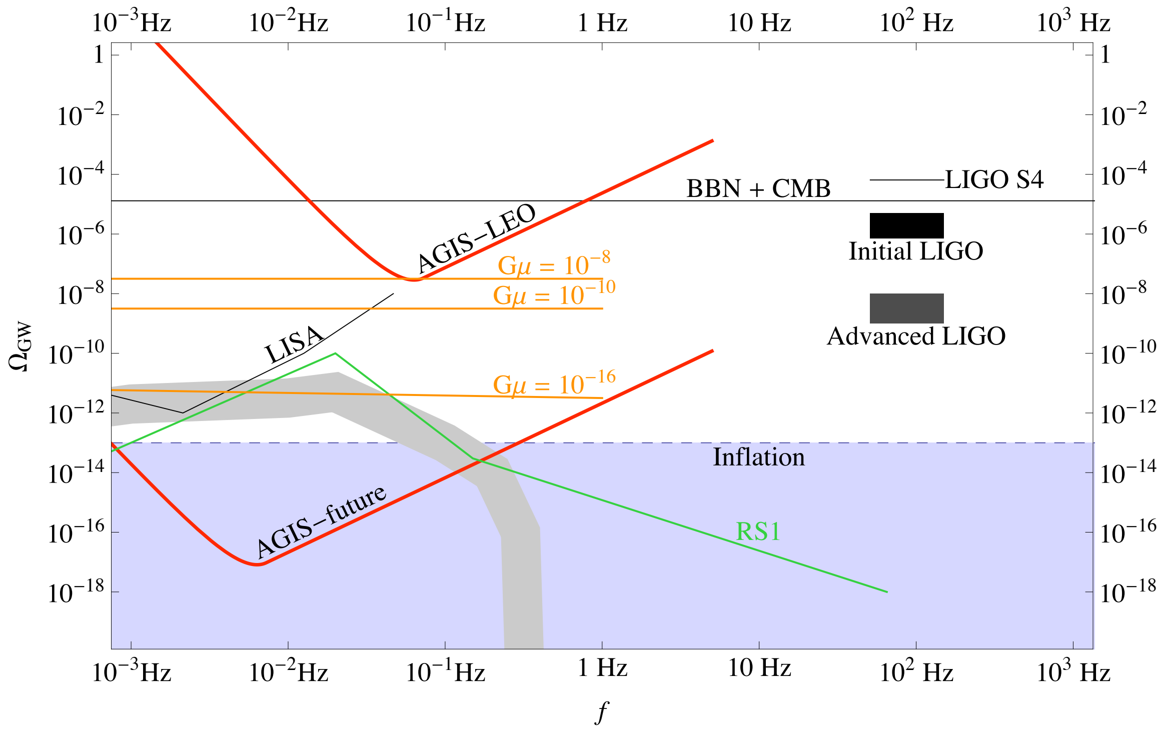

Figure 3.9: The 95% confidence limits on the sensitivity of the instrument in terms of gravitational wave energy density, Ω_{GW}, for present and future parameters are plotted as thick (red) curves. These sensitivities are for a five-pulse interferometer sequence. The limit from the LIGO Science Run 4 and the projected limits from initial and advanced LIGO are shown [88]. The limits from models of big-bang nucleosynthesis (BBN) [89] and the cosmic microwave background (CMB) [90] apply to the integral of the stochastic gravitational wave background over frequency. The possible region of gravitational waves produced by a period of inflation (not including reheating) is shown. The upper limit on this region is set by the COBE bound [91]. The gray band shows a prediction for the stochastic gravitational wave background from extragalactic white dwarf binaries; its width shows an expected error [92]. The (green) curve labeled RS1 corresponds to an example spectrum of gravitational waves from a TeV scale phase transition, in this case from the RS1 spacetime geometry [56]. The (orange) lines labeled with $G\mu$ correspond to one prediction for a network of cosmic strings with tensions $G\mu = 10^{-8}$, $G\mu = 10^{-10}$ and 10^{-16} (with $\alpha = 0.1$ and $\gamma = 50$) [58]. Note that these cosmic string estimates have large uncertainties and may be optimistic assumptions.

Parameter	Specification	Location in Text
Baseline	30 km	Gravitational Wave Sensitivity (Sec. 3.2.3)
Atom Flux	$10^8\ \frac{\text{atoms}}{\text{s}}$	Atom Source (Sec. 3.3.1)
Cloud Temperature	~ 100 pK	Atom Source & Wavefront (Secs. 3.3.1 & 3.4.2)
Cooling Laser Powers	~ 100 mW	Atom Source (Sec. 3.3.1)
Interferometer Sequence	Three loop	AI Pulse Sequence (Sec. 3.2.1)
Beamsplitter Momentum	$\sim 200\hbar k$	Atom Optics (Sec. 3.3.2)

Table 3.1: Instrument parameter summary. For laser requirements, see Table 3.2.

sensitivity, the instrument must be designed to provide, maintain, and measure a cold atom flux of 10^8 atoms/s. Instrument parameters resulting from the discussion below are summarized in Table 3.1.

3.3.1 Cold Atom Source

The suppression of velocity-dependent backgrounds requires the atoms to be cold, with RMS velocity widths as small as ~ 100 μm/s, corresponding to cloud temperatures of ~ 100 pK. Cold atom clouds with 10^8 to 10^{10} atoms are readily produced using modern laser cooling techniques [93]. Furthermore, it has been demonstrated that the hardware necessary for these standard techniques is simple and can be made robust and compact - enough to perform well in drop towers and on airplanes [94, 95, 70]. It is possible to make cold atom sources with fluxes of 10^{10} atoms/s that fit in a volume of less than 20 cm^3 and require ~ 100 mW of laser power [96].

Although high-flux atom sources are already available, a source capable of producing the necessary sub-nanokelvin cloud temperature at these fluxes has not yet been demonstrated. However, there are many promising cooling technologies in this rapidly evolving field. The required ~ 100 μm/s wide cloud can be extracted from a large ($\gtrsim 10^{10}$ atom) μK-temperature thermal cloud by applying a velocity-selective cut using Doppler sensitive two-photon transitions [97]. Alternatively, the necessary velocity spread can be achieved by adiabatically expanding a cold atom ensemble in a harmonic trap or by implementing delta kick cooling [98], both of which are aided by being in microgravity, or by cooling the atoms in an optical lattice [99]. Adiabatic expansion of a cold atom ensemble in an all-optical trap, which requires only simple cavities and low power lasers, can be realized with a setup similar to that described in [100]. In any of these cases, low densities are desirable to mitigate possible systematic noise sources associated with cold collisions. If necessary, an array of multiple atom sources can be placed in each satellite in order to meet the atom flux requirement.

After preparation, the cold atom ensemble must be positioned outside the satellite in the interferometer region (see Figs. 3.1 and 3.10). This can be accomplished using laser light along the primary interferometer beam axis. Counter-propagating light from two satellites is used to form an optical standing wave, trapping the cold atoms. Such an optical lattice trap can then shuttle the atom cloud out to the appropriate initial position [27].

In leader-follower low Earth orbits, positioning the atoms outside the satellites must account for significant Coriolis forces. For example, moving an atom cloud 10 m from the satellite in a time of 1 s can lead to transverse accelerations of $\left|\mathbf{v} \times \overrightarrow{\Omega}\right| \sim 10^{-3} g$. Using a red–detuned optical lattice can provide transverse confinement during cloud positioning, but this may require small lattice beam diameters, potentially limiting the baseline due to laser light diffraction effects. To avoid this baseline constraint, a more tightly focused beam may be sent out from each satellite in addition to the larger diameter optical lattice, acting as a source of transverse confinement for the atoms near the satellite from which the small beam originates.

3.3.2 Atom Optics

Conventional Raman atom optics transfer momentum $\hbar k_{\text{eff}} = 2\hbar k$ to the atom at each beamsplitter [1]. To achieve the desired sensitivity, it is necessary to increase k_{eff} using large momentum transfer (LMT) beamsplitters. An Nth order LMT beamsplitter (corresponding to an $\hbar k_{\text{eff}} = 2N\hbar k$ momentum separation between the arms of an interferometer) yields an N-fold enhancement of the phase difference in Eq. (3.1). LMT beamsplitters can be achieved by adiabatically accelerating the atoms with an optical lattice [101, 27, 102, 28] or by applying sequential two-photon Bragg pulses.

The AGIS-LEO design calls for $200\hbar k$ LMT beamsplitters. This exceeds the current state-of-the-art for a beamsplitter implemented in an interferometer by about an order of magnitude [28]. However, the coherent transfer of more than $1000\hbar k$ to an atom has already been demonstrated using an accelerated optical lattice [103].

In addition to the desired stimulated transitions, the atoms can undergo unwanted spontaneous two-photon transitions when in the presence of the laser field, consisting of the stimulated absorption of a photon from the laser field followed by the spontaneous emission of a photon in a random direction. Atoms that undergo spontaneous emission are lost and do not contribute to the interferometer signal, so this effect must be minimized. The spontaneous emission rate can be written as $R_{\text{sc}} = \frac{2\Omega_{\text{st}}^2}{\Gamma(I/I_{\text{sat}})}$, where $\Omega_{\text{st}} = \frac{\Gamma^2(I/I_{\text{sat}})}{4\Delta}$ is the Rabi frequency of the desired stimulated two-photon transitions, I is the laser intensity, Γ is the spontaneous decay rate of the excited state, Δ is the laser detuning, and $I_{\text{sat}} = \frac{\hbar\omega\Gamma}{2\sigma_0}$ is the saturation intensity of the atomic transition with resonant cross section $\sigma_0 \sim \lambda^2$ and frequency $\omega = 2\pi c/\lambda$. The time t_N needed for an Nth order LMT beamsplitter is given by $t_N \sim N\frac{\pi}{\Omega_{\text{st}}}$. In order for spontaneous emission not to degrade the signal significantly for q concurrent interferometers, we therefore require that $qN\frac{\pi}{\Omega_{\text{st}}} \lesssim \frac{1}{R_{\text{sc}}}$. For example, for rubidium atoms ($I_{\text{sat}} \approx 2.5\ \frac{\text{mW}}{\text{cm}^2}$ and $\Gamma \approx 2\pi \cdot 6.066$ MHz) we can meet this requirement for $N \sim 100$ LMT beam splitters and $q \sim 10$ concurrent interferometers by choosing $\Omega_{\text{st}} \sim 2\pi \times 0.4$ kHz and laser powers of ~ 1 W with a beam waist of 10 cm.

In addition to minimizing spontaneous emission, we must ensure that noise in the differential phase between interferometers due to AC Stark shifts [93] (including those due to off-resonant two-photon transitions) is below 10^{-4} rad/$\sqrt{\text{Hz}}$. This can be achieved with the parameters given above by stabilizing fluctuations in the laser intensity at the $\sim 0.01\%$ level. Off-resonant shifts are discussed in detail in [101].

When multiple concurrent Doppler-multiplexed interferometers are used, for certain parameters it may be optimal to choose the launch momentum difference between successive interferometers to be smaller than the momentum splitting between the arms of a single interferometer. If this is the case, then there will be instances when two arms of two different interferometers will cross each other in momentum space. The existence of such crossings will require only minor alterations to the pulse sequence to avoid interactions with the undesired arm.

3.3.3 Detection

A normalized phase measurement at the end of a pulse sequence consists of counting the number of atoms at each output port of the interferometer. Since the interferometry takes place outside the satellite, this counting must be done remotely. Fluorescence imaging, commonly used in laser cooling and trapping experiments, is an inherently remote state-sensitive atom-counting technique. A fluorescence-based detection scheme can be implemented with CCD cameras on the satellites and the same telescopes used to enlarge the primary atom interferometry beams (see Sec. 3.4.1). For instance, a detection beam tuned to resonance with an optically allowed atomic transition can be generated and enlarged to a 10 cm waist with the telescope on one satellite. Backscattered fluorescent light can be collected using the same telescope.

The details of such a scheme are as follows. Prior to detection, atoms exiting a port of the atom interferometer can be transferred from the interferometer atomic state (one of the ground state hyperfine levels) to the imaging state (the other ground state hyperfine level) using a Doppler sensitive two-photon Raman transition. Here the Doppler sensitivity of the transition is used to distinguish the output ports of the interferometer, since the atoms in the different output ports have different longitudinal velocities (see for instance, Fig. 3.3). A subsequent short detection pulse (~ 1 ms) (tuned to be resonant with the transition between the imaging state and an optical excited state) scatters fluorescent photons. A fraction of these photons is collected by the imaging optics and focused onto a CCD detector. With $N_a \sim 10^8$ atoms in the cloud, a phase sensitivity of $\sim 10^{-4}$ rad requires the ability to measure changes in atom number at the $\sqrt{N_a} \sim 10^4$ level (the atom shot-noise limit). This limit can be achieved by scattering enough photons from the atomic ensemble so that at least 1 photon per atom is detected. The required number of scattered photons is determined by the solid angle subtended by the light collecting optics, which for a 30 cm primary telescope mirror at a distance of 10 m from the atom cloud is $\Delta\Omega \sim 6 \times 10^{-5}$ sr. For the above parameters, the detection laser must have $I/I_{\text{sat}} \sim 10$ to achieve 1 photon per atom in 1 ms. The outputs from multiple Doppler-multiplexed interferometers run concurrently can be detected with a sequence of appropriately tuned velocity-selective Raman pulses. Finally, the fact that the detection is spatially resolved allows for characterization of possible optical wavefront curvature errors in the sensor (see

Parameter	Specification	Location in Text
Power	~ 1 W	Atom Optics (Sec. 3.3.2)
Waist	~ 10 cm	Atom Optics & Laser Pointing (Secs. 3.3.2 & 3.5.1)
Intensity Stability	$< 0.01\%$	Atom Optics (Sec. 3.3.2)
Wavefront	$< \frac{\lambda}{1000}$	Wavefront Aberrations (Sec. 3.4.2)
Phase Noise	$< -120\ \frac{\text{dBc}}{\text{Hz}}$	Laser Phase Noise (Sec. 3.4.3)
Frequency Stability	$< 10^{-15}$	Laser Phase Noise (Sec. 3.4.3)

Table 3.2: Interferometer laser pulse requirements.

Sec. 3.4.2).

Having shown that it is possible to produce, maintain and detect the necessary cold atom flux, we now consider the noise sources that may compete with gravitational waves at this level of phase sensitivity. These sources fall into two general categories: noise of the laser beam wavefronts and noise from the environment. The former affects the atoms' reference frame while the latter affects the atoms' motion relative to that reference frame.

3.4 Atom Interferometer Laser Beam Considerations

The laser wavefronts form the reference frame against which the motion of the atoms is compared. Thus, it is important that any non-common noise in the wavefronts that is imprinted on the interferometers falls below the phase sensitivity limit in the frequency band of interest. In this section, we discuss telescope design, wavefront aberrations, and laser phase noise as they relate to this requirement. For a summary of the resulting interferometer laser pulse parameters in particular, see Table 3.2.

3.4.1 Telescope Design

To generate interferometry beams with Rayleigh ranges equal to half our satellite separations (see Sect. 3.5.1), we will need telescopes to magnify the beam waists to $5 - 10$ cm. These beam waists should also be large compared to the typical sizes of the atom clouds, ~ 1 cm FWHM. The telescopes serve multiple purposes in the instrument: their beam line can be used to shuttle atoms into the interferometry region before a measurement (Fig. 3.1), they magnify the primary atom optics beams (Sec. 3.3.2), and they form a part of the detection imaging system (Sec. 3.3.3). As discussed in the remainder of this section, the atom optics beams place the most stringent requirements on the telescope design. Each beam is used in two atom interferometers, one a few meters outside the laser's originating satellite, and the other ~ 30 km away. Thus, the telescope's design must account for both small-scale imperfections and large-scale aberrations. For instance, the secondary mirror must not obstruct the path of the beam and the primary mirror must be large enough to avoid aperture diffraction effects. A ~ 30 cm diameter off-axis Schiefspiegler or off-axis Gregorian telescope could

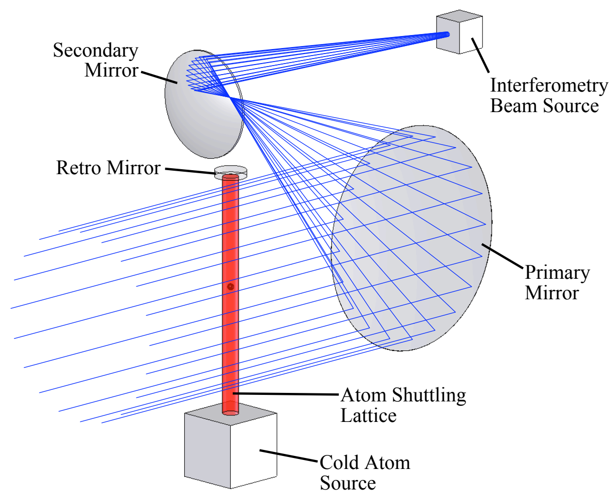

Figure 3.10: Conceptual diagram of a single AGIS telescope as a 30 cm off-axis Gregorian system. Note the intermediate focus of the Gregorian telescope, which allows for the use of pinhole spatial filtering to remove wavefront aberrations from upstream optics and lasers. Also shown approximately to scale are the cold atom source and a lattice beam that could be used to shuttle the atoms into the middle of the primary interferometry beam. The primary beam line can in turn be used to shuttle the atoms away from the satellite prior to an interferometer sequence.

be capable of meeting these requirements [104, 105]. A Gregorian telescope has the additional advantage that its real intermediate focus could be used for pinhole spatial filtering, which would eliminate wavefront errors from all optics and lasers before the primary mirror (see Fig. 3.10).

3.4.2 Wavefront Aberrations

Wavefront aberration of the interferometer laser beams is a potential source of noise. Conceptually, if the interferometer laser beam were an ideal plane wave, then the phase shift measured by the atom interferometer would only depend on the longitudinal position of the atom along the beam and would be insensitive to any transverse position offsets. However, in addition to having a Gaussian spatial profile, the physical laser beam inevitably contains additional aberrations that cause the transverse phase to be nonuniform. Since the local phase of the laser is imprinted on the atom during each interferometer light pulse, any transverse phase variation of the beam results in a phase shift that depends on the transverse position of the atom with respect to the laser beam spatial profile. This effect can result in phase noise (i) due to satellite transverse position noise during the interferometer pulse sequence, (ii) due to fluctuations in the mode itself induced by temporal variations in the optics path or (iii) due to jitter in initial atom cloud position or width from shot to shot.

It is reasonable to guess that this source of phase noise might be suppressed by the differential measurement strategy employed by AGIS-LEO. Naively, both atom interferometers are manipulated by the same laser beam (see Fig. 3.1), so any phase aberration will affect both atom interferometers symmetrically. However, the effect is in general not common-mode suppressed since wavefront perturbations with high transverse spatial frequency diffract out of the beam over distances smaller than the satellite separation. As a result, the transverse spatial profile of the beam can differ substantially between the two atom interferometers. The perturbations only become common-mode suppressed at transverse wavelengths λ_t longer than the corner wavelength $\lambda_c = \sqrt{L\lambda} \sim 17\ \text{cm}\left(\frac{L}{30\ \text{km}}\right)^{1/2}$, where λ_c is defined as the transverse spatial length scale which has a collimation distance (i.e., Rayleigh range) equal to the baseline separation L.

Wavefront aberration noise is complicated by the fact that the ensemble of atoms has a finite spatial extent, so different parts of the atom cloud generally experience different local laser phases in the aberrated beam. At the end of the atom interferometer sequence, the final atom ensemble contains a distribution of phase shifts from which a single phase measurement, corresponding to the GW signal, must be extracted. In the simplest case, the interferometer phase measurement is determined by taking the average phase of the ensemble. Spatially averaging over the atom cloud leads to a suppression of aberrations with transverse wavelengths smaller than the size of the cloud. However, by performing this average, the phase measurement becomes sensitive to changes in the atom cloud size from shot to shot, a noise source which we also discuss below.

Wavefront Aberration Phase Shift Calculation

The phase shift imparted to the atoms as a result of non-uniform transverse spatial phase of the laser wavefronts can be calculated by summing the local phase $\phi(x, z, t)$ imparted to the atom at each of the light pulses:

$$\Phi_t = N \sum_{j=1}^{n} \kappa_j\, \phi(x_j, z_j, t_j) = N \sum_{j=1}^{n} \kappa_j\, \theta(x_j, k_t, t_j) \sin[k_t(z_j - \delta z(t_j)) + \Theta_t] \tag{3.2}$$

Here the position of the atom is referenced to the coordinate system shown in Fig. 3.1, where x_j is the longitudinal position of the atom along the direction of k_{eff} at the time t_j of the jth pulse and z_j is the transverse position at time t_j (for simplicity we suppress the transverse coordinate y_j). The sum is taken over the n light pulses of the interferometer sequence; in the cases considered here $n = 5$ and the coordinates (x_j, z_j, t_j) are taken to be those shown in Fig. 3.3. Each phase $\phi(x_j, z_j, t_j)$ is weighted by the normalized wavevector κ_j for the pulse (see Fig. 3.3), and $N = k_{\text{eff}}/2k$ is the order of the LMT atom optics. For simplicity, we consider the response to a single Fourier component of amplitude $\theta(x, k_t, t)$ and transverse spatial wavevector $k_t = 2\pi/\lambda_t$. The general phase Θ_t allows modeling of both quadratures. The time-dependent reference position $\delta z(t)$ accounts for vibrational noise that affects the relative transverse position of the atom and the satellite. In the paraxial limit, the amplitude of each Fourier component evolves as the beam propagates along the x direction according to the paraxial propagator given by $\theta(x, k_t, t) = \theta(k_t, t) \exp\left[-i\frac{k_t^2}{2k}(x - x_0)\right]$

where x_0 is the reference plane at which $\theta(k_t, t)$ is known. Finally, in order to model time-dependent wavefront aberrations, we decompose each Fourier coefficient into a static and a dynamic component: $\theta(k_t, t) \equiv \theta_0(k_t) + \delta\theta(k_t, t)$. Here the term $\delta\theta(k_t, t)$ represents an explicitly time-dependent variation of the amplitude of the aberration with transverse wavevector k_t.

The phase shift Φ_t given by Eq. (3.2) depends on the initial coordinates (x_1, z_1, t_1) of the atom cloud at the first laser pulse. Since the atom wavefunction is spatially delocalized, the final phase shift will vary with position across the atom. These position dependent phase shifts lead to an atom probability distribution (in one of the output ports) that may be written as $|\Psi(\mathbf{x})|^2 = A(\mathbf{x}) + B(\mathbf{x}) \cos(\phi_0 + \phi(\mathbf{x}))$, where $\phi(\mathbf{x})$ is the spatially varying phase shift. Assuming $|\phi(\mathbf{x})| \ll 1$, the total probability of detecting an atom in this port then depends on the spatial average of the phase over the atom wavefunction:

$$P = \int_{-\infty}^{\infty} |\Psi(\mathbf{x})|^2 dx \approx A + B \cos(\phi_0 + \langle \phi(\mathbf{x}) \rangle) \qquad |\phi(\mathbf{x})| \ll 1 \tag{3.3}$$

where $A \equiv \int_{-\infty}^{\infty} A(\mathbf{x}) dx$, $B \equiv \int_{-\infty}^{\infty} B(\mathbf{x}) dx$, and $\langle \phi(\mathbf{x}) \rangle \equiv \frac{1}{B} \int_{-\infty}^{\infty} B(\mathbf{x}) \phi(\mathbf{x}) dx$.

For this noise analysis, we want to calculate the differential response between two atom interferometers separated by a distance L along the laser propagation direction (see Fig. 3.1). This gradiometer phase is denoted $\Delta\Phi_t \equiv \langle \Phi_t(x_1 = L) \rangle - \langle \Phi_t(x_1 = 0) \rangle$, where the expression given in Eq. (3.2) has been evaluated at initial positions 0 and L along the x axis and the angle brackets indicate a spatial average over the atom distribution. Note that in practice the spatially resolved atom detection scheme discussed above (see Sec. 3.3.3) allows direct resolution of the transverse spatial dependence of the phase shift. In this section, we assume that we simply average over this information in order to set conservative upper limits on the allowed aberrations of the optical system.

Since we are only concerned with calculating time-dependent phase shifts originating from the jitter amplitudes $\delta z(t)$ and $\delta\theta(k_t, t)$, we may ignore any constant phase shift piece in Eq. (3.2) that does not depend on either jitter amplitude. To enforce this condition, we consider only the first-order jitter susceptibilities $\partial(\Delta\Phi_t)/\partial(\delta z)$ and $\partial(\Delta\Phi_t)/\partial(\delta\theta)$. We then consider the response at frequency ω of Eq. (3.2) for the position and phase jitter amplitudes, $\delta z(t) = \widetilde{\delta z}(\omega) e^{-i\omega t}$ and $\delta\theta(k_t, t) = \widetilde{\delta\theta}(k_t, \omega) e^{-i\omega t}$. Assuming that the system can be linearized for small jitter, the transfer function $H_n(\omega)$ that relates each noise source $\delta_n(t)$ to the phase error is simply the first-order susceptibility evaluated at zero jitter amplitude.

With these transfer functions, we can relate the amplitude spectral density of the measured phase $\overline{\delta\phi}$ to the amplitude spectral density of noise sources $\overline{\delta_n}$. In general, the spectral densities are matrices, but because the measured phase is a scalar quantity and the noise sources are assumed to be uncorrelated, the relationship reduces to

$$\overline{\delta\phi}^{\,2} = \sum_n |H_n(\omega)|^2 \, \overline{\delta_n}^{\,2} \tag{3.4}$$

where $|H_n(\omega)|$ is the magnitude of the transfer function. As noted above, the transfer function $H_n(\omega)$ of the particular noise source $\delta_n(t)$ takes the form $\partial(\Delta\Phi_t)/\partial(\delta_n)|_{\{\delta_n\}=0}$ in this analysis.

By taking the square root of Eq. (3.4), we find the amplitude spectral density of the phase noise imprinted on the atom (in [rad/$\sqrt{\text{Hz}}$]) in terms of the amplitude spectral density of the transverse position jitter $\overline{\delta z}(\omega)$ (in [m/$\sqrt{\text{Hz}}$]) and of the wavefront variation $\overline{\delta\theta}(k_t,\omega)$ (in [rad/$\sqrt{\text{Hz}}$]):

$$\begin{aligned}\overline{\delta\phi}(k_t,\omega) &\equiv \sqrt{\left(\left|\frac{\partial(\Delta\Phi_t)}{\partial(\delta z)}\right|\overline{\delta z}(\omega)\right)^2+\left(\left|\frac{\partial(\Delta\Phi_t)}{\partial(\delta\theta)}\right|\overline{\delta\theta}(k_t,\omega)\right)^2} \\ &= 2\,N\sqrt{\left(\tfrac{2\pi}{\lambda_t}\theta(k_t)\overline{\delta z}\right)^2+\overline{\delta\theta}^2}\,\sin\left[\frac{\pi L\lambda}{2\lambda_t^2}\right]\exp\left[-\frac{\pi^2\sigma^2}{4\lambda_t^2\ln 2}\right] \\ &\times\left(5-9\cos\left[\frac{2\pi\Delta z}{\lambda_t}\right]\cos[2\omega T]+4\cos[3\omega T]\right)\end{aligned} \tag{3.5}$$

where the noise contributions from transverse motion and dynamic wavefront variation have been summed in quadrature since they are assumed to be uncorrelated. The atom position distribution is taken to be a Gaussian with width σ (FWHM), $N = k_{\text{eff}}/2k$ is the order of the LMT beamsplitter, and T is the interrogation time. The Coriolis-induced transverse position separation of the upper and lower interferometer paths is given by Δz. A five-pulse sequence is assumed in obtaining Eq. (3.5). Finally, Eq. (3.5) represents the magnitude of the response with respect to the aberration phase Θ_t (i.e., the quadrature sum of both sine and cosine aberrations).

We note that as $\Delta z \to 0$, the frequency dependence of the wavefront noise given in Eq. (3.5) is identical to the GW phase response for a five-pulse sequence (see Eq. (3.1)). This is expected since both effects arise from what amounts to a time variation in the local phase of the laser. However, as opposed to the GW signal, the response to wavefront aberration is suppressed to the extent that the aberrations are common between the two interferometers.

Transverse Vibration Induced Phase Noise

Here we focus on the wavefront-induced noise due to stochastic variation of the relative transverse position δz of the atom with respect to the aberrated beam during the interferometer pulse sequence. This type of transverse position jitter may result from variations in the satellite position. Sensitivity to transverse position jitter of the satellite is given by Eq. (3.5) with $\overline{\delta\theta}(\omega) = 0$. In order for this noise source to be less than the atom shot noise $\overline{\delta\phi}$, the maximum allowable wavefront aberration (in fractions of an optical wave) is given by

$$\frac{\delta\lambda}{\lambda} \equiv \frac{\theta_0(k_t)}{2\pi} = \frac{\overline{\delta\phi}\,\lambda_t}{N\pi^2\,\overline{\delta z}}\csc\left[\frac{\pi L\lambda}{2\lambda_t^2}\right]\exp\left[\frac{\pi^2\sigma^2}{4\lambda_t^2\ln 2}\right]\left(62+63\cos\left[\frac{2\pi\Delta z}{\lambda_t}\right]\right)^{-1} \tag{3.6}$$

assuming $k_{\text{eff}} = 2Nk = 200k$ LMT atom optics, $\overline{\delta\phi} = 10^{-4}$ rad/$\sqrt{\text{Hz}}$, and motion δz at the GW peak sensitivity frequency $\omega_c = \frac{2}{T}\cos^{-1}\sqrt{\frac{3}{8}}$. Figure 3.11(a) shows the general wavefront constraint

Parameter	Specification	Location in text
Transverse Position	< 10 nm/$\sqrt{\text{Hz}}$	Wavefront Aberration (Sec. 3.4.2)
Angle Jitter	$\lesssim 1$ nrad/$\sqrt{\text{Hz}}$	Laser Pointing Angle Jitter (Sec. 3.5.1)
Angular Rate	$\lesssim 1$ μrad/s	Rotational Effects (Sec. 3.5.1)

Table 3.3: Satellite position and angle control requirements. These requirements form constraints on one satellite's laser position and angle relative to the counter-propagating beam from the opposing satellite. We therefore can split the requirements into a coarse constraint on the satellite and a fine constraint on the laser's mounting platform. The constraints apply to both leader-follower and inclined-great-circle orbits.

for all spatial wavelengths due to δz vibration with amplitude $\overline{\delta z} = 10$ nm/$\sqrt{\text{Hz}}$ for both the leader-follower and inclined-great-circle orbits. Note the expected exponential relaxation of the wavefront constraint due to spatial averaging for $\lambda_t < \sigma$ as well as the expected corner wavelength at $\lambda_t \sim \sqrt{L\lambda}$ above which the perturbations start to be common for the two interferometers. For wavelengths in the intermediate region $\sigma < \lambda_t < \lambda_c$, which are not suppressed by spatial averaging and are not common between the interferometers, the maximum allowable wavefront aberration in waves is $\delta\lambda \sim \frac{\lambda}{1000}\left(\frac{\lambda_t}{\text{cm}}\right)\left(\frac{10\ \text{nm}/\sqrt{\text{Hz}}}{\overline{\delta z}}\right)$. For the leader-follower orbit, however, the larger transverse position separation ($\Delta z \sim 4$ cm versus 80 μm) results in resonant cancellation of certain transverse wavelengths. In the vicinity of these particular wavelengths, the wavefront requirements are significantly relaxed.

Dynamic Wavefront Induced Phase Noise

Temporal fluctuations of the spatial mode can also be a problem if the mode varies at frequencies within the GW detection band. The phase noise induced by such dynamical wavefront variations is given by Eq. (3.5) with $\overline{\delta z}(\omega) = 0$. Figure 3.11(b) shows the required constraint on the fluctuation amplitude $\overline{\delta\theta}(k_t, \omega_c)$ in order for this noise source to be less than the atom shot noise. To be maximally conservative, the wavefront fluctuation is assumed to be at the GW peak sensitivity frequency ω_c.

To put the dynamical wavefront constraints given in Fig. 3.11(b) into perspective, we consider several well-known sources of wavefront fluctuation associated with aberrations caused by the mirrors in the optics system [106]. To be concrete, we model the wavefront aberrations caused by vibrations and thermal fluctuations of the 30 cm primary telescope mirror. The primary mirror is significant since it is the last element that the beam encounters before propagating to the atoms, and no wavefront spatial filtering can be done after this element. We emphasize that the same analysis can easily be extended to other elements in the optics system as necessary.

Fluctuations of wavefront aberrations are a source of stochastic noise in the interferometer. As a result, they contribute to the strain noise floor of the gravitational wave detector. Figure 3.12 shows the effective strain noise contributed by the various sources of dynamical wavefront aberrations that

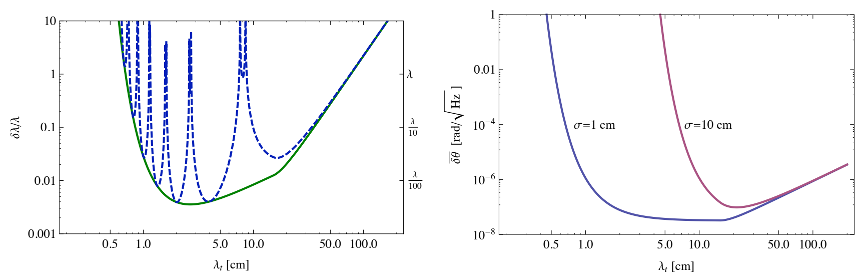

(a) (b)

Figure 3.11: Transverse spatial phase profile requirement as a function of transverse wavelength. Both figures assume the five-pulse interferometer of Fig. 3.3 with $200\hbar k$ atom optics, $L = 30$ km baseline, and a $\overline{\delta\phi} = 10^{-4}$ rad/$\sqrt{\text{Hz}}$ phase noise requirement. (a) Static wavefront requirements due to satellite transverse motion with amplitude $\overline{\delta z} = 10$ nm/$\sqrt{\text{Hz}}$ at frequencies within the GW detection band versus spatial wavelength of the perturbation. The atom cloud size is taken to be $\sigma =$ 1 cm (FWHM). The blue (dashed) curve assumes a leader-follower orbit, while the solid (green) curve assumes an inclined-great-circles orbit. The large notches in the dashed (blue) envelope arise from resonant cancellations of the wavefront aberration with the large transverse separation of the leader-follower interferometer paths. Both curves have been partially enveloped. (b) Temporal stability requirements for wavefront perturbation amplitude versus spatial wavelength of the perturbation. This constraint applies only to perturbation amplitude variations that occur at temporal frequencies within the GW detection band. The pair of curves shows the effect of averaging over the atom cloud size. As labeled, the curves assume atom cloud widths of $\sigma = 1$ cm (FWHM) and $\sigma = 10$ cm (FWHM). Both curves have been enveloped.

we model. For comparison, the atom shot-noise limited strain sensitivity from Fig. 3.4 is reproduced in blue (labeled “Shot”).

The first effect we consider is dynamic wavefront aberrations caused by vibrations of the primary telescope mirror. We consider a mirror modeled as a solid cylindrical disk with diameter $D = 30$ cm. The mirror surface is nominally polished so as to produce a collimated Gaussian beam upon reflection (see Fig. 3.10 for example). However, the surface of the mirror can become deformed as a result of vibrations of the mirror substrate, and these aberrations will be imprinted onto the laser beam as a spatially dependent phase. It is convenient to decompose the mirror vibrations into a set of normal modes and consider the effect of each mode separately. For example, the lowest frequency mode corresponds to the bowl-shaped deformation given by $\delta x(r) = \delta x_0 J_0(2 j_{0,1} r/D)$, where J_0 is the zeroth order Bessel function of the first kind, and $j_{0,1}$ is the first zero of J_0. (For simplicity, we assume Dirichlet boundary conditions on the mirror edge.) In order to leverage the results of Fig. 3.11(b), this deformation can be rewritten as a discrete Fourier series. Since the first term in the series has a 97% overlap with the deformation, the mirror surface profile is well-approximated by the truncated series $\delta x(r) \approx \delta x_0 \cos(\pi r/D)$. Therefore, excitations of the lowest order normal mode of the mirror correspond to dynamic wavefront perturbations with transverse spatial wavelength $\lambda_t = 2D \approx 0.6$ m in Fig. 3.11(b).

The normal modes of the mirror can be excited by vibration of the satellite. We conservatively assume a white noise acceleration spectrum of $\overline{\delta a} = 10^{-7} g/\sqrt{\text{Hz}}$ which acts as a driving force for this motion [107]. The vibration response to acceleration of one of the normal modes is a narrow Lorentzian centered at the resonant frequency of the mode:

$$V_n(\omega) \equiv \frac{\delta x_0}{\delta a} = \frac{1}{\sqrt{(\omega_n^2 - \omega^2)^2 + (2\gamma_n \omega_n)^2}} \tag{3.7}$$

Here γ_n is the (structural) damping of the mode with resonance frequency ω_n. Even the lowest of the mirror resonances (corresponding to the $\lambda_t = 2D$ mode discussed above) tends to be relatively high frequency ($\omega_n > 100$ Hz) and sharply peaked ($Q_n \equiv \frac{1}{2\sqrt{3}\gamma_n} > 10^2$), so the response of the mirror surface in the GW detection bandwidth is relatively flat, corresponding to the tail of the Lorentzian. The phase noise induced by these off-resonant mechanical vibrations is given by Eq. (3.5) with $\overline{\delta z}(\omega) = 0$ and $\overline{\delta\theta} = k\, V_n(\omega)\overline{\delta a}$. The resulting effective strain noise from mechanical vibrations of the lowest mirror mode (with $\lambda_t = 2D$) is shown in red (labeled “Mechanical”) in Fig. 3.12.

The effect of higher order normal mode oscillations is generally suppressed since these modes have correspondingly higher resonance frequencies, and the size of the mechanical response in the GW band scales as $\sim \omega_n^{-2}$. However, since they are also associated with smaller λ_t, these modes can induce a larger intrinsic response (see Fig. 3.11(b)). Specifically, for long spatial wavelengths $\lambda_t > \lambda_c$, the phase response scales as λ_t^{-2}. Assuming a linear dispersion relation for the mirror modes, the transverse wavelength of the mode goes as $\lambda_t \sim \omega_n^{-1}$. As a result, the first several

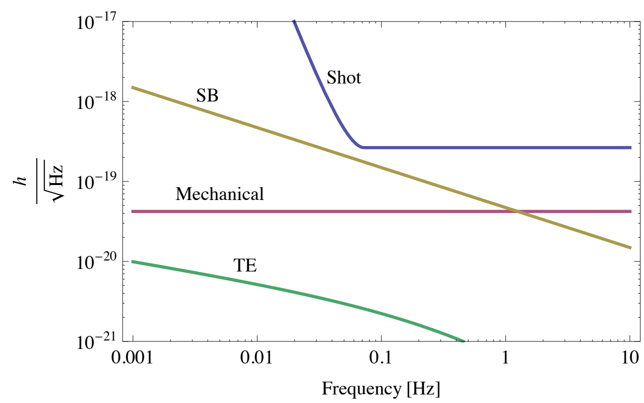

Figure 3.12: Effective strain induced by several sources of wavefront noise. Each of the effects is caused by a phase aberration that is imparted to the laser beam as it reflects off the final telescope mirror. “Mechanical” (red) is the off-resonant mechanical acceleration noise of the mirror assuming a normal mode with $\omega_n = 2\pi \times 2$ kHz and $Q_n = 10^2$ and a white noise acceleration amplitude spectral density of $10^{-7} g/\sqrt{\mathrm{Hz}}$. “SB” (yellow) is substrate Brownian noise at temperature $\tau = 300$ K for a $M = 1$ kg mirror assuming a normal mode with $\omega_n = 2\pi \times 2$ kHz, $Q_n = 10^2$, and with effective mass fraction $\alpha_n = 10^{-3}$. “TE” (green) is thermoelastic noise for a beryllium mirror substrate at temperature $\tau = 300$ K assuming a Gaussian deformation scale factor of $w_0 = 1$ cm.

normal modes make roughly equal contributions to the noise floor. Beyond this, higher frequency normal modes with transverse wavelengths $\lambda_t < \lambda_c$ are suppressed by the decreased mechanical response and (eventually, for $\lambda_t < \sigma$) by spatial averaging over the atom cloud.

In addition to external mechanical excitation, the normal modes of the mirror also vibrate as a result of the finite temperature of the substrate[106, 108]. These thermally excited mirror vibrations are known as substrate Brownian (SB) noise. The phase noise caused by substrate Brownian noise is given by Eq. (3.5) with $\overline{\delta z} = 0$ and $\overline{\delta \theta} = \delta\theta_{\mathrm{SB}}$ where

$$\delta\theta_{\mathrm{SB}} = k\,\overline{\delta x}_n \approx k\sqrt{\frac{4k_B\tau}{M\alpha_n\omega_n^2}}\sqrt{\frac{\gamma_n}{\omega}} \tag{3.8}$$

and where $\overline{\delta x}_n$ is the amplitude spectral density of the vibration amplitude of the mirror mode with resonance frequency ω_n, structural damping γ_n, and effective mass $M\alpha_n$ [108]. This result follows from the fluctuation dissipation theorem for a mirror substrate of temperature τ [109]. Here M is the mass of the mirror itself and α_n quantifies the effective fraction of the mirror mass that oscillates for a particular mode. Values for the effective mass $M\alpha_n$ of the mode depend on the details of the mode shape as well as material properties of the substrate and the geometry of the mirror. As an example, the values of α_n reported by [108] range from 10^{-3} to 10 for a 10 cm diameter, 8.8 cm thick fused silica mirror with mass $M = 1.6$ kg. For a conservative estimate, we take $M = 1$ kg and $\alpha = 10^{-3}$. The resulting effective strain noise caused by substrate Brownian noise of the lowest mirror mode (with $\lambda_t = 2D$) is shown in yellow (labeled 'SB') in Fig. 3.12. As with mechanical noise above, the first few higher mirror modes (which have smaller λ_t) are potentially of comparable size since the phase response scales as λ_t^{-2}. However, the overall response for the different modes will depend on the detailed scaling of α_n for the specific mirror geometry, and this is beyond the scope of our analysis.

Thermal fluctuations in the mirror substrate can also cause wavefront aberrations by inducing non-uniform thermal expansion, an effect known as thermoelastic noise [110, 106, 109]. This effect depends strongly on the material properties of the mirror substrate, such as the coefficient of thermal expansion, thermal conductivity, density, and specific heat. As an example, we consider the thermoelastic noise for a beryllium mirror at 300 K. Note that this is a conservative estimate, since the James Webb Space Telescope (JWST) will use a beryllium mirror at 30 K, and the lower temperature significantly reduces thermoelastic noise [111]. The power spectral density of thermoelastic vibrations has previously been calculated for a Gaussian deformation with scale factor w_0 [106, 110]. To be maximally conservative, we choose the scale factor to be $w_0 = 1$ cm since this maximizes the thermoelastic amplitude spectral density in the GW frequency band for the beryllium mirror. Represented in the Fourier domain, a Gaussian surface deformation contains a distribution of transverse wavelength components $\lambda_t \gtrsim w_0$. The total phase noise response to the Gaussian deformation is then the sum of each spectral component weighted by the amplitude of the phase noise spectral response $\overline{\delta\phi}(k_t, \omega)$ given by Eq. (3.5). As a conservative simplification, we take the value of the peak response of $\overline{\delta\phi}(k_t, \omega)$ which occurs at $\lambda_t \approx 10$ cm as the response for the entire Gaussian spectrum. The resulting effective strain due to thermoelastic noise is shown in green (labeled "TE") in Fig. 3.12.

The dynamic wavefront noise effects considered so far are all caused by a phase aberration that is imparted to the laser beam as it reflects off the final telescope mirror. Alternatively, wavefront noise may be present on the beam as a result of the intrinsic limitation of the spatial mode purity of the laser source. To eliminate this type of aberration, the laser beam can be sent through a high finesse mode-scrubbing optical cavity prior to delivery to the final telescope optics. In this case, the wavefront noise upon exiting the optical cavity will be limited by the same sources of dynamic wavefront noise shown in Fig. 3.12 as applied to the mirrors in the optical cavity, rather than by anything intrinsic to the laser.

Atom Distribution Induced Phase Noise

Quantum projection noise and variation in the preparation of the atom sample result in shot-to-shot fluctuations of the atom cloud size and its position in the interferometer beam. At zero temperature, if the two paths of the interferometer do not separate transversely, the cloud size variation induces no net phase error as both paths average over the same wavefront error at each pulse. Such a constant phase is perfectly cancelled by the pulse sequence. Due to finite temperature, however, the expansion of the cloud will result in a phase error since the cancellation will no longer be perfect between pulses. Moreover, the interferometer paths do separate in general, and each path explores a slightly different region of the wavefront aberration, resulting in non-zero phase shift. If the cloud size or transverse position changes from shot to shot, this phase-shift varies in time, mimicking a gravitational wave.

We compute the phase noise due to stochastic variations in the atom clouds' initial size and position. The size during the interferometer is determined by the linear expansion associated with the clouds' finite temperature. Assuming the typical FWHM cloud size and the corresponding amplitude spectral density for the two distant atom clouds are identical, but uncorrelated, and assuming that the amplitude spectral density of the transverse position is also identical and uncorrelated, we find that the maximum allowable wavefront aberration is proportional to $\csc[\pi \Delta z/\lambda_t]^2$. This scaling factor indicates that transverse wavelengths longer than the transverse separation will be quadratically suppressed, as expected. As previously mentioned, the cloud expands linearly once released from the trap due to its non-zero temperature. At temperatures of 100 pK and corresponding velocities of 100 μm/s, this effect is negligible.

The transverse trajectory deflection of the inclined-great-circles orbit is suppressed by $\sim L/R \sim 10^3$ compared to that of the leader-follower orbit (see Fig. 3.7). Consequently, for the same wavefront aberration, the inclined-great-circle orbit has a significantly relaxed cloud size and position constraint, as shown in Fig. 3.13. The size of the transverse wavepacket separation ($\Delta z = 4$ cm for the LF orbit, and $\Delta z = 80$ μm for the IGC orbit) is the only difference between the two curves. In both cases, the shot-to-shot atom cloud size and centroid position variation are assumed to be at the shot-noise limit. For an atom cloud of $N_a \sim 10^8$ atoms with size $\sigma = 1$ cm, the centroid transverse position uncertainty due to shot-noise is $\delta z \sim \sqrt{N_a}\sigma \sim 1$ μm. Likewise, the cloud size uncertainty is given by $\delta\sigma \sim \sqrt{N_a}\sigma \sim 1$ μm . These values represent a lower bound on the centroid position and cloud size variation from shot-to-shot. Since the results of Fig. 3.13 assume spatially averaged detection, the curves represent maximally conservative wavefront requirements that can be relaxed through the use of more sophisticated detection schemes (see Sec. 3.4.2).

Wavefront Aberration Noise Mitigation

In an alternative interferometer beam geometry, the atom optics laser beams can be made to first propagate between two satellite stations along a path that is displaced from the atoms before being redirected to interact with the atoms. As a consequence, the first propagation segment would serve as a spatial filter, allowing high frequency wavefront noise to diffract out of the beam. This arrangement is still sensitive to wavefront aberrations caused by the beam steering optics, but these are potentially better defined than the mode of the laser source. If needed, this alternative beam geometry could be used in conjunction with a mode-scrubbing cavity.

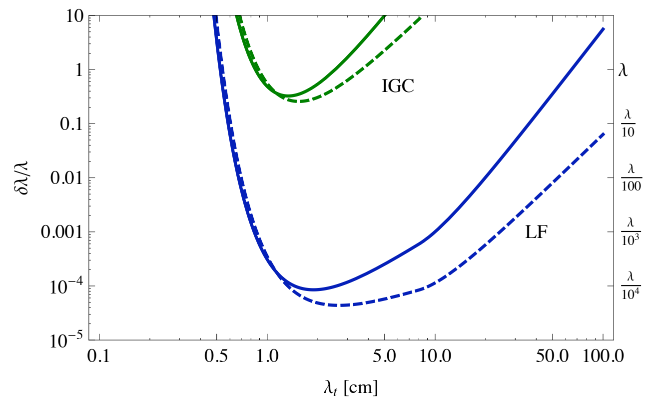

Figure 3.13: Transverse spatial phase profile requirement as a function of transverse wavelength for shot-to-shot noise sources. All curves have been enveloped and assume the five-pulse interferometer of Fig. 3.3 with $200\hbar k$ atom optics, an $L = 30$ km baseline, and a $\overline{\delta\phi} = 10^{-4}$ rad$/\sqrt{\text{Hz}}$ phase noise requirement. The clouds are assumed to be 1 cm FWHM, and the amplitude spectral densities of the cloud size and cloud centroid position noise are both 1 μm$/\sqrt{\text{Hz}}$ (This assumes a $f_r = 1$ Hz repetition rate with 10^8 atoms per shot). The curves associated with the leader-follower and inclined-great-circles orbits have been labeled "LF" (blue) and "IGC" (green) respectively. The solid curves are associated with fluctuations in initial atom cloud size, while the dashed curves are associated with fluctuations in initial cloud position in the beam. Note that the wavefront requirements shown here are maximally conservative; those from the LF curves in particular can be significantly mitigated by characterizing the wavefront and by reducing the transverse separation through multiple-pulse interferometer sequences (see Sec. 3.4.2).

Commercially available interferometers can measure wavefront aberration with RMS repeatability $\delta\lambda < \lambda/1000$ [112]. In principle, the aberration can be monitored in real time using such a wavefront sensor. Furthermore, we can leverage the fact that the detection is spatially resolved (see Sec. 3.3.3) to perform *in situ* measurements of the wavefront. If necessary, in addition to measuring the spatial phase distribution for each of the five-pulse sequences, we can interleave alternate pulse sequences tailored to be maximally wavefront-sensitive. The initial transverse velocity and the time between pulses can be varied to gain maximal information about the wavefront.

Spatially resolved detection of the atom cloud can help mitigate the wavefront requirements that result from spatial averaging (Fig. 3.13). In particular, the requirements shown in Fig. 3.13 for the leader-follower orbit exceed the range of current technologies [112]. However, these bounds are maximally conservative since with spatially resolved detection we can directly measure the shot-to-shot variations of the cloud size and centroid position in real time. The additional information provided by spatially resolved detection should allow the wavefront requirements shown in Fig. 3.13 to be relaxed by an amount that depends on how well the shot-to-shot wavefront characterization can be performed.

Finally, to lessen the impact of shot-to-shot noise sources, it is possible to use interferometer sequences with more pulses than considered here to reduce the transverse wavepacket separation. The quadratic scaling of the phase noise with respect to this separation makes it a useful tool for shot-to-shot phase noise management.

3.4.3 Laser Phase Noise

As the laser phase is the reference used to measure the atom motion, noise of the laser phase is imprinted as noise on the atom phase. We consider two sources of laser phase noise: noise of the laser phase $\delta\phi$ and noise δk of the center frequency of the laser. For both, the noise is significantly suppressed by using common lasers.

Despite the use of common lasers, the laser phase $\delta\phi$ is asymmetrically imprinted due to the finite light travel time $\frac{L}{c} \sim 100\ \mu$s between the two interferometers [34]. If the laser pulses are longer than the light travel time, however, the phase noise at only the beginning and end of the pulses will be non-common; the rest of the pulse overlaps in time and is common. Modeling the LMT pulse as a large number of consecutive pulses, we note that if the total pulse is longer than the finite light travel time, laser phase noise at frequencies lower than $\sim \frac{c}{L}$ will be suppressed.

In addition, each interferometer is insensitive to static phase offsets and will suppress phase noise at frequencies below $\sim \frac{1}{T}$. Moreover, the atoms average the laser phase at frequencies higher than the effective Rabi frequency of the LMT pulse, Ω_{st}/N, where N is the number of pulses. By calculating the quantum evolution of the phase of the atom states through the interferometer sequence, we can analytically find the transfer function $T_\phi(f)$ that relates laser phase noise to atom phase noise for a five-pulse sequence (see [113] for a similar calculation).

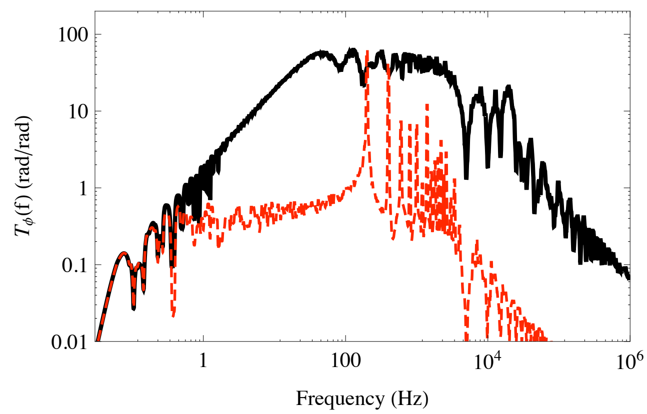

Figure 3.14: The transfer function relating laser phase noise to atom phase noise as a function of frequency. Many of the sharp oscillations have been enveloped. Both curves have fixed interferometer time $T = 4$ s and baseline $L = 30$ km. The solid (black) and dashed (red) curves represent interferometers that use pulses of Rabi frequency 10 and 0.1 kHz respectively; each LMT beamsplitter contains 100 of these pulses. Below $\frac{1}{T} = 0.25$ Hz, phase noise is suppressed due to the interferometer's insensitivity to static phase offsets and overlapping pulse suppression. The sensitivity to phase noise rises until the atom begins to average the phase roughly at the effective LMT Rabi frequency. This effective low pass filter balances the high pass filter of the overlapping pulse suppression, producing the flat region. Above the corner frequency of the overlapping pulse suppression, $\frac{c}{2\pi L} \sim 1.6$ kHz, the atom's averaging suppresses phase noise.

In Fig. 3.14 we show this transfer function for two sample Rabi frequencies $\Omega_{\rm st} = 10$ and 0.1 kHz for fixed baseline $L = 30$ km and interferometer time $T \sim 4$ s. Note the corner frequencies at $\frac{1}{T} = 0.25$ Hz, $\frac{\Omega_{\rm st}}{N\pi}$, and $\frac{c}{2\pi L} = 1.6$ kHz corresponding to the interferometer low-frequency suppression, the atom's LMT Rabi frequency suppression, and the overlapping pulse suppression respectively. In general, the transfer function will exhibit the sharp, narrow spikes that appear in the $\Omega_{st} = 0.1$ kHz curve. These resonances arise at multiples of Ω_{st} due to the accumulative imprint of phase at the beginning and end of each component pulse of the LMT. For special values of the Rabi frequency where Ω_{st} is a multiple or fraction of $\pi c/2L = 10$ kHz, these resonances match with nulls in the overlapping pulse suppression and are greatly reduced. This is shown in the $\Omega_{st} = 10$ kHz curve.

Drift δk in the central frequency of the laser between pulses differentially changes the position of the laser phase reference relative to the near and far interferometers. While the near interferometer sees negligible motion of the phase front, the far interferometer sees a wavefront shifted by $\delta k\, L$. Equivalently, this effect can be understood in the time domain as an extra phase accumulated at a rate δk during the time L/c during which the light pulse is not common for the two interferometers.

To place an upper bound on laser phase noise control constraints, consider the conservative Rabi frequency ~ 10 kHz . As the $\delta\phi$ and $\delta k\, L$ phase noise sources are uncorrelated, they must be added in quadrature as $\sim \sqrt{(\delta\phi)^2 + (\delta k\, L)^2}$. For this contribution to be $< 10^{-4}$ rad/$\sqrt{\text{Hz}}$, the phase noise of the laser must be $< -120\ \frac{\text{dBc}}{\text{Hz}}$ at an offset of ~ 10 kHz , and the fractional stability in the laser frequency must be $\sim 10^{-15}$ over time scales $\sim T$. Both requirements are experimentally achievable [114, 115].

Alternatively, in a three or more satellite configuration, two non-parallel beam lines designed to share a common laser can be used to greatly suppress the laser phase noise due to frequency instabilities. Noise from each beam's counter-propagating laser is common to both interferometers on that line, while noise on the shared laser is common to both pairs of interferometers. The gravitational wave signal remains present because it has different components along the two beam lines. The residual phase shift is $\sim \delta k \delta L$, where δL is the difference in length of the two beam lines. If the arm lengths are known to ~ 1 m, then the laser frequency drift may be relaxed to a fractional stability of $\sim 10^{-11}$.

3.5 Environmental Considerations in Low Earth Orbit

The sensitivity curve in Fig. 3.4 assumes atom shot noise limited phase noise. Many environmental effects, however, can contribute noise and spurious signals. We analyze these potential backgrounds for both leader-follower and inclined-great-circle orbits. The error model we consider includes the effects of rotation (Coriolis and centrifugal forces), non-uniform gravity, and magnetic fields. In addition, the analysis accounts for laser pointing errors and the effect of the finite curvature of the laser wavefronts.

3.5.1 Error Model

To calculate the phase shifts caused by these effects on the atom interferometer, we follow the standard procedure of calculating the contributions from propagation phase, separation phase, and laser phase [70, 5] for both orbital configurations shown in Fig. 3.6. In both cases, we compute the phase shift between two atom interferometers separated by baseline L. The light pulse sequence used for the calculations is the five-pulse, double-diffraction interferometer shown in Fig. 3.3. The two interferometers are assumed to have an uncontrolled initial velocity difference of $(\delta v_x, \delta v_y, \delta v_z)$ and position offset of $(\delta x, \delta y, \delta z)$, both of which may vary stochastically shot to shot. Since we are concerned only with terms that vary in time in the detection band, all constant offset phase shift terms have been subtracted. Here $T_{ij} \equiv -\partial_j g_i$ is the gravity gradient tensor and g_i is the gravitational field vector. Values for T_{ij} assume a spherical Earth.

A realistic analysis must account for the finite curvature of the laser wavefronts at the position of the atoms and temporal stability of the laser propagation axes. For a Gaussian beam propagating along the x-direction, the phase imprinted during one of the light pulses on an atom at position

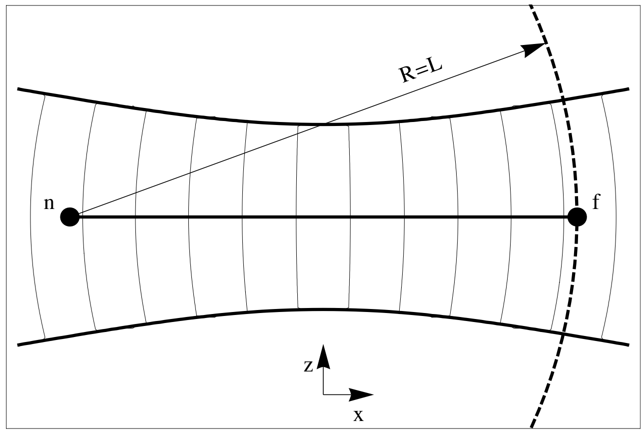

Figure 3.15: Pointing jitter insensitivity due to wavefront curvature of the Gaussian interferometer beams. The two satellites are depicted as points labeled "n" and "f" for near and far, respectively. The satellites are separated by a baseline L, and the atom interferometers are assumed to operate in close proximity to each of the satellites. The interferometer beam waist is located at the midpoint between the two satellites and the beam has a Rayleigh range $x_R = L/2$ so that the radius of curvature of the wavefront at the far satellite is $R = L$. As a result, the laser phase at the location of the far satellite is insensitive to pointing jitter caused by uncontrolled rotation of the beam about the near satellite.

(x, y, z) is given by

$$\Phi(x, y, z) = kx + \frac{k}{2R_c(x)}(y^2 + z^2) - \arctan(x/x_R) \tag{3.9}$$

where $R_c(x) \equiv x + x_R^2/x$ is the radius of curvature of the wavefronts and x_R is the Rayleigh range of the beam. Notice that for an atom at position $x = x_R$, the radius of curvature is $R_c(x_R) = 2x_R$.

To reduce sensitivity to pointing jitter (see below), the curvature of the wavefronts may be chosen so that the phase of the beam at the position of the atom is minimally sensitive to beam rotation. Therefore, a beam with $x_R = L/2$ and a beam waist placed at the midpoint between the two atom interferometers separated by baseline L will have a radius of curvature equal to the baseline distance at the position of the atoms (this is the same configuration as a confocal cavity). This configuration results in a substantial suppression of laser pointing sensitivity because the wavefront curvature at the position of the atom is locally the same as the curvature of a spherical wave originating from the point of rotation of the beam (see Fig. 3.15).

In our model, we assume the laser configuration shown in Fig. 3.15 in which the beam waist is placed at the midpoint between the satellites and the Rayleigh range is $x_R = L/2$. We consider the phase shift resulting from a two-photon transition (Raman or Bragg), and we assume that each of the two satellites serves as the point of origin for one of the two (counter-propagating) frequencies involved in the transition. For simplicity, we further assume that only one of the satellites experiences pointing error, so that only one of the counter-propagating Gaussian beams will have a variable propagation direction, parameterized by the angle $\theta(t)$. For the case that the near satellite is rotating (point "n" in Fig. 3.15), the phase of the two-photon phase shift at the location of the far satellite (point "f" in Fig. 3.15) is approximately given by

$$\Delta\Phi_{\mathrm{f}}(x_{\mathrm{f}}, z_{\mathrm{f}}) = 2kx_{\mathrm{f}} + kz_{\mathrm{f}}\left(\frac{x_{\mathrm{f}}}{L}\right)\theta(t) + kL \cdot O\Big[(x_{\mathrm{f}}/L)^2, (z_{\mathrm{f}}/L)^2, (kL)^{-1}\Big] \tag{3.10}$$

where x_{f} is the longitudinal position offset of the far atom away from the ideal Rayleigh range position, z_{f} is the transverse offset, and we take $y = 0$. Notice that at the exact position of the Rayleigh range ($x_{\mathrm{f}} = 0$) the phase is (to first-order) independent of $\theta(t)$. Likewise, the two-photon phase shift at the location of the near satellite is approximately

$$\Delta\Phi_{\mathrm{n}}(x_{\mathrm{n}}, z_{\mathrm{n}}) = 2kx_{\mathrm{n}} + kz_{\mathrm{n}}\left(1 + \frac{x_{\mathrm{n}}}{L}\right)\theta(t) + kL \cdot O\Big[(x_{\mathrm{n}}/L)^2, (z_{\mathrm{n}}/L)^2, (kL)^{-1}\Big] \tag{3.11}$$

where in this case x_{n} is the longitudinal position of the near atom with respect to the point of rotation of the beam and z_{n} is the transverse offset. Note that the phase $\Delta\Phi_{\mathrm{n}}$ at the position of the near atom is still sensitive to $\theta(t)$ since the wavefront curvature cannot be matched at that distance from the rotation point.

We parameterize laser pointing errors with two terms: a low frequency uniform pointing drift $\delta\Omega$, which persists over the duration of the interferometer sequence, and a high frequency pulse-to-pulse angular jitter $\delta\theta(t)$; *e.g.* $\theta(t) = \delta\Omega t + \delta\theta(t)$. The rms deviation for $\delta\theta(t)$ is taken to be $\overline{\delta\theta}$.

In addition to the phase, the curved wavefronts also affect the atom trajectories during the interferometer. As a result of wavefront curvature, the momentum transferred to the atom at each interaction point is position dependent. Since the wavevector $\mathbf{k}$ is the gradient of the local phase at the position of the atom-light interaction, we have $\mathbf{k}(x, y, z) = \nabla\Phi(x, y, z)$, where $\Phi(x, y, z)$ is given by Eq. 3.9. The finite beam curvature causes the effective wavevector $\mathbf{k}_{\mathrm{eff}}$ to point slightly away from the interferometer axis at off-axis points. This changes the direction of the momentum kick imparted to the atom, modifying its trajectory. A modified atom trajectory in turn leads to a correction to the phase shift. We account for this effect by calculating the actual momentum transfer $\hbar\mathbf{k}_{\mathrm{eff}}$ at each interaction point based on the local phase gradient at the position of the atom. As far as we know, this is the first time that the effect of wavefront curvature on the atom trajectory has been included in an AI phase shift analysis.

We perform the phase shift calculation outlined above using the rotated Gaussian beam phase shifts of Eqs. (3.10) and (3.11). The results of this analysis for the LF orbital configuration are shown in Table 3.4. The LF calculation uses the coordinate system shown in Fig. 3.1, where the x-axis is the chord connecting the two satellites. Table 3.5 shows the results of the same analysis for the IGC configuration; here the connecting chord is along the y-axis as shown in Fig. 3.6. As mentioned before, we only show the differential phase terms that can cause time-dependent background signals and so we have subtracted the unimportant constant phase.

	Differential phase shift	Size (rad)	Constraint
1	$\frac{1485 k_{\mathrm{eff}}^3 \hbar^2}{4 L m^2} T^6 T_{\mathrm{xx}} \Omega_{\mathrm{or}} \delta\Omega$	$(180\ \mathrm{s})\delta\Omega$	$\delta\Omega < 0.57\ \mu\mathrm{rad/s}$
2	$\frac{1485 k_{\mathrm{eff}}^3 \hbar^2}{2 L m^2} T^6 \Omega_{\mathrm{or}}^3 \epsilon_{\mathrm{zz}} \delta\Omega$	$(350\ \mathrm{s})\epsilon_{\mathrm{zz}}\delta\Omega$	$\epsilon_{\mathrm{zz}} < 0.50$
3	$\frac{15}{2} k_{\mathrm{eff}} T^4 R \Omega_{\mathrm{or}}^2 \left(15T\left(T_{\mathrm{zz}} + 3\Omega_{\mathrm{or}}^2\right) + 8\Phi\Omega_{\mathrm{or}}\right) \epsilon_g \delta\Omega$	$\left(3 \times 10^9\ \mathrm{s}\right) \epsilon_g \delta\Omega$	$\epsilon_g < 5.8 \times 10^{-8}$
4	$30 k_{\mathrm{eff}} T^4 \Omega_{\mathrm{or}}^4 \epsilon_{\mathrm{xx}} (\delta x_{\mathrm{n}} - \delta x_{\mathrm{f}})$	$(22\ \mathrm{m}^{-1})\epsilon_{\mathrm{xx}}(\delta x_{\mathrm{n}} - \delta x_{\mathrm{f}})$	$(\delta x_{\mathrm{n}} - \delta x_{\mathrm{f}})\epsilon_{\mathrm{xx}} < 4.5\ \mu\mathrm{m}$
5	$15 k_{\mathrm{eff}} T^4 T_{\mathrm{xx}} \Omega_{\mathrm{or}} \left(\frac{k_{\mathrm{eff}}\hbar}{Lm} + 9T\Omega_{\mathrm{or}}^2\right)(\delta z_{\mathrm{f}} - \delta z_{\mathrm{n}})$	$(0.84\ \mathrm{m}^{-1})(\delta z_{\mathrm{f}} - \delta z_{\mathrm{n}})$	$(\delta z_{\mathrm{f}} - \delta z_{\mathrm{n}}) < 120\ \mu\mathrm{m}$
6	$30 k_{\mathrm{eff}} T^4 \Omega_{\mathrm{or}}^3 \left(\frac{k_{\mathrm{eff}}\hbar}{Lm} + 9T\Omega_{\mathrm{or}}^2\right) \epsilon_{\mathrm{zz}}(\delta z_{\mathrm{f}} - \delta z_{\mathrm{n}})$	$(1.7\ \mathrm{m}^{-1})\epsilon_{\mathrm{zz}}(\delta z_{\mathrm{f}} - \delta z_{\mathrm{n}})$	$\epsilon_{\mathrm{zz}} < 0.49$
7	$\frac{45}{2} k_{\mathrm{eff}} T^5 \left(T_{\mathrm{xx}}^2 + 6T_{\mathrm{xx}}\Omega_{\mathrm{or}}^2 + 4T_{\mathrm{zz}}\Omega_{\mathrm{or}}^2 + 5\Omega_{\mathrm{or}}^4\right)\Delta v_x$	$(270\ \mathrm{s/m})\Delta v_x$	$\Delta v_x < 370\ \mathrm{nm/s}$
8	$3 k_{\mathrm{eff}} T^4 \Omega_{\mathrm{or}} \left(\frac{9 k_{\mathrm{eff}}^2 \hbar^2}{L^2 m^2} - 5T_{\mathrm{xx}}\right)\Delta v_z$	$\left(9.6 \times 10^3\ \mathrm{s/m}\right)\Delta v_z$	$\Delta v_z < 10\ \mathrm{nm/s}$
9	$30 k_{\mathrm{eff}} T^4 \epsilon_{\mathrm{zz}} \Omega_{\mathrm{or}}^3 \Delta v_z$	$\left(1.9 \times 10^4\ \mathrm{s/m}\right)\epsilon_{\mathrm{zz}}\Delta v_z$	$\epsilon_{\mathrm{zz}} < 0.52$
10	$60 \frac{\hbar k_{\mathrm{eff}}^2}{L^2 m} T^4 T_{\mathrm{yy}} \delta v_{\mathrm{yn}} \delta y_{\mathrm{n}}$	$\left(4.3 \times 10^{-2}\ \mathrm{s/m^2}\right)\delta v_{\mathrm{yn}}\delta y_{\mathrm{n}}$	$\delta v_{\mathrm{yn}}\delta y_{\mathrm{n}} < 23\ \mathrm{cm^2/s}$
11	$36 k_{\mathrm{eff}}^3 \frac{\hbar^2}{L m^2} \Omega_{\mathrm{or}} T^3 (7 + 8\cos(\omega T)) \sin^4\left(\frac{\omega T}{2}\right) \overline{\delta\theta}$	$\left(3.9 \times 10^5\right)\overline{\delta\theta}$	$\overline{\delta\theta} < 0.26\ \mathrm{nrad}$
12	$4\, k_{\mathrm{eff}} \delta z_n (7 + 8\cos(\omega T)) \sin^4\left(\frac{\omega T}{2}\right) \overline{\delta\theta}$	$\left(1.3 \times 10^{10}\ \mathrm{m}^{-1}\right)\delta z_n \overline{\delta\theta}$	$\overline{\delta\theta} < 0.77\ \mathrm{nrad}$
13	$\frac{27\sqrt{2}}{4} k_{\mathrm{eff}} x_n \frac{L}{R} \Omega_{\mathrm{or}}^2 T^2 \chi(\omega T) \overline{\delta\theta}$	$\left(1.1 \times 10^4\right) x_n \overline{\delta\theta}$	$\overline{\delta\theta} < 0.91\ \mathrm{nrad}$

Table 3.4: Differential phase shift error budget for AGIS-LEO in a leader-follower orbit. The results are written in terms of the LF coordinate system shown in Fig. 3.6. The shot-to-shot constraints on the noise amplitudes shown in the far right column are determined by the $10^{-4}\ \mathrm{rad}/\sqrt{\mathrm{Hz}}$ phase noise requirement with an assumed 1 Hz repetition rate. The labels "n" and "f" refer to the near and far atom clouds, respectively, using the convention described in Fig. 3.15. Here ω is the frequency of the pointing jitter $\delta\theta(t)$. In term 13, $x_n = 10$ m is the longitudinal distance of the near atom cloud from the center of rotation of the near satellite, and the frequency response is given by $\chi(\omega T) \equiv \sqrt{155 + 100\cos(\omega T) - 39\cos(2\,\omega T) - 16\cos(3\,\omega T)}\, \sin^2\left(\frac{\omega T}{2}\right)$. The longitudinal and transverse velocity differences are given by $\Delta v_x \equiv (\delta v_{\mathrm{xf}} - \delta v_{\mathrm{xn}})$ and $\Delta v_z \equiv (\delta v_{\mathrm{zf}} - \delta v_{\mathrm{zn}})$. In term 12 we take $\delta z_n = 10\ \mu$m.

	Differential phase shift	Size (rad)	Constraint
1	$\frac{15k_{\mathrm{eff}}^2\hbar}{Lm}T^4T_{\mathrm{yz}}\Omega_{\mathrm{or}}(\delta x_{\mathrm{f}}+\delta x_{\mathrm{n}})$	$\left(2.3\times10^{-3}\ \mathrm{m}^{-1}\right)(\delta x_{\mathrm{f}}+\delta x_{\mathrm{n}})$	$(\delta x_{\mathrm{f}}+\delta x_{\mathrm{n}})<43$ mm
2	$\frac{15}{2}k_{\mathrm{eff}}T^4T_{\mathrm{yy}}^2(\delta y_{\mathrm{n}}-\delta y_{\mathrm{f}})$	$\left(5.6\ \mathrm{m}^{-1}\right)(\delta y_{\mathrm{n}}-\delta y_{\mathrm{f}})$	$(\delta y_{\mathrm{n}}-\delta y_{\mathrm{f}})<18\ \mu$m
3	$\frac{45}{2}k_{\mathrm{eff}}T^4T_{\mathrm{zz}}\Phi\Omega_{\mathrm{or}}^2(\delta z_{\mathrm{f}}+\delta z_{\mathrm{n}})$	$\left(6.9\times10^{-2}\ \mathrm{m}^{-1}\right)(\delta z_{\mathrm{f}}+\delta z_{\mathrm{n}})$	$(\delta z_{\mathrm{f}}+\delta z_{\mathrm{n}})<1.5$ mm
4	$15k_{\mathrm{eff}}T^4\Omega_{\mathrm{or}}\left(T_{\mathrm{yz}}+\Phi\left(T_{\mathrm{yy}}+T_{\mathrm{xx}}+2\Omega_{\mathrm{or}}^2\right)\right)(\delta v_{\mathrm{xf}}+\delta v_{\mathrm{xn}})$	$(140\ \mathrm{s/m})(\delta v_{\mathrm{xf}}+\delta v_{\mathrm{xn}})$	$(\delta v_{\mathrm{xf}}+\delta v_{\mathrm{xn}})<0.72\ \mu$m/s
5	$\frac{45}{2}k_{\mathrm{eff}}T^5T_{\mathrm{yy}}^2(\delta v_{\mathrm{yf}}-\delta v_{\mathrm{yn}})$	$(67\ \mathrm{s/m})(\delta v_{\mathrm{yf}}-\delta v_{\mathrm{yn}})$	$(\delta v_{\mathrm{yf}}-\delta v_{\mathrm{yn}})<1.5\ \mu$m/s
6	$\frac{135}{2}k_{\mathrm{eff}}T^5T_{\mathrm{yz}}\Omega_{\mathrm{or}}^2(\delta v_{\mathrm{zf}}+\delta v_{\mathrm{zn}})$	$(1.2\ \mathrm{s/m})(\delta v_{\mathrm{zf}}+\delta v_{\mathrm{zn}})$	$(\delta v_{\mathrm{zf}}+\delta v_{\mathrm{zn}})<81\ \mu$m/s
7	$30k_{\mathrm{eff}}T^4\left(T_{\mathrm{zz}}+3\Omega_{\mathrm{or}}^2\right)\delta\Omega\delta v_{\mathrm{zn}}$	$\left(1.7\times10^{7}\ \mathrm{s}^2/\mathrm{m}\right)\delta\Omega\delta v_{\mathrm{zn}}$	$\delta\Omega<6.0\ \mu$rad/s
8	$\frac{171k_{\mathrm{eff}}^3\hbar^2}{2L^2m^2}\Omega_{\mathrm{or}}T^4\delta\Omega\delta x_{\mathrm{n}}$	$(63\ \mathrm{s/m})\delta\Omega\delta x_{\mathrm{n}}$	$\delta\Omega<37\ \mu$rad/s
9	$30k_{\mathrm{eff}}T^4\Omega_{\mathrm{or}}^2\left(\frac{k_{\mathrm{eff}}\hbar}{Lm}+\frac{45}{4}T\Omega_{\mathrm{or}}^2\right)\delta\Omega\delta z_{\mathrm{n}}$	$\left(1.7\times10^{3}\ \mathrm{s/m}\right)\delta\Omega\delta z_{\mathrm{n}}$	$\delta\Omega<40\ \mu$rad/s
10	$\frac{27k_{\mathrm{eff}}^3\hbar^2}{4L^2m^2}T^3\delta\Omega\delta z_{\mathrm{f}}$	$\left(1.1\times10^{3}\ \mathrm{s/m}\right)\delta\Omega\delta z_{\mathrm{f}}$	$\delta\Omega<62\ \mu$rad/s
11	$\frac{225}{2}k_{\mathrm{eff}}RT^5\Omega_{\mathrm{or}}^2\left(T_{\mathrm{zz}}+3\Omega_{\mathrm{or}}^2\right)\epsilon_g\delta\Omega$	$\left(2.4\times10^{9}\ \mathrm{s}\right)\epsilon_g\delta\Omega$	$\epsilon_g<6.8\times10^{-9}$
12	$\frac{30k_{\mathrm{eff}}^2\hbar}{L^2m}T^4R\Omega_{\mathrm{or}}^3\epsilon_g(\delta x_{\mathrm{n}}-\delta x_{\mathrm{f}})$	$\left(180\ \mathrm{m}^{-1}\right)\epsilon_g(\delta x_{\mathrm{n}}-\delta x_{\mathrm{f}})$	$\epsilon_g<1.3\times10^{-5}$
13	$60k_{\mathrm{eff}}T^4\Omega_{\mathrm{or}}^3\epsilon_{\mathrm{xx}}\delta\Omega\delta x_{\mathrm{n}}$	$\left(3.8\times10^{4}\ \mathrm{s/m}\right)\epsilon_{\mathrm{xx}}\delta\Omega\delta x_{\mathrm{n}}$	$\epsilon_{\mathrm{xx}}<1.0\times10^{-2}$
14	$60k_{\mathrm{eff}}T^4\Omega_{\mathrm{or}}^2\left(\frac{k_{\mathrm{eff}}\hbar}{Lm}+\frac{15}{2}T\Omega_{\mathrm{or}}^2\right)\epsilon_{\mathrm{zz}}\delta\Omega\delta z_{\mathrm{n}}$	$\left(2.6\times10^{3}\ \mathrm{s/m}\right)\epsilon_{\mathrm{zz}}\delta\Omega\delta z_{\mathrm{n}}$	$\epsilon_{\mathrm{zz}}<4.2$
15	$9\sqrt{2}k_{\mathrm{eff}}y_n\frac{L}{R}\Omega_{\mathrm{or}}^2T^2\chi(\omega T)\overline{\delta\theta}$	$\left(1.4\times10^{4}\ \mathrm{m}^{-1}\right)y_n\overline{\delta\theta}$	$\overline{\delta\theta}<0.71$ nrad
16	$4\,k_{\mathrm{eff}}\delta z_n(7+8\cos(\omega T))\sin^4\left(\frac{\omega T}{2}\right)\overline{\delta\theta}$	$\left(1.3\times10^{10}\ \mathrm{m}^{-1}\right)\delta z_n\overline{\delta\theta}$	$\overline{\delta\theta}<0.77$ nrad

Table 3.5: Differential phase shift error budget for AGIS-LEO in an inclined-great-circle orbit. The results are written in terms of the IGC coordinate system shown in Fig. 3.6. The shot-to-shot constraints on the noise amplitudes shown in the far right column are determined by the 10^{-4} rad/$\sqrt{\mathrm{Hz}}$ phase noise requirement with an assumed 1 Hz repetition rate. The labels "n" and "f" refer to the near and far atom clouds, respectively, using the convention described in Fig. 3.15. In term 15, $y_n = 10$ m is the longitudinal distance of the near atom cloud from the center of rotation of the near satellite. In term 16 we take $\delta z_n = 10\ \mu$m.

The terms in Tables 3.4 and 3.5 are not ranked by size. Instead, each term represents the largest contribution to the susceptibility of the instrument to a particular noise source. The relative sign between the different terms in the list is suppressed in order to avoid the appearance of spurious cancelations. As an example, term 5 in Table 3.4 quantifies the instrument sensitivity to the transverse positions δz_n and δz_f of the near and far atom clouds.

The size of each term in the list is written as a product of the susceptibility and the associated noise amplitude. The size of the phase response determines a constraint on the noise amplitude based on the desired instrument phase noise requirement of 10^{-4} rad/$\sqrt{\text{Hz}}$. The constraint values shown in the far right column of Tables 3.4 and 3.5 should be interpreted as the allowed tolerance per shot (not per $\sqrt{\text{Hz}}$), with an assumed instrument repetition rate of 1 Hz.

Note that the IGC analysis does not account for the non-zero relative velocity between the two atom clouds or the required orbital phase offset between the satellites to avoid collision (see discussion in Sec. 3.2.4). Both of these effects can be addressed by rotating and chirping the interferometer laser to make it remain inertial during the interferometer sequence [70, 15].

Many of the phase shifts calculated in this analysis cancel as a result of the symmetry of a perfect circular orbit around a spherical Earth. In order to account for possible deviations from this condition, these terms appear in Tables 3.4 and 3.5 along with a set of ϵ-parameters that quantify the imperfections of the orbit. The orbital acceleration condition imperfection is quantified by $\epsilon_g \equiv (g - R\Omega_{\text{or}}^2)/R\Omega_{\text{or}}^2$, where R is the radius of the orbit. Deviations in the gravity gradients caused by the non-spherical Earth are parameterized by $\epsilon_{\text{xx}} \equiv (T_{\text{xx}} - \Omega_{\text{or}}^2)/\Omega_{\text{or}}^2$, $\epsilon_{\text{yy}} \equiv (T_{\text{yy}} - \Omega_{\text{or}}^2)/\Omega_{\text{or}}^2$, and $\epsilon_{\text{zz}} \equiv (T_{\text{xx}} + 2\Omega_{\text{or}}^2)/2\Omega_{\text{or}}^2$. Note that these gravity gradient ϵ-parameters are not exactly zero even for a perfectly spherical Earth as a result of the small tilt of the coordinate system (see Fig. 3.6).

Several of the noise constraints in Tables 3.4 and 3.5 are interdependent. For example, term 9 in Table 3.5 puts a constraint on the product of transverse position error δz_{n} and the rotation rate error $\delta\Omega$. In situations like this, the constraint reported in the table generally accounts for the other existing constraints on the same noise parameters. In the case of term 9 in Table 3.5, we use the pre-existing constraint on δz_{n} from term 3 in Table 3.5 in order to calculate the implied constraint on $\delta\Omega$. A counterexample of this is term 7 in Table 3.5, where we assume $\delta v_{\text{vn}} = 1$ μm/s in order to mitigate the requirement on $\delta\Omega$.

The largest terms due to angle jitter $\overline{\delta\theta}$ appear at the end of Tables 3.4 and 3.5 (terms 11-13 and 15-16, respectively). The angle jitter results have been simplified by enforcing the ideal circular orbit condition $g = \Omega_{\text{or}}^2 R$ as well as the spherical Earth approximation $T_{\text{xx}} = T_{\text{yy}} = -\frac{1}{2}T_{\text{zz}} = \Omega_{\text{or}}^2$. In Table 3.4, term 13 appears as the dominant response for the same analysis performed with plane waves instead of Gaussian beams; here the term is suppressed by $x_n/L \sim 3 \times 10^{-4}$ compared to the plane wave case as a result of the curved wavefront at the far satellite. Term 11 consists of equal contributions from both the near and the far satellite and is a result of the finite longitudinal extent of each atom interferometer coupled with transverse Coriolis deflections. The main difference for the inclined-great-circles orbit configuration is that the transverse Coriolis deflections are greatly suppressed. As a result, term 11 in Table 3.4 does not appear in the IGC results.

The precise physical origin and interpretation of the individual terms in Tables 3.4 and 3.5 is often obscure. To first order, the five-pulse sequence combined with the differential measurement strategy is immune to many of the readily identifiable error sources such as the Coriolis force and the gravity gradient. The remaining noise susceptibilities arise due to higher order cross-couplings between rotation, gravity gradients, laser wavefront curvature and pointing jitter. For example, the second half of term 8 in Table 3.4 results from a coupling between rotation and the longitudinal gravity gradient. Specifically, a nonzero velocity v_z in the z direction combined with the y-axis rotation bias Ω_{or} causes an x-axis Coriolis deflection given by $\Delta x \sim \Omega_{\mathrm{or}} v_z T^2$. This longitudinal deflection leads to an acceleration $a \sim T_{\mathrm{xx}} \Delta x$ due to the gravity gradient. Roughly speaking, the phase shift response to this acceleration is then $k_{\mathrm{eff}} a T^2 \sim k_{\mathrm{eff}} T_{\mathrm{xx}} \Omega_{\mathrm{or}} v_z T^4$, and term 8 is recovered in the case when the near and far atom clouds have different v_z.

Terms that arise due to the Gaussian wavefront curvature typically scale as $\sim 1/L$ since the radius of curvature of the wavefronts is $R_c(x) \sim L$ at the position of the interferometers (see Eq. 3.9). These terms are proportional to the dimensionless ratio $\rho \equiv v_r T/L$, where $v_r = \hbar k_{\mathrm{eff}}/m$ is the atom recoil velocity. This ratio appears in terms 1, 2, 5, 6, 8, 10, and 11 in Table 3.4. Since $v_r T$ is the rough longitudinal extent of the interferometer, the curvature terms depend on the ratio of the size of the interferometer to the collimation length of the laser. In the limit that $\rho \to 0$, the laser wavefronts act effectively as plane waves and the wavefront curvature error terms become negligible. For comparison, we include the results of the phase shift analysis for plane waves (ignoring the Gaussian wavefront curvature) in Appendix 3.9.1.

3.5.2 Orbital Altitude

The error model in the previous section indicates that many systematic effects are mitigated at larger orbital radii, where the rotation rate, Ω_{or}, is smaller. In general, the selection of the satellites' orbital altitude needs to account for several factors, including vacuum requirements, magnetic fields, gravity gradients, rotational effects, and the fraction of an orbital period spent in the Earth's shadow. Of these, only the last favors lower orbits, since without a satellite sunshield we rely on the Earth's umbra to protect the interferometers' atom clouds from a large $\sim 10\ \mathrm{s}^{-1}$ scattering rate with solar photons [43]. For an equatorial orbit at 1000 km, a satellite spends approximately one third of its time in the Earth's shadow. All of the other altitude-influencing factors mentioned improve with increasing height, and in the following subsections we argue that their effects are acceptable at or even below 1000 km.

3.5.3 Vacuum Requirements

Collisions with particles in the background gas cause the decoherence of atoms in the interferometer region. While this is not a source of noise and does not introduce a spurious phase shift, the removal of atoms from the cloud reduces the sensitivity of the instrument. Around 700 km above the Earth's surface, the atmosphere is dominated by H, He, and atomic O with densities of $10^6 \ \frac{\text{atoms}}{\text{cm}^3}$ at temperatures near 1000 K [116]. This corresponds to pressures of 10^{-10} Torr, typical for low-background-collision-loss atom interferometry experiments. At this pressure, the average time between collisions of an atom with the background gas is ~ 100 s [43], several times longer than the proposed interferometer time of ~ 10 s. Thus, satellites constructed with low-outgassing materials at altitudes above 700 km are sufficient to meet the vacuum requirements of AGIS-LEO.

3.5.4 Magnetic Fields

Since the atom interferometers operate outside of the AGIS-LEO satellites, the atoms are in a region of space with uncontrolled magnetic fields. The local magnetic field defines the quantization axis used for the atom-laser interaction and shifts the atomic energy levels, which can cause an unwanted phase shift in the atom interferometer. Variations in the direction of the magnetic field have been addressed in the original AGIS proposal [43], and similar strategies to define the quantization axis, such as a large permanent magnet, can be employed in AGIS-LEO.

Phase shifts due to time-dependent variation in the atomic energy levels are more concerning. In order to mitigate these effects, we perform interferometry on atoms in the $|m = 0\rangle$ clock state which is first-order insensitive to energy shifts from magnetic fields. However, alkali atoms have a second order energy shift $\Delta E = \frac{1}{2}\alpha \vec{B}^2$, and since $\vec{B}$ varies between the two AGIS-LEO interferometers, there is a differential phase shift between them. We analyzed the in-situ magnetic field for various AGIS-LEO orbits using the World Magnetic Model 2010 (WMM2010) [117]. The WMM2010 models the Earth's magnetic field using a 12th order spherical harmonic expansion with the minimum wavelength of 28.8°. This spatial resolution is much larger than the orbital arc traversed during a 24 second interferometer sequence, so we use a linear expansion of the field about the position of the atom given by $\vec{B_0} + \frac{\partial \vec{B}}{\partial \lambda}\lambda(t - t_0)$, where $\lambda(t)$ is the orbital longitude of the atom at time t and the start of the interferometer sequence is denoted by time t_0. The leading order response of a single three-pulse atom interferometer to these spatial gradients is:

$$\Delta\phi_{\mathrm{B}} \simeq \frac{\hbar k_{\mathrm{eff}}}{m}\alpha T^2 B_0 \frac{\partial B}{\partial \lambda}\frac{\partial \lambda}{\partial x} \simeq \frac{\hbar k_{\mathrm{eff}}}{m}\alpha T^2 B_0 \frac{\partial B}{\partial \lambda}\frac{1}{R_E} \tag{3.12}$$

For this long-wavelength model, variations are suppressed by the ratio of the interferometer size to the radius of the Earth, R_E. The differential phase shift between two interferometers spaced by a baseline L (assuming the LF orbital configuration) is

$$\Delta\phi_{\mathrm{B,gradiometer}} \simeq \frac{\hbar k_{\mathrm{eff}}}{m}\alpha T^2 \frac{(\partial_\lambda B)^2 L}{R_E^2} \tag{3.13}$$

This is further suppressed by the ratio of the baseline L to the radius of the Earth. For an example 850 km equatorial orbit, using WMM2010, we expect $\sim$ 200 mG magnetic fields with $\frac{\partial B}{\partial \lambda} \sim 1$ mG/(deg. longitude). This results in a differential phase shift of $\Delta\phi_{\mathrm{B,gradiometer}} \sim 10^{-8}$ rad, which is well below atom shot noise. Magnetic field dependent phase shifts should be further suppressed by a five-pulse AI sequence (see discussion in Sec. 3.5.1).

Additionally, there may exist finer local variation in the magnetic field due to the presence of the spacecraft. However, the interferometer sequence takes place a distance $d \gtrsim 3$ m from the spacecraft, and local spurious magnetic fields fall off as $\sim d^3$. Therefore the magnetic field gradient differences between the two spacecraft can likely be controlled below the required level of 10^{-8}G/m. Furthermore, pulse sequences that use larger numbers of pulses, such as the four- and five-pulse sequences discussed above, will result in additional suppression of this effect.

3.5.5 Newtonian Gravitational Fields

Fluctuations in the Newtonian gravitational field can masquerade as a gravitational wave. To avoid these spurious signals, the AI pulse sequence and the differential measurement satellite configuration are designed to minimize phase shifts caused by the Newtonian gravitation potential ϕ. As discussed in Sec. 3.5.1, AGIS-LEO benefits from many exact cancellations due to its pulse sequence symmetry and the circular orbit condition which implies $\Omega_{\mathrm{or}}^2 = g/R = T_{\mathrm{xx}} = \frac{1}{2}T_{\mathrm{zz}}$. However, the analysis in Sec. 3.5.1 assumes a perfectly spherical Earth and does not take into account the local gravitational field due to the satellites. We now analyze the impact of a large local mass with position jitter as well as the effect of the Earth's inhomogeneous gravitational field.

Satellite Position Jitter

The motion of the satellite is typically a major source of concern for any space-based precision test mass experiment, since this motion can cause large time-dependent gravity gradients that can be a significant source of error. We take advantage of the disposability of the atomic test masses and perform our precision measurement outside of the satellite in order to decrease our sensitivity to these spurious gravity gradients. Since we are searching for time-dependent gravitational wave signals, we are concerned only with background noise in our detection band; overall static phase offsets between the two atom interferometers are not important. Therefore we require the satellite jitter position $\delta x_{\mathrm{sat}}(\omega)$ to be small enough in our measurement frequency band so that spurious phase shifts are below 10^{-4} rad/$\sqrt{\mathrm{Hz}}$.

We calculate the response of a single atom interferometer to satellite position jitter using perturbation theory [118]. To first order in the jitter amplitude $\delta x_{\mathrm{sat}}(\omega)$ the perturbing Lagrangian ΔL_{sat}

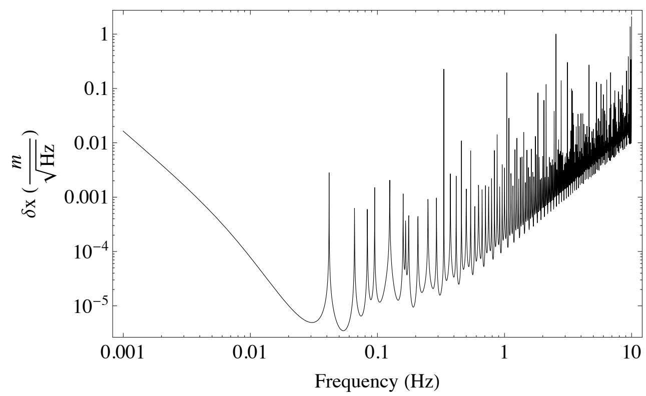

Figure 3.16: Satellite position jitter requirements as a function of frequency. This plot is for a single interferometer and assumes the five-pulse sequence of Fig. 3.3 with $200\hbar k$ atom optics, a satellite mass $M_{\text{sat}} = 1000$ kg, an initial atom cloud distance $d = 10$ m from the satellite, and a $\delta\phi = 10^{-4}$ rad/$\sqrt{\text{Hz}}$ phase noise requirement.

is

$$\Delta L_{\text{sat}}(\omega) \simeq \frac{GM_{\text{sat}}m_a}{d^2}\delta x_{\text{sat}}(\omega) \tag{3.14}$$

where M_{sat} is the mass of the satellite, m_a is the atom mass, and d is the distance between the atom and the satellite. The resulting propagation phase shift [70] from this perturbation as a function of frequency determines the interferometer's sensitivity to satellite position jitter. The required satellite control to reduce the position jitter phase shift below atom shot noise is shown in Fig. 3.16. The peak control requirement is ~ 1 μm/$\sqrt{\text{Hz}}$ at $\omega \sim 1/T$ where T is the interrogation time. This requirement is significantly less stringent than that imposed by wavefront aberrations and thus is not a driving factor in the mission design.

Non-Spherical Earth Gravity Inhomogeneities

The proximity of the Earth's non-uniform gravitational field to the instrument is also a source of non-trivial Newtonian gravitational backgrounds. The near-Earth gravity field is well known and is characterized in the Earth Gravity Model 2008 (EGM2008) by a spherical harmonic expansion [119]. EGM2008 resolves spatial wavelengths as small as 5 arcminutes, which is approximately the same size as the orbital arc subtended by AGIS-LEO during a single AI sequence. If the two interferometers are in different gravitational fields, they will record different phase shifts. These phase shifts will change over the course of the orbit, producing a time-dependent background signal.

At the position of the atom, the non-spherical Earth produces a time-varying gravitational potential ϕ which can be characterized by spherical harmonics. These spherical harmonics form a natural basis to use for computing the AGIS-LEO response to gravitational deviations from the static, spherical Earth. Expanded in this basis, the total gravitational potential Φ at the atom's position $\{r(t), \theta(t), \lambda(t)\}$ is

$$\Phi(r(t), \theta(t), \lambda(t)) = \frac{GM_E}{r(t)} + \phi(r(t), \theta(t), \lambda(t)) \tag{3.15}$$

where the deviation $\phi(r(t), \theta(t), \lambda(t))$ from a spherical Earth is

$$\begin{aligned} \phi(r(t), \theta(t), \lambda(t)) = \frac{GM_E}{r(t)} \sum_{n=2}^{n_{\max}} \left(\frac{R_E}{r(t)}\right)^n \\ \times \sum_{m=0}^{n} \Big(\bar{c}[n,m] \cos m\lambda(t) + \bar{s}[n,m] \sin m\lambda(t)\Big) \bar{P}^n_m[\cos\theta(t)] \end{aligned} \tag{3.16}$$

where M_E is the mass of the Earth, $r(t)$, $\theta(t)$, and $\lambda(t)$ are the radial distance, northern spherical polar distance (co-latitude), and longitude of the atom cloud, $\bar{c}[n,m]$ and $\bar{s}[n,m]$ are the expansion coefficients from the EGM2008 model, and $\bar{P}^n_m[x]$ are the fully normalized associated Legendre functions of the first kind [119]. The effects of gravity gradients from a perfectly spherical Earth are already accounted for in the phase shift calculation presented in Sec. 3.5.1. Here we compute the atom interferometer response to an arbitrary spherical harmonic component of ϕ using perturbation theory. We compute the response to the following perturbing Lagrangian for the AGIS-LEO leader-follower equatorial orbit

$$\Delta L[n,m,t] = GM_E\, m_a \frac{R_E^n}{r(t)^{n+1}} \Big(\bar{c}[n,m] \cos m\lambda(t) + \bar{s}[n,m] \sin m\lambda(t)\Big) \tag{3.17}$$

We then sum the response of each individual spherical harmonic to produce the total phase shift $\Delta\phi_{\text{EGM}}$.

$$\Delta\phi_{\text{EGM}} = -\frac{1}{\hbar} \sum_{n=2}^{n_{\max}} \sum_{m=0}^{n} \left(\int_{\text{upper}} \Delta L[n,m,t] dt - \int_{\text{lower}} \Delta L[n,m,t] dt \right) \bar{P}^n_m[0] \tag{3.18}$$

The Ω_{or}^{-1} ~100 minute orbital period of the AGIS-LEO constellation causes spherical harmonics of order m to induce time-dependent phase shifts with frequency $\omega_m = m(\Omega_{\text{or}} - \Omega_E)$, where Ω_E is the Earth's rotation rate. In order to find the full frequency response due to Earth gravity for a 1 year AGIS-LEO science run, we sum the phase shift response for all spherical harmonic components with effective frequency ω_m, and divide by the AGIS-LEO frequency response for a 1 year^{-1} bandwidth. This procedure is identical to that used for plotting gravitational wave sources in Figs. 3.4 and 3.5. The effective strain response is plotted with circles (blue) in Fig. 3.17. At low frequencies ($f < 30$ mHz), this Newtonian gravity background is the dominant systematic.

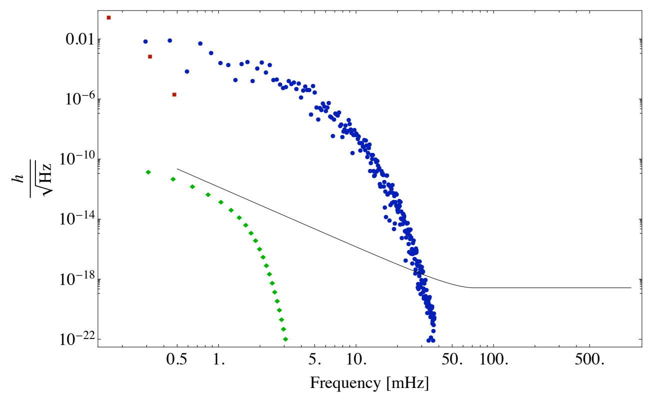

Figure 3.17: Newtonian Gravity Backgrounds for a Leader-Follower Satellite Configuration. Each point on this plot is the effective strain $h/\sqrt{\text{Hz}}$ for a Newtonian gravity background frequency component with a measurement bandwidth of 1 year^{-1}. The circles (blue) are EGM2008 components, the squares (red) are phase shifts due to lunar tidal forces, and the diamonds (green) are from free earth oscillations normalized to a $s_{0,2,0}$ peak amplitude of 1 cm. These results are for the LF orbital configuration with a baseline of $L = 30$ km. Shown for reference, the AGIS-LEO GW sensitivity curve (solid line) assumes the five-pulse sequence of Fig. 3.3 with $200\hbar k$ atom optics, a $T = 4$ s interrogation time, and a $\delta\phi = 10^{-4}$ rad$/\sqrt{\text{Hz}}$ phase sensitivity.

Lunar and Solar Tidal Forces

Tidal forces from the Moon and the Sun can also contribute to the AGIS-LEO Newtonian gravity background. We treat both bodies as point sources, and the perturbing effects arise from their gravity gradients (since AGIS-LEO is in free-fall with respect to both the Sun and the Moon it is not affected by a direct gravitational force from either body). While the five-pulse interferometer sequence is insensitive to gravity gradient phase shifts from uniform gradients (e.g., those from a perfectly spherical Earth at constant orbital altitude), it is sensitive to gradient differences between the satellites. As the satellite constellation orbits the Earth, the orientation of the sensors changes with respect to the vector pointing from AGIS-LEO to the Moon (or Sun). Therefore the orientation of the Moon gravity gradients M_{ij} changes with respect to the atom interferometers' sensitive axes. The two satellite orientations are also generally different with respect to the Moon gravity gradients. Similarly, one satellite is often farther from the Moon than the other, and this creates a difference in size between the gradients experienced by the two sensors, even though their orientations may be the same. Thus $M_{\text{ij},1}$ is not the same as $M_{\text{ij},2}$, and this induces a phase shift that is modulated by both the AGIS-LEO and Moon orbital velocities. We compute the differential phase shift response for a gravity gradient difference $\delta M_{\text{ij}} = M_{\text{ij},1} - M_{\text{ij},2}$ using perturbation theory. The 1 year science run results are shown as squares (red) in Fig. 3.17. The frequency of the response is determined by the relative orbital positions of the satellite constellation and the Moon. The plotted response is shown for the point of the Moon's closest approach to AGIS-LEO and thus is a conservative estimate for the size of the perturbation. The Sun produces a similar response, but smaller by $\frac{M_{\mathrm{Sun}}}{M_{\mathrm{Moon}}}\frac{D_{\mathrm{Moon}}^3}{D_{\mathrm{Sun}}^3} \sim 45\%$.

Free Earth Oscillations

The gravitational potential of the Earth is not constant in time and undergoes free Earth oscillations at frequencies which could potentially interfere with the search for gravitational waves [120]. Free earth oscillations are spherical harmonic eigenmodes of the Earth which ring at frequencies <10 mHz and can be excited by major earthquakes, in some cases producing >1 cm height deviations at all points on the Earth's surface [120]. Not only are the gravitational fields from these events naturally time-dependent, the orbital motion of the satellite also splits their frequency spectrum, making spectral data analysis techniques more difficult (see Fig. 3.18). However, it is important to note that free Earth oscillations are already well studied and that they can be independently measured during AGIS-LEO science runs.

The standard formalism for computing the shape and oscillation frequency for the normal modes of the Earth is well documented in the geophysics community, so we lean heavily on this preexisting analysis [121]. The Earth can oscillate in three fundamental classes of modes: radial, toroidal, and spheroidal. Radial modes are characterized by expansion and contraction of the Earth in a purely radial direction with $\mathbf{s} = f(r)\hat{r}$, where $\mathbf{s}$ is the displacement of the Earth's surface. Toroidal modes produce shearing motion between sections of the Earth's surface and have angular variation, but do not contain radial displacement. For a spherical Earth, radial and toroidal modes do not produce a Newtonian gravitational signal that is detectable by AGIS-LEO. On the other hand, spheroidal modes have angle-dependent radial deviations which create ripples in the Newtonian gravitational field that AGIS-LEO can detect during its orbit.

In order to determine the AGIS-LEO response to free Earth oscillations, we compute the perturbing potential $\phi(\mathbf{x})$ resulting from an arbitrary surface displacement, which gives

$$\phi(\mathbf{x}) = -G \int_V \frac{\rho(\mathbf{x}')\mathbf{s}(\mathbf{x}') \cdot (\mathbf{x} - \mathbf{x}')}{\|\mathbf{x} - \mathbf{x}'\|^3} dV' \tag{3.19}$$

where ρ is the density profile of the Earth and $\mathbf{s}(\mathbf{x}') = \mathbf{s}_{n,l,m}(r, \theta, \lambda, t)$ is the displacement associated with the $(n, l, m)^{\mathrm{th}}$ Earth eigenmode [121]. Therefore the gravitational potential ϕ only depends on (r, θ, λ, t) and the initial amplitude and phase of the displacement. The $(n, l, m)^{\mathrm{th}}$ component of the potential is given by

$$\phi_{n,l,m}(r, \theta, \lambda, t) = \frac{R_E^{l+1}}{r^{l+1}} e^{i\omega_{n,l}(t+t_0)} Y_{l,m}(\theta, \lambda) f(R_E) \tag{3.20}$$

where $f(R_E)$ is the integrated scale factor, t_0 is an arbitrary time offset, $\omega_{n,l}$ is the eigenfrequency, and $Y_{l,m}(\theta, \lambda)$ is shorthand notation for the fully normalized spherical harmonics that were previously defined, with $\bar{c}[m, n]$ and $\bar{s}[m, n]$ set to unity. The normal mode eigenfrequencies strongly depend on the Earth composition and are not well approximated by a simple analytic model. Therefore we use the MINEOS geophysics software package to calculate eigenfrequencies using the Preliminary Reference Earth Model (PREM) for $\rho(\mathbf{x})$ [122].

We compute the phase shift for this time-dependent ϕ using perturbation theory, as described in the previous sections. The response is shown as diamond (green) markers in Fig. 3.17. For these results, we chose $f(R_E)$ such that the peak displacement for mode $s_{0,2,0}$ is 1 cm, corresponding to the minimum displacement at any location on Earth immediately after the 2004 Sumatra-Andaman earthquake [120]. This provides a conservative upper bound to the 1 year integrated coherent signal response.

Newtonian Gravity Spectral Signatures

It is important to note that AGIS-LEO GW sources will persist over many satellite orbital periods, and can thus potentially be discriminated from signals of terrestrial origin by their spectral signature (see Fig. 3.18). For simplicity, we consider only an equatorial LF configuration here. Signal sources which are fixed to the Earth (such as those arising from the non-uniform geoid characterized by EGM2008) are shifted in frequency by the Earth rotation rate, Ω_E. However, signal sources which are not fixed to the rotating Earth, such as gravitational waves, have their signal split into a doublet spaced by twice the orbital frequency. Since free Earth oscillations are decoupled from the Earth rotation rate, they also form doublets which are spectrally indistinguishable from gravitational waves. However, free Earth oscillations are independently monitored by ground-based seismic networks and can be subtracted during signal post-processing. This picture is strictly correct only for an equatorial orbit, since a finite inclination angle would lead to doublet signatures even for Earth-fixed sources.

3.6 Short-Baseline Configurations

The AGIS instrument concept can also be adapted to configurations which have baselines significantly shorter than the 30 km version discussed above. For example, a simplified, single satellite configuration could use extendable, deployable booms to house the interferometer regions. Such a proposal would avoid the additional cost and complication of formation flying with multiple satellites. It is also worth evaluating the merits of a proof-of-principle AGIS instrument that could be space tested on the International Space Station (AGIS-ISS). While such instruments necessarily give up sensitivity due to their relatively short baselines, they could still allow for scientifically interesting gravitational wave detection, and they would facilitate technology development for a longer baseline AGIS-LEO mission.

In addition to serving as a testbed for the AGIS-LEO light-pulse atom interferometry sequences, a short-baseline instrument could be used to implement lattice-hold atom interferometry, a recently

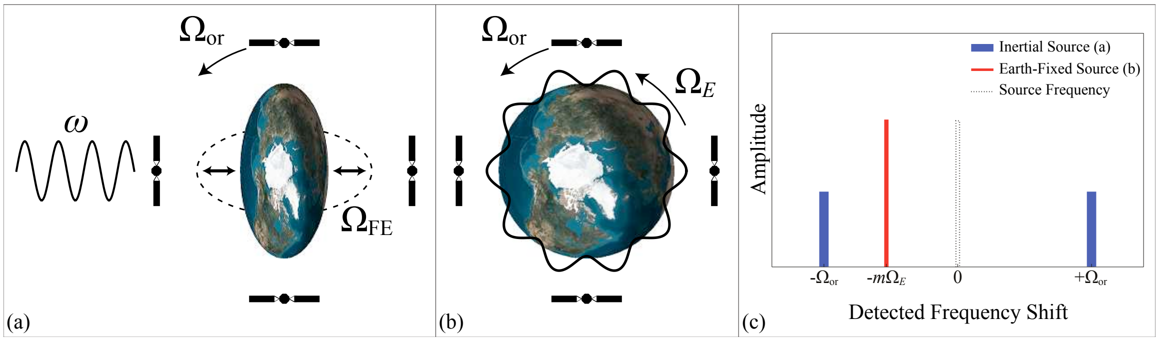

Figure 3.18: Discrimination of Newtonian gravity gradient signals by their spectral signatures for an equatorial LF configuration. (a) Inertial signal sources include gravitational waves (at frequency ω) and free-earth oscillations (Ω_{FE}, "football mode" shown). (b) Non-inertial Earth-fixed sources include the non-uniform geoid and terrain features, which rotate with the planet at a frequency of Ω_E (the spherical harmonic component pictured has $m = 10$). (c) Amplitude spectral signatures of the effects pictured in (a) and (b). For an equatorial orbit, inertial sources (blue, thicker bars) are amplitude modulated by the satellite's orbital frequency, Ω_{or}, and appear as doublets at $\pm\Omega_{\mathrm{or}}$ relative to the source frequency (gravitational wave or free-Earth oscillation). For satellites in equatorial orbits, Earth-fixed sources would appear as signals at $m\Omega_{\mathrm{or}}$ were it not for the rotation of the Earth. With the Earth rotating at Ω_E, these sources are downshifted by $m\Omega_E$ (red, thinner bar).

proposed alternative to light-pulse atom interferometry [101]. Lattice-hold atom interferometry involves the use of optical lattices to continuously control the trajectories of the two arms of an atom interferometer. Ideally, optical-lattice manipulations of the atoms require a larger two-photon Rabi frequency (which is proportional to laser intensity) than may be available in the AGIS-LEO configuration. The long AGIS-LEO baseline requires a relatively large laser beam diameter in order to mitigate diffraction, limiting the beam intensity. With the smaller beam waist allowed by a shorter baseline, the necessary intensity for lattice-hold interferometry is more easily attainable. The potential improvements in sensitivity of this promising new approach to atom interferometry suggest that it might be possible to observe gravitational waves using $10 \sim 100$ m baselines. In principle, lattice-hold AI requires the same hardware arrangement as light-pulse AI, with only software modifications needed to change the durations and frequencies of the applied laser light. Thus, a short-baseline mission could provide a space demonstration of both light-pulse atom interferometry and lattice-hold atom interferometry, with the possibility of gravitational wave detection.

In what follows, we provide a brief conceptual overview of lattice-hold interferometry as it could be applied to an ISS or other short-baseline satellite mission. Much of the analysis discussed above for a light-pulse interferometer geometry also applies to the case of a confined lattice-hold interferometer. In this section, we focus on the aspects of the analysis that must be modified for confined interferometer geometries and short baselines.

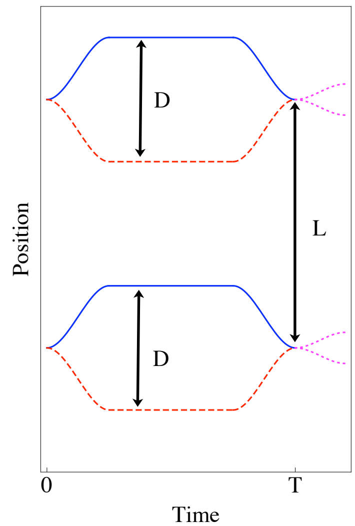

Figure 3.19: Lattice-hold interferometers in a differential configuration that could be realized on the ISS. The baseline distance between the two interferometers is L, the wavepacket separation in each interferometer is D, and the interrogation time is T. The gravitational-wave-induced phase difference between the interferometers is given by Eq. (3.21).

3.6.1 Lattice-Hold Interferometer Gravitational Wave Sensitivity

A conceptual lattice-hold atom interferometer configuration is illustrated in Fig. 3.19. Consistent with the AGIS differential measurement concept, two independent atom interferometers are separated by a baseline L along the measurement axis. Both atom clouds are continuously confined in the same one dimensional optical lattice (not shown in Fig. 3.19) such that vibration noise and laser phase noise are largely suppressed as a common mode. Using a series of accelerating and decelerating lattice potentials, each atom is coherently split and then pulled apart by a distance D. Once they are divided, a stationary optical lattice potential holds the atoms for a time T. During this hold time, the phase shift accrued by the atom is sensitive to the presence of gravitational waves.

To estimate the GW sensitivity scaling of this configuration, we only consider the effects of the gravitational wave during the hold time T and ignore the periods of acceleration and deceleration at the beginning and end of the sequence. We also point out that methods exist for calculating the classical phase shift in an accelerating optical lattice [101], and that these methods are essential for understanding the classical background effects that arise during the acceleration and deceleration phases of the interferometer sequence. However, in this section we are only interested in the phase shift caused by a GW, and the dominant contribution to this phase often arises during the hold time. Thus, we perform a relativistic phase shift calculation [34] for an atom held in a stationary optical lattice, and we ignore the effect of the GW during the times when the atoms are being accelerated.

In the relativistic calculation, we leverage the fact that in the local Lorentz frame at the initial spacetime location of the atom, the effects of curvature are negligible and we can use the classical phase shift results from [101]. While confined in the nominally stationary optical lattice, the atom's spacetime trajectory is determined by the motion of the minima of the optical standing wave. We determine the location of the standing wave minima by finding the intersection points of the counter-propagating null geodesics of the light that makes up the optical lattice. These intersection points are calculated using the general method presented in [34]. The null geodesic intersections are affected by the GW, and since we assume that the atom is rigidly confined to the optical lattice potential minima, the atom's position is also affected. The resulting gravitational-wave-induced phase difference between two lattice-hold interferometers is:

$$\Delta\phi_{\mathrm{GW}} = \frac{m}{\hbar}\omega L D h \sin(\omega T/2)\sin\theta_{\mathrm{GW}} \tag{3.21}$$

where m is the mass of the atom, h is the gravitational wave strain, ω is the gravitational wave frequency, T is the hold time, L is the baseline distance between the two interferometers, and D is the wavepacket separation in each interferometer. The phase θ_{GW} of the gravitational wave is defined as in Sec. 3.2.2. We note that the acceleration and deceleration stages of the interferometer sequence can make contributions to the phase difference that exhibit the same scaling with the relevant experimental parameters as Eq. (3.21). The size of these contributions relative to Eq. (3.21) depends on how much time is spent during the acceleration and deceleration stages as compared to the hold time.

In comparison to light-pulse atom interferometry, lattice-hold atom interferometry offers improved gravitational wave sensitivities for short baseline configurations. An interesting feature of the phase difference in Eq. (3.21) is that if we set $D \sim L$, the result scales quadratically with L, suggesting the possibility of scientifically interesting sensitivities with only modest baselines.

To illustrate the promising sensitivity of lattice-hold interferometer gravitational wave detectors over short baselines, we consider an example set of feasible parameters. For a baseline of $L =$ 10 m, a wavepacket separation of $D = 10$ m, a hold time of $T = 0.5$ s, and a phase sensitivity of 10^{-4} rad/$\sqrt{\mathrm{Hz}}$, a lattice-hold interferometer configuration using Rb atoms would have a strain sensitivity of $h \sim 10^{-16}/\sqrt{\mathrm{Hz}}$ for a gravitational wave frequency of $\omega = 2\pi \times 1$ Hz.

The GW sensitivity of the lattice-hold geometry is maximized when the wavepacket separation extends to fill the baseline ($D \sim L$). However, achieving this condition within a time that is short enough to achieve a sufficiently high repetition rate places requirements on the optical lattice intensity and detuning. Specifically, the lattice light must be far detuned in order to avoid loss from spontaneous emission, and it also must be deep enough to keep the atoms confined during acceleration. For laser intensities of ~ 100 W/cm^2, a sufficiently large two-photon Rabi frequency ($\sim 2\pi \times 60$ kHz) can be achieved to efficiently pull apart and recombine the atomic wavepackets over the ~ 10 m baseline in ~ 1 s, while the spontaneous emission rate can be kept sufficiently low (~ 0.1 s^{-1}) to avoid significant losses.

The parameters described allow for a scientifically interesting gravitational wave strain sensitivity for an instrument that could be built as a single satellite or that could fit on the ISS. Lattice-hold schemes are also promising for ground-based detectors. For instance, such a detector could be realized in the 10 m atomic fountain currently being constructed at Stanford [70].

3.6.2 Environmental Considerations near the International Space Station

As discussed previously, most of the environmental factors of Sec. 3.5 become more problematic at lower altitudes. At 350 km above the surface of the Earth (the approximate orbital altitude of the ISS), the number density of atomic oxygen is around $10^8 \ \frac{\text{atoms}}{\text{cm}^3}$, at least two orders of magnitude worse than required [116]. Thus, the AGIS-ISS instrument would need to have an enclosed vacuum tube to protect the interferometry region from background gas collisions. However, the exterior vacuum environment would allow the tube walls to be thin and non-structural. This enclosing tube would also act as a sunshield, increasing uptime by eliminating the need to operate only in the Earth's shadow. Such an enclosure could potentially even serve as a magnetic shield if one were necessitated by the lower orbit or the proximity of the ISS. Further study is required to determine the impact of gravity gradients and Coriolis effects in the lower ISS orbit and to develop a full systematic error budget for AGIS-ISS. Preliminary calculations indicate that the transverse confinement provided by red-detuned optical lattice laser beams (assuming a lattice-hold configuration) can significantly suppress the sensitivity of the instrument to Coriolis effects. AC Stark shifts are an important factor in choosing the optimal beam waist and in specifying the necessary level of intensity stabilization. These shifts can cause non-common noise in the interferometer phase that depends both on the degree of divergence of the laser beams over the interferometer region and on the size of laser intensity fluctuations. Finally, to mitigate vibrational noise from the ISS itself, one option might be to tether the instrument to the space station, rather than to attach it directly.

3.6.3 Three-Axis Boom-Based Configuration

A promising configuration for short-baseline detectors involves the use of long, extendable booms that can deploy in orbit. For example, six booms originating from a central module would allow for interferometer pairs along three orthogonal axes, with each boom housing a single interferometer (see Fig. 3.20). The interferometer laser beams could be aligned in a racetrack configuration through the use of retroreflecting elements (e.g., corner cubes) at the end of each boom and pentaprisms inside the central module to allow for 90 degree turns, enabling the control of all six interferometers with a common set of laser beams. The central module would house the atom source, optics, and control electronics. To establish the necessary vacuum level, each boom could contain an extendable (non-structural) vacuum bellows. As mentioned in the previous section, a lightweight magnetic shield could enclose each of the interferometer regions or be incorporated directly into the vacuum bellows. Compact, self-deploying booms are available from commercial manufacturers [123].

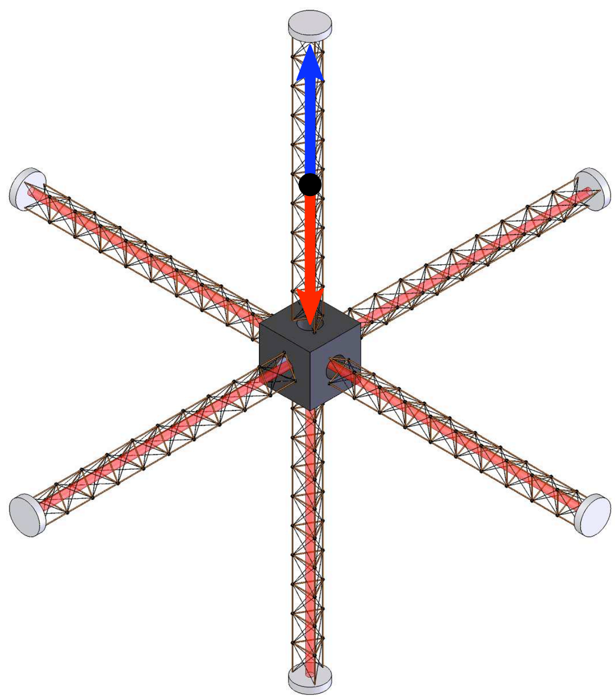

Figure 3.20: Boom configuration for a three-axis gravitational wave detector, with retroreflectors mounted at the end of each boom. The shaded red lines indicate the optical path, and the arrows on the upper boom indicate an atom interferometer fully filling its boom (specifically the trajectories of the atoms during the first half of the first loop of a sequence, as in Fig. 3.3 or Fig. 3.19). The advantages of a boom configuration are discussed in Sec. 3.6.3. Note that the figure is not to scale, that the vacuum bellows in each boom are hidden for clarity, and that only one of the six interferometer regions is indicated with arrows.

A key advantage of the boom configuration in Fig. 3.20 is that the three-axis measurement allows for the mitigation of background signals due to fluctuations in centrifugal force gradients, which arise from rotational noise. For a single-axis sensor, such fluctuations have the same signature as a gravitational wave. However, correlation measurements between three orthogonal sensors would allow for these fluctuations to be distinguished from a gravitational wave.

A boom configuration can either operate as a free flyer or be attached to the ISS. The instrument would naturally provide the flexibility to implement both light-pulse and lattice-hold interferometry due to their identical hardware requirements. In either case, it is desirable for each interferometer to fill its respective boom (i.e., the wavepacket separation should be nearly the length of a boom), so that the maximal gravitational wave sensitivity allowed by the instrument is achieved. For a free flyer configuration, it may be possible to use boom lengths of up to $L \sim 100$ m. With such a boom length and with $\hbar k_{\text{eff}} = 200\hbar k$ beam splitters, a light-pulse atom interferometer would fill the boom given interrogation times of $T \sim 100$ s.

For a free flyer, one axis could be made effectively inertial by orienting it perpendicular to the orbital plane, suppressing effects arising from rotation bias and Coriolis deflections as is done for the IGC configuration (see Sec. 3.2.4). With the use of thrusters, it may be possible to make all three axes effectively inertial.

An additional scientific motivation for the boom configuration is that the three-axis measurement can aid in the performance of certain tests of general relativity. For example, the boom configuration appears to be an ideal platform for measuring the nonzero divergence of the gravitational field in free space predicted by general relativity [34].

3.7 Secondary Objectives

Although variations in Earth's gravity gradients are a noise source for the GW detector, pulse sequences designed to measure these variations can provide important geophysical information including data regarding the flow of water on the Earth that is important for understanding weather, the climate, and oceans. Similar arguments have provided motivation for satellite missions such as the CHAllenging Mini-satellite Payload (CHAMP), the Gravity Recovery and Climate Experiment (GRACE), and the Constellation Observing System for Meteorology, Ionosphere and Climate (COSMIC). The method works: GRACE's monthly gravitational maps have improved our understanding of ocean circulation [124], polar ice loss [125], and water flow of large basins like the Amazon [126].

Furthermore, if the atom clouds are prepared in a magnetically-sensitive state, $|m \neq 0\rangle$, the gravitational wave detector becomes a detector of the local magnetic field gradient at each of the interferometers. Similar to the gravity gradiometer, the transformed detector would produce a map of the temporal variation in the Earth's magnetic gradients. Data collected by previous satellite missions–Ørsted, CHAMP, and Satellite Argentina Cientificas - C (SAC-C)–have been used to model the static and time-varying geomagnetic field [127] and mantle conductivity [128].

The modifications necessary to make these measurements do not require different hardware, only different software. With this unique flexibility of atom interferometric measurements, AGIS-LEO can provide complementary maps of temporal variations in the Earth's gravity gradients and magnetic field gradients.

3.8 Summary

Atom interferometry performed in low Earth orbit is a promising method to detect gravitational waves. The AGIS-LEO instrument has a projected strain sensitivity of $< 10^{-18}/\sqrt{\mathrm{Hz}}$ in the 50 mHz - 10 Hz frequency band, offering the possibility of detecting GW signals in an intermediate frequency band that is complementary to the coverage of laser interferometers. Previous theoretical analysis of noise sources performed for solar orbits [43] has been adapted to the conditions of low Earth orbit. We have analyzed possible environmental backgrounds from the Earth's magnetic field and inhomogeneous gravity field and find that it should be possible to keep them below atom shot noise. Laser requirements–phase noise, frequency stability, and wavefront–all appear achievable with current technology. The strong rotational effects present in low Earth orbit can be mitigated using symmetric multiple pulse sequences.

Many of the instrument and system level requirements of AGIS-LEO can be validated on the ground in terrestrial laboratories. For example, the Stanford 10 m free fall tower will allow characterization of LMT beamsplitter sequences, wavefront and laser jitter requirements, magnetic field sensitivity, laser cooling protocols, kinematic requirements for the atom sources, Coriolis and gravity gradient sensitivities, etc.

Finally, atom interferometry has the potential for even higher precision. As the technology evolves, longer baselines, higher atom flux, and larger atom beamsplitters could increase the gravitational wave sensitivity of light-pulse atom interferometers by multiple orders of magnitude beyond AGIS-LEO. Complementary to this, lattice-hold atom interferometers have the potential to provide high sensitivity for short baseline configurations. In fact, an ISS-scale test mission for AGIS using lattice-hold atom interferometry might be capable of directly detecting gravitational waves, while simultaneously flight testing the technology for a full AGIS-LEO science mission.

In short, AGIS is a sensitive and scalable instrument that can help open the door to information about our universe that is encoded in gravitational waves.

3.9 Appendix

3.9.1 Error Model using Plane Waves

Here we give the results of the AGIS-LEO error model using plane waves instead of Gaussian beams for the interferometer lasers. This is equivalent to setting $x_R \to \infty$ in Eq. 3.9. While it is necessary to include the effects of the realistic Gaussian wavefront curvature in the final analysis, the plane wave calculation serves as an important reference result since it is expected to have many features in common with the full Gaussian phase shift. Also, the plane wave analysis is much less complicated, so it is a useful diagnostic for the Gaussian calculation and provides a simplified physical picture of the results.

We follow the same calculation procedure outlined in Sec. 3.5.1. In this case, the phase of the laser at each interaction point $\mathbf{r}$ is just $\mathbf{k}_{\text{eff}} \cdot \mathbf{r}$, and the effective laser wavevector $\mathbf{k}_{\text{eff}}$ is now independent of position. As before, we allow $\mathbf{k}_{\text{eff}}$ to rotate at a fixed rate $\delta\Omega$ during the pulse sequence. The results of the LF calculation appear in Table 3.6 and the IGC results are in Table 3.7.

The LF and IGC plane wave calculations share many commonalities, and we can make the following general statements. Since the atoms are in free fall, many of the largest phase shift terms cancel in pairs. This occurs because in circular orbit around a spherical Earth, $\Omega_{\text{or}}^2 = g/R = T_{\text{xx}} = -\frac{1}{2}T_{\text{zz}}$, where R is the radius of the orbit and g is the local gravitational acceleration at altitude. As an example, see terms 1 and 2 in both Table 3.6 and Table 3.7. There are also more elaborate cancelations, such as terms 5, 8 and 9 in Table 3.6, which together all cancel as a result of $\nabla \cdot \mathbf{g} = -T_{ii} = 0$. Additionally, the orbital rotation of the laser axis, Ω_{or} can couple with variations in the atoms' launch velocity and position as well as variations in the laser axis orientation itself to produce many of the dominant backgrounds.

The T^4 scaling of the leading order phase shifts in Table 3.6 is expected given the symmetric triple-loop geometry of the five-pulse sequence (see Fig. 3.3). The conventional three-pulse, single-loop ($\frac{\pi}{2}$-π-$\frac{\pi}{2}$) gravimeter sequence is sensitive to the quadratic $\sim \frac{1}{2}gt^2$ motion of the atom, resulting in a T^2 leading-order scaling of the phase shift. Additional symmetric loops have the effect of making the interferometer sensitive instead to higher-order components of the motion $\sim t^{(\Lambda+1)}$, yielding a $T^{(\Lambda+1)}$ leading phase shift, where Λ is the number of loops. As a consequence, symmetric multiple-loop interferometers respond to higher-order derivatives of the gravitational potential, since the first non-trivial correction due to the derivative $\partial^{(\Lambda)}\phi$ appears as motion proportional to $t^{(\Lambda+1)}$. A corollary of this is that the response to lower-order derivatives $\partial^{(i)}\phi$ with $i < \Lambda$ is highly suppressed. Therefore, the use of multiple-loop interferometers can substantially reduce the interferometer's sensitivity to background effects due to inhomogeneous fields, since the perturbations caused by higher-order derivatives are significantly smaller for typical fields.

From the phase shift list for the LF orbit, we note that the largest term that does not cancel

	Differential phase shift	Size (rad)
1	$-60 k_{\mathrm{eff}} L \Omega_{\mathrm{or}}^3 \delta\Omega\, T^4$	-1.15
2	$60 k_{\mathrm{eff}} L\, T_{\mathrm{xx}} \Omega_{\mathrm{or}} \delta\Omega\, T^4$	$+1.15$
3	$888 k_{\mathrm{eff}} L \Omega_{\mathrm{or}}^5 \delta\Omega\, T^6$	$+3.67 \times 10^{-4}$
4	$444 k_{\mathrm{eff}} L\, T_{\mathrm{zz}} \Omega_{\mathrm{or}}^3 \delta\Omega\, T^6$	-3.67×10^{-4}
5	$-444 k_{\mathrm{eff}} L\, T_{\mathrm{xx}} T_{\mathrm{zz}} \Omega_{\mathrm{or}} \delta\Omega\, T^6$	$+3.67 \times 10^{-4}$
6	$30 k_{\mathrm{eff}} \Omega_{\mathrm{or}}^3 \delta v_z T^4$	$+1.92 \times 10^{-4}$
7	$15 k_{\mathrm{eff}} T_{\mathrm{zz}} \Omega_{\mathrm{or}} \delta v_z T^4$	-1.92×10^{-4}
8	$-444 k_{\mathrm{eff}} L\, T_{\mathrm{xx}} \Omega_{\mathrm{or}}^3 \delta\Omega\, T^6$	-1.84×10^{-4}
9	$-444 k_{\mathrm{eff}} L\, T_{\mathrm{xx}}^2 \Omega_{\mathrm{or}} \delta\Omega\, T^6$	-1.84×10^{-4}
10	$-225 k_{\mathrm{eff}} L\, T_{\mathrm{xz}} \Omega_{\mathrm{or}}^2 \delta\Omega\, T^5$	-1.24×10^{-4}
11	$-\frac{225}{2} k_{\mathrm{eff}} L\, T_{\mathrm{xz}} T_{\mathrm{zz}} \delta\Omega\, T^5$	$+1.24 \times 10^{-4}$
12	$15 k_{\mathrm{eff}} T_{\mathrm{xx}} \Omega_{\mathrm{or}} \delta v_z T^4$	$+9.62 \times 10^{-5}$
13	$-\frac{225}{2} k_{\mathrm{eff}} L\, T_{\mathrm{xz}} T_{\mathrm{xx}} \delta\Omega\, T^5$	-6.19×10^{-5}
14	$-\frac{45}{2} k_{\mathrm{eff}} \Omega_{\mathrm{or}}^4 \delta x T^4$	-1.67×10^{-5}
15	$15 k_{\mathrm{eff}} T_{\mathrm{xx}} \Omega_{\mathrm{or}}^2 \delta x T^4$	$+1.11 \times 10^{-5}$

Table 3.6: Differential phase shift error budget for AGIS-LEO in a leader-follower orbit using plane waves. The results are based on the triple-loop interferometer in Fig. 3.3, and the phase difference assumes the gradiometer configuration shown in Fig. 3.1 between two atom interferometers separated by $L = 30$ km in a leader-follower orbit (see Fig. 3.6). An unimportant constant phase shift has been subtracted. The jitter in the atom initial position and velocity are taken to be $\delta x = 1\ \mu$m and $\delta v_z = 10$ nm/s, respectively. The laser pointing error rotation rate is $\delta\Omega = 1$ nrad/s and the laser light is treated as an ideal plane wave. The orbital radius R and rotational rate Ω_{or} are for a 1000 km orbital altitude. Interrogation time is $T = 4$ s and the LMT beamsplitters have $\hbar k_{\mathrm{eff}} = 200\hbar k$.

	Differential phase shift	Size (rad)
1	$-\frac{225}{2}k_{\text{eff}}\frac{L^2}{R}T_{\text{zz}}\Omega_{\text{or}}^2\delta\Omega\, T^5$	$+8.25\times10^{-5}$
2	$-\frac{225}{4}k_{\text{eff}}\frac{L^2}{R}T_{\text{zz}}^2\delta\Omega\, T^5$	-8.25×10^{-5}
3	$\frac{675}{4}k_{\text{eff}}\frac{L^2}{R}\Omega_{\text{or}}^4\delta\Omega\, T^5$	$+6.19\times10^{-5}$
4	$\frac{225}{4}k_{\text{eff}}\frac{L^2}{R}T_{\text{yy}}^2\delta\Omega\, T^5$	$+2.06\times10^{-5}$
5	$\frac{15}{2}k_{\text{eff}}T_{\text{yy}}^2\delta y\, T^4$	$+5.57\times10^{-6}$
6	$45\, k_{\text{eff}}LT_{\text{yy}}\delta\Omega^2\, T^4$	$+7.47\times10^{-7}$
7	$\frac{45}{2}k_{\text{eff}}T_{\text{yy}}^2\delta v_y\, T^5$	$+6.69\times10^{-7}$
8	$15k_{\text{eff}}\frac{L}{R}\delta v_x\Omega_{\text{or}}^3T^4$	$+3.95\times10^{-7}$
9	$\frac{15}{2}k_{\text{eff}}\frac{L}{R}T_{\text{zz}}\delta v_x\Omega_{\text{or}}T^4$	-3.95×10^{-7}
10	$\frac{15}{2}k_{\text{eff}}\frac{L}{R}T_{\text{xx}}\delta v_x\Omega_{\text{or}}T^4$	$+1.98\times10^{-7}$

Table 3.7: Differential phase shift error budget for AGIS-LEO in an inclined-great-circle orbit using plane waves. The results are based on the triple-loop interferometer in Fig. 3.3, and the phase difference assumes the gradiometer configuration shown in Fig. 3.1 between two atom interferometers separated by maximum distance $L = 30$ km in an inclined-great-circles orbit (see Fig. 3.6). The required ~ 100 m offset to avoid collision is not included. An unimportant constant phase shift has been subtracted. The jitter in the atom initial position and velocity are taken to be $\delta y = 1$ μm and $\delta v_x = \delta v_y = 10$ nm/s, respectively. The separation distance L is taken to be 30 km, which is the maximum separation over the course of the orbit, and the point of largest systematic error. The laser pointing error rotation rate is $\delta\Omega = 1$ nrad/s and the laser light is treated as an ideal plane wave. The orbital radius R and rotational rate Ω_{or} are for a 1000 km orbital altitude. Interrogation time is $T = 4$ s and the LMT beamsplitters have $\hbar k_{\text{eff}} = 200\hbar k$.

(term 12 in Table 3.6) results from a coupling between gravity gradients and rotation. This term sets the limit of the size of the transverse velocity jitter δv_z for the LF configuration: for $200\hbar k$ beamsplitters, 4 s interferometer time, and $\Omega_{\rm or} \sim 10^{-3}$ rad/s, $\delta v_z \lesssim 10^{-8}$ m/s. This shot-to-shot velocity variation is equivalent to $\lesssim 100$ pK in atom cloud temperature for 10^8 atoms. Term 13 constrains the laser pointing error rotation rate to $\delta\Omega \sim 1$ nrad/s. Term 14 limits the longitudinal position variation of the atoms to $\delta x \sim 10$ μm. Term 15 is reduced below the shot noise limit by this latter constraint. These requirements are not strict, but imply a trade-off with several other parameters. If necessary, a small reduction in interferometer time at the expense of bandwidth, for example, could significantly relax these constraints, as both terms scale with multiple orders of the time.

The IGC orbit offers many systematic benefits over the LF orbit. Note that the largest non-canceling term in the IGC list (term 3 in Table 3.7) is the leading term from Table 3.6 but suppressed by an additional power of $\Omega_{\rm or}T$. This suppression arises because the IGC interferometer axis is nearly parallel with the orbital rotation axis. Therefore the transverse trajectory deviations due to the Coriolis force are suppressed by $\frac{L}{R}$ as compared to the LF orbit, which leads to a phase shift suppression $\propto \Omega_{\rm or}T$. For the same stochastic atom position jitter, the contributions to the IGC phase shifts are smaller than those of the LF. The largest phase shift due to atom jitter for the IGC is term 5 in Table 3.7, which remains below atom shot noise for position fluctuations of $\delta y \sim 20$ μm instead of ~ 10 μm. Similarly, velocity variations of ~ 1 μm/s instead of ~ 10 nm/s could be tolerated. However, if the tight atom kinematic constraints can be achieved and $\delta\Omega$ effects can be controlled, then the interrogation time could be doubled to bring these systematics to the same size as atom shot noise. This would lower the instrument corner frequency by a factor of two, potentially allowing for the detection of lower frequency GWs.

Chapter 4

ATEP

This chapter describes the Atomic Test of the Equivalence Principle (ATEP) effort at Stanford to build an 8.8 m baseline atom interferometer to test the weak equivalence principle (WEP). This apparatus will use atom interferometry to simultaneously measure the acceleration of samples of ^{85}Rb and ^{87}Rb falling in the Earth's gravitational field.

4.1 Experiment Overview

The ATEP apparatus consists of an 8.8 m vacuum chamber housed in the basement of the Varian Physics Building at Stanford. A sample of $\sim 10^6$ ^{85}Rb and ^{87}Rb atoms is cooled to $< 1\mu$K at the bottom of the vacuum chamber. In order to cool and trap Rb atoms at sub μK temperatures, we standard evaporative cooling techniques [129]. First we create a collimated beam of Rb atoms using a 2D-MOT created in a relatively high vacuum vapor cell (see Sec. 4.2.1). This beam of atoms is directed through a differential pump into an ultra-high vacuum (UHV) evaporative cooling chamber (see Sec. 4.2.2). There the atoms are evaporated in an Optically Plugged Trap (OPT) to $< 1\mu$K as detailed in Sec. 4.6.

The cooled sample is then launched into the interferometry region (See Sec. 4.2.3) via an optical lattice acceleration to an initial velocity $v_i = 13$ m/s. The atoms then follow an approximately Newtonian trajectory for ~ 2.3 s, reaching a height of 8.8 m. During their approximately inertial flight, a standard 3-pulse $\pi - \frac{\pi}{2} - \pi$ interferometer sequence is applied to measure the acceleration of each isotope (See Sec. 4.4.3 and 2.6.1). The atoms are then imaged using a CCD camera (See Sec. 4.4.2) at the bottom of the fountain.

Provided the above technical challenges are met, the experiment will only reach its design sensitivity of $\eta \simeq 10^{-15}$, if the systematic accelerations on the two isotopes due to gravity gradients, coriolis forces, and magnetic field gradients are successfully reduced below $10^{-15}g$.

Systematic Control Overview

Control over species-dependent accelerations requires a specific subsystems of the experiment to address each potential source of error. The error model is presented in great detail in section 2.6, and only the dominant sources are highlighted here. Fundamentally, any non-WEP violating species dependent acceleration has to arise from the two isotopes following different geodesics and exploring different potentials. Since the isotopes have different masses and different electronic structure, they can easily follow different trajectories (e.g. two differing mass objects suspended from springs with the same spring constant k will have different displacements in the same gravitational field). If we could enforce identical trajectories on the two isotopes during their inertial flight, then we would not need to reduce inhomogeneities in the potential. However, there will always be a small difference in the isotope trajectories due to experimental imperfections, and these lead to spurious accelerations which have three dominant sources. First, if the isotopes explore different locations in the gravitational field, then they will experience different accelerations due to gravity gradients, since the gravitation field of the earth is not constant. These are planned to be restricted via tight kinematic control over the isotopes as described in Section 2.6.1. Secondly, because the lasers are tied to the earth, interacting with the atoms via laser pulses to perform the interferometry will impart an effective differential coriolis force on the atoms due to their different transverse velocities. We combat this with kinematic constraint (reducing $\delta \mathbf{v}$) and by rotating the interferometer lasers counter to the Earth's rotation (reducing $\mathbf{\Omega}$). The theoretical basis of this technique is described in Section 2.6.1, while the experimental implementation is presented below in section 4.3.2. Lastly, since ^{85}Rb and ^{87}Rb have different electronic structures, they have different second order Zeeman coefficients, which lead to different energy shifts on the $|m_f = 0>$ state in a magnetic field. Therefore magnetic field gradients impart different forces on the two isotopes. We reduce this effect by reducing the size of the magnetic field gradient in the interferometer region as described in section 4.3.1.

4.2 Vacuum System

Since ATEP is just a large atom interferometer, we need to create a standard ultra-high vacuum environment. The vacuum system for the ATEP experiment consists of a vapor cell 2D-MOT, a condensation chamber, the large 8.8 m interferometer region, the RCS chamber, and the interferometer beam injection chamber. These systems are all connected together, using differential pumps where appropriate. It is important to note that the majority of the apparatus, by volume, is the interferometer region, which is simply a clear aperture tube, connecting the upper window of the beam injection chamber down to the retro-reflection mirror of the RCS.

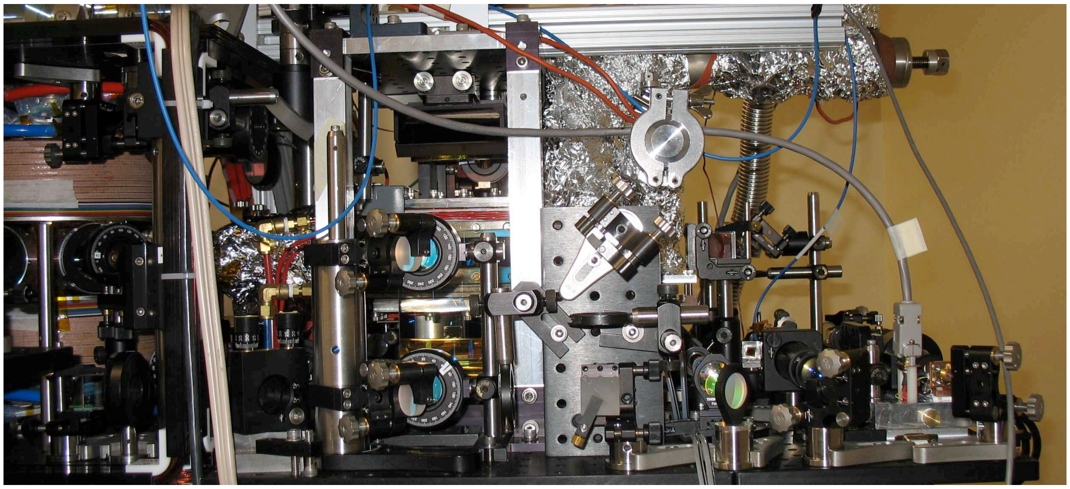

Figure 4.1: The vapor cell surrounded by optics. In this image, the light delivery optics and the 2D MOT vapor cell are shown. The evaporative cooling chamber is visible on the left, and the light amplification stage (see sec. 4.4.4) is on the right.

4.2.1 Vapor Cell

For loading our initial sample of *Rb* atoms we use a 2D-MOT in a high vacuum ($\sim 10^{-7}$ Torr) Rb-vapor cell. The chamber is formed from a 55 mm by 55 mm by 200 mm long 304 stainless block. The central region was machined out to form a cell with optical access of 125 mm by 25 mm on all four sides. The four rectangular windows which provide access for the 2D MOT beams are attached via Indalloy seals bonded to the glass and the chamber via Epotex 353. The Indalloy seals are then compressed after bonding via 14 bolt flanges which are torqued to $3-5$ N-m. These seals achieve leak rates $< 10^{-10}$ Torr-L/s. When viewed from the side, these large windows are clearly visible, as shown in Fig. 4.1.

The 2D MOT vapor cell is connected to the ultra-high vacuum (UHV) interferometer region via a 10^{-3} differential pump, which is composed of a 1 cm long by 2 mm diameter stainless steel tube and an 8 cm long by 3 mm diameter POCO graphite tube. The POCO also serves as a getter pump for Rb. A single 45 degree protected gold mirror was bonded to the end of the chamber with a 3 mm hole bored down its center. The first stainless steel section of the differential pump is bonded into this mirror. The atoms exit the vapor cell through this aperture.

The stainless steel chamber has optical access down the axis of the 2D-MOT via a 1 inch aperture window connected to the end of the chamber via an Indalloy seal. This allows for imaging the axis of the 2D MOT and imaging of the 3D MOT through the differential pump aperture.

Some trouble occurred during the initial implementation of this chamber when the rectangular seals were not bonded on to the chamber with Epotek 353-ND epoxy. Without the epoxy bond, the chamber flexed under the higher required seal torque, causing leaks.

4.2.2 Cooling Chamber

The evaporative cooling chamber consists of an octagon with internal water cooling channels that is vacuum brazed to two 6 inch long by 4 inch OD stainless steel vacuum pipes. Each of these pipes as a rotatable 6 inch ConFlat (CF) flanged welded to it. Coil supports are welded to each of the tubes to provide mechanical attachment points for the wire forming the quadrupole coils. The octagon has 8 faces machined flat with doubled bolt patterns for a 2.125 CF flange. Six windows which are AR coated on both sides at 780 nm with 0.96 inch clear aperture are attached via Indalloy seals. On one of the vacant sealing surfaces, the differential pump connecting to the 2D MOT vapor cell is attached using a 0.06 inch thick Indalloy seal. The other seal surface connects to a 2.75 inch CF reducer which ultimately connects to a tee containing a 2 inch clear aperture window AR coated at 780 nm and 532 nm as well as an all-metal tee valve connecting to a turbo-mechanical pump for initial pumpdown. This window is used to deliver light for the blue detuned plug of the OPT (see Sec. 4.6).

The 3D MOT and evaporative cooling apparatus share the same chamber and the same quadrupole coils. These coils are each composed of 12 layers with 13 turns per layer of 0.25 inch by 0.25 inch square cross section copper wire with a 0.125 inch diameter hole extruded down the center. The coils are plumbed in parallel, with every 2 layers forming a separate sub-coil for cooling. All 12 sub-coils are connected in parallel so as to reduce the pressure drop across the coil. We typically flow 2 gallons/minute (GPM) of water through the coil and another 2 GPM through the main chamber water cooling manifold. This water cooling keeps the total water temperature increase $<$ 5 degrees C while running at 300 Amps with a $< 10\%$ duty cycle.

The quadrupole coils provide the 300 Gauss/cm magnetic field gradient required for the OPT. Each of the coils has a resistance of $\sim 1/12$ Ohm, which combined with the voltage drop of the switching circuitry allows 300 Amps to be delivered with a $\sim$ 55 V drop at room temperature. Based on simulation of these coils, 1 Amp corresponds to 1 Gauss/cm gradient at the magnetic field zero.

4.2.3 Atomic Fountain

The interferometer region consists of a 5 inch OD, 4 inch ID, 28.8 foot long aluminum vacuum chamber with 6.75 inch CF flanges attached to either end. This chamber is wrapped in a 12 segment, single layer solenoid designed to provide a 1 mG bias field to define the quantization axis for the atom optics. Around this solenoid, 20 type-E (non-magnetic) thermocouples are installed with uniform spacings along the tube. The chamber is wrapped with four resistive heat tapes which are used during bakeout. Finally the entire structure is wrapped in 0.5 inch thick ceramic insulation. All material used during construction is non-ferrous and bake-able to 150 degrees-C.

The pressure in the interferometer region is measured via two ion gauges located in crosses just below and above the aluminum vacuum chamber. These crosses also house titanium sublimation pumps (TiSub pumps) with custom baffles on translatable bellows. These TiSub pumps can be moved to have line-of-sight access to the inside of the interferometer region, sublimate titanium, and then be moved clear of the interferometer beam axis.

Two all-metal VAT gate valves are used to isolate various parts of the system. One is located just above the top of the aluminum vacuum chamber. The other is located directly above the RCS chamber. The first all-metal gate valve serves to shield the upper window from titanium being sublimated by the lower titanium sublimation pump. The second gate valve is used to shield the waveplate and retro-reflection mirror from titanium sublimated by both sublimation pumps.

The main interferometer region is pumped by two 300 L/s Gamma Vacuum ion pumps. One pump is located at the very top of the apparatus, right above the injection chamber. The second pump is located immediately below the cooling chamber and directly above the RCS chamber. Additionally, the RCS chamber is pumped via a 35 L/s ion pump and the 2D MOT vapor cell is pumped via a 10 L/s ion pump, both from Gamma Vacuum.

4.3 Systematic Control

4.3.1 Magnetic Shield

The magnetic shield for the experiment consists of three concentric mu-metal cylinders (0.05 inch thick by 7.5, 11.5, and 15 inches in diameter respectively which are 28.8 feet long. The size of the shielded region defines the interferometer region. As described in section 2.6, we need 10 μG/m magnetic field gradients in order to meet our systematic error requirements. Given an estimated background of 100 mG with order 1 variation, this requires a shielding ratio of 10^4, which is generally achievable with mu-metal shields.

Before installing any shielding the background magnetic field was measured using an Applied Physics 540 flux-gate magnetometer, and $|\mathbf{B}|$ is shown in black (dotted) in Fig. 4.2. This result matched our expectations and we proceeded to install a segmented version of the shield described above. This resulted in an undesirable magnetic field profile as shown in blue (dashed) in Fig. 4.2. The regular spikes are caused by the gaps between the segments of the shield, where the permeability changes from $\mu \sim 10^5$ to $\mu = 1$. This results in trapped magnetic flux leaking into the shielded region and causing a large spike in $|\mathbf{B}|$. This behavior was verified in finite element simulations when they included the physical gaps in the magnetic shield.

In order to remove these undesirable spikes, the only solution was to weld all 14 segments into a single continuous shield and thus eliminate the gaps between the segments. However, if the shield is not annealed after welding, then the permeability at the heat-affected weld joint will be lower than that of the annealed metal around it and this creates an effective gap, albeit not as bad as air

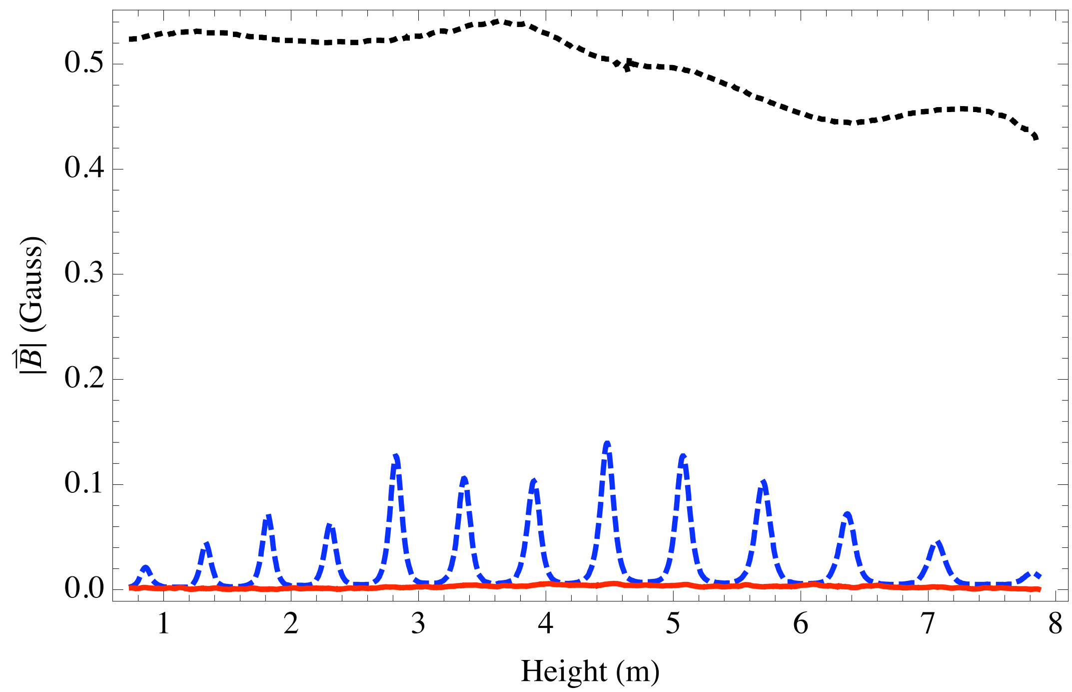

Figure 4.2: Magnetic Shield Performance in the Interferometer Region. Before Shielding (Black, Dotted), With Segmented Shield (Blue, Dashed), With Welded and Annealed Shield (Red, Solid).

with $\mu = 1$. Therefore the shield must be annealed after welding to create a uniform permeability throughout the shield. We welded the shield segments together into continuous tubes and annealed them in a hydrogen atmosphere at 2100 deg. F using a 36' x 6' x 6' furnace at Solar Atmospheres in Pennsylvannia. The resulting magnetic field profile is shown in red (solid) in Fig. 4.2. It appears to be within a factor of 10 of the require specification and stable from measurement to measurement. Therefore it may be possible to measure its effect and post-subtract it from the final result.

4.3.2 Rotation Compensation System

The rotation compensation system uses three piezo-actuators to move the retro-reflection mirror so that the interferometer beam is no longer rotating with the Earth. By rotating the mirror equal and opposite to the Earth's rotation rate, k_{eff}, which is the vector sum of the downward, k_1 and upward, k_2 going beams, no longer rotates at Ω_E in the freely falling frame. In order to do so, the retro-reflection mirror must sweep 167 μrad per experiment and then reset back to its original position for the start of the next trial.

The dominant systematic error from rotation is due to the coriolis force. The differential force between the two isotopes is proportional to $\mathbf{\Omega_E} \times \delta\mathbf{v}$ where $\delta\mathbf{v}$ is the difference in transverse velocity between the two isotopes. After constraining the transverse velocity difference between ^{85}Rb and ^{87}Rb to be ~ 1 μm/s, the systematic error phase shift from the coriolis force is $\sim 10^{-11}$ rad which is 5 orders of magnitude larger than our desired measurement accuracy. Therefore we need to reduces Ω_E by 10^5. Practically this requires positioning the retro-reflection mirror at three positions for the three pulses during the sequence with an accuracy of $10^{-5} \cdot 167$ μrad which is ~ 1 nrad.

Modern piezo-actuators with resistive readouts are capable of accuracies of 10^{-5}, and we outsourced the fabrication of these actuators to Mad City Labs in Madison, WI. They provided three actuators with a noise floor of 2.7 pm/$\sqrt{Hz}$. These actuate a triangular base platform with a 10 cm lever arm, on which the retro-reflection mirror is mounted. This gives an angular noise floor of ~ 0.5 nrad/$\sqrt{Hz}$ which meets our requirements.

In order to avoid spatial phase variation across the laser beam, all interfaces between the retro-reflection mirror and the atoms need to be eliminated. If the laser beam passes through a window, then temperature gradients in the glass can cause variations in the index of refraction which cause phase variations in the beam (i.e. the window can act as a very weak lens). Secondly, if the atom optics laser beam passes through air, then temperature gradients again cause variations in index of refraction which leads to spatial phase variation across the beam.

Therefore the RCS system is housed in a vacuum chamber located directly below the cooling chamber. The piezo-actuators are mounted directly to the walls of this vacuum chamber. This chamber is attached via flexible bellows to the bottom of the lower VAT all-metal gate valve. These bellows allow the entire chamber to tip and tilt in order to grossly adjust the position of the retro-reflection mirror. The piezo-actuators allow fine adjustment. The lower chamber is attached to the concrete floor via three set-screw-opposed 3/4 inch bolts. These opposed fasteners allow for tip-tilt actuation of the entire chamber while also opposing the 600 lbs restoring force on the chamber from atmospheric pressure.

4.4 Laser System

As stated above, we use laser cooling and trapping techniques to cool our test masses, and atom interferometry to measure their accelerations, and detect them using near resonance absorption imaging. Since both ^{85}Rb and ^{87}Rb can be treated as effective three level systems, we need four distinct laser frequencies for the cooling phase. For manipulating the atoms during the interferometer, we use a far-detuned, $\Delta \sim 200$ GHz laser in order to reduce losses due to spontaneous emission and to create a similar Rabi frequency for both isotopes. We have to use a single laser frequency to manipulate both species in the atom interferometer in order to eliminate systematic errors caused by laser phase noise. We also need spectrally pure frequencies to measure large optical depths during imaging. Therefore we need to generate three clean frequencies (one for interferometry, and one for imaging each species), and four less-clean frequencies for cooling. We do this with four separate diode lasers

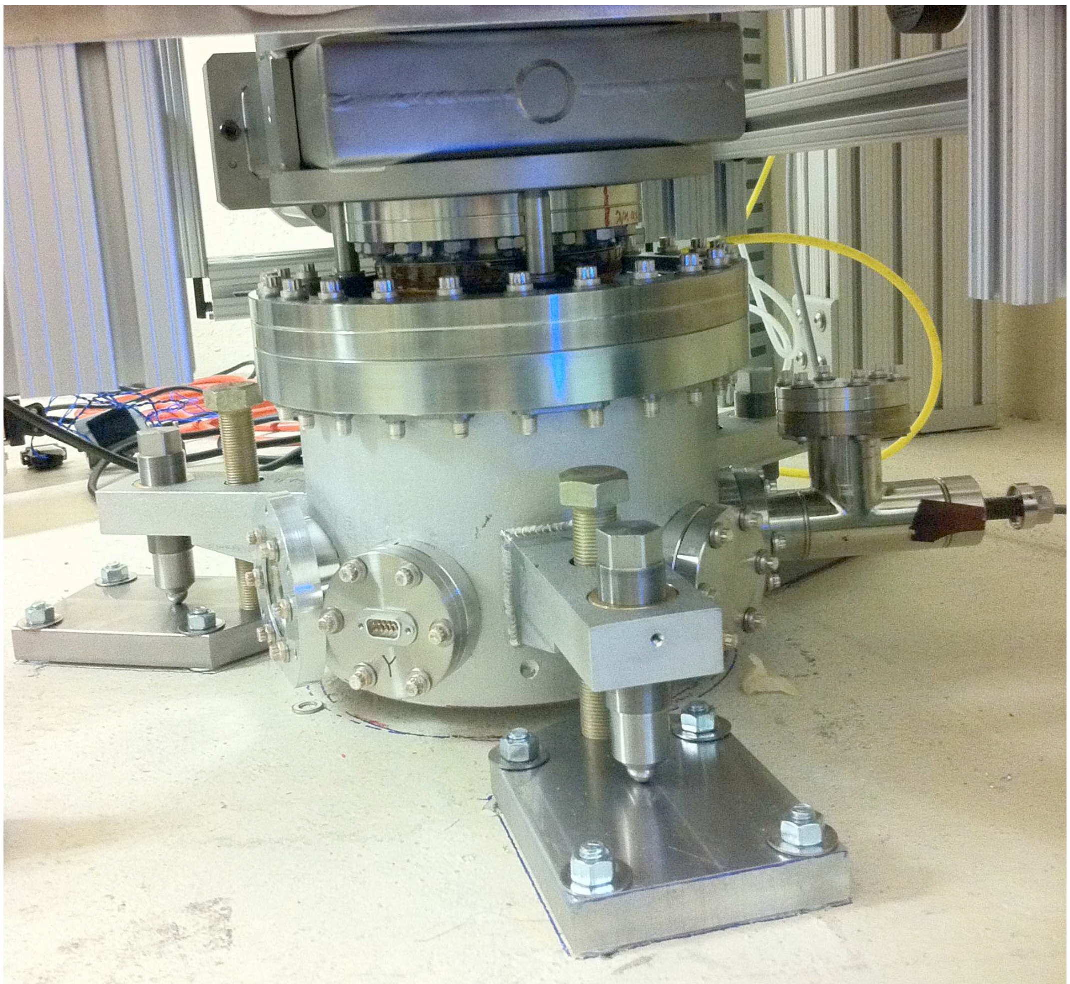

Figure 4.3: The installed rotation compensation system.

(one for cooling, one for interferometry, and two for imaging) along with fiber phase modulators as described below.

4.4.1 Laser Cooling

To create an efficient Magneto-Optical Trap (MOT), the four cooling frequencies need to be referenced to within ~ 1 MHz of their atomic transitions, and thus need to be actively stabilized. We use a single master laser locked 449 MHz blue of the $F = 3 \rightarrow 4$ and $F = 3 \rightarrow 3$ cross-over transition in ^{85}Rb via an AOM and saturation absorption spectroscopy (the lock bandwidth is ~ 5 kHz). From this master laser, the two cooling frequencies are generated via NLTL-driven fiber phase modulators (see section A.1) and the repump frequencies are generated from the first and second sidebands of another fiber phase modulator. The output of the three phase modulators is then combined using two PBS cubes and fed into a tapered amplifier. This method allows for efficient modulation with $\sim 80\%$ of the light in the desired sideband without the complication of trying to span ~ 7 GHz with AOMs or using separate lasers for cooling and repump.

In order to generate both cooling and repump frequencies for ^{85}Rb and ^{87}Rb using three modulators (instead of four), we must lock to a specific frequency relative to the Rb frequency manifold. We use a single modulator to generate repump light for both ^{85}Rb and ^{87}Rb by using the 1st-order sideband for ^{85}Rb and the second order sideband for ^{87}Rb. This modulator must be run at 2.526 GHz which is the distance between the repump lines for the two isotopes. Therefore the master laser must be located 2.526 GHz red of the ^{85}Rb repump line or 449 MHz blue of the cooling line. This also determines the frequency for the ^{87}Rb cooling modulator, which happens to be 1.067 GHz. Though this is a little complicated to get all of the various frequency offsets correct, it proves to be very stable and easy to adjust the relative detunings for each cooling frequency.

4.4.2 Detection

We use absorption imaging to detect the atoms during experiments. This requires an agile laser frequency source to sweep the detuning of the imaging light as desired, and a spectrally pure laser source to measure large optical depths. In order to do this we use a separate diode laser (New Focus Vortex) locked to a sideband of the master laser generated by another fiber phase modulator. We can adjust the frequency of the imaging laser by changing the frequency of the phase modulator or by adjusting the lock parameters. This allows variation of the imaging laser light frequency by up to 20 GHz while still referenced to an atomic line.

We lock the two lasers together by measuring their relative frequency on a photodiode (Thorlabs DET-210) and a frequency-to-voltage converter to create an error signal. The frequency-to-voltage converter consists of power splitter, two different path-length transmission lines and a mixer. The voltage at the output of the mixer depends on the phase difference between to the two paths and this is frequency dependent. This is a microwave interferometer and the output of the mixer is a sinusoidal error signal which, when combined with an offset voltage has a large locking range using standard PID control techniques. An interferometer with a free-spectral range (FSR) of ~ 30 MHz works nicely for tight locks around a specific offset frequency. For a large tuning range without mode-hopping, we were able to attain reliable lock performance with an FSR of 300 MHz. The lock point can be easily scanned by changing the offset voltage as the PID controller naturally locks to the zero crossing of the error signal.

We create absorption images by illuminating an Andor iXon camera with two low-power imaging pulses from the imaging laser. The first pulse is taken with atoms present and their spatially dependent absorption is measured on the camera. The second pulse is taken without atoms present and serves as the reference image. The difference between these two images is the desired atomic signal. We also take a background image with no imaging light to compensate for the dark counts on the camera. The three images are automatically processed by our custom driver software and the processed image, as well as the raw data is saved to a networked file server. This allows for rapid post processing by an experimenter.

Various systematic errors are present in this method and they have been nicely detailed in Ref. [130].

4.4.3 Interferometry Lasers

This laser is detuned by ~ 100GHz from any atomic resonance to avoid atom loss due to spontaneous emission during the interferometer sequence. Large frequency variations (~ 1 GHz) of this laser do not cause first order phase differences between the species. However, frequency variation of this laser changes the relative Rabi frequency between the two species, which, for fixed pulse lengths, adjusts the relative pulse areas. By not applying a precise π or $\pi/2$ pulse, we lose contrast in the interferometer. Therefore this laser need only be tightly referenced (~ 1 MHz)for repeatability between experimental trials.

We generate this far detuned laser light by a separate diode laser (New Focus SWL, 75 mW at 780.65 nm) that is offset locked to the master laser via a transfer cavity. We lock the cavity to the master light and then lock the interferometer laser to the cavity, thus tightly referencing the interferometer laser to a Rb atomic line.

In order to lock the cavity to the master laser, we inject master laser light that is phase-modulated at 10 MHz, the same light that is used for the probe in our saturation-absorption spectroscopy, into the cavity and measure the cavity transmittance on a photodiode. This is a standard dither-lock error signal and we use it to apply feedback to a piezo-actuator bonded to one of the cavity mirrors to maximize the transmittance.

Now that we have a referenced cavity, we can apply the same technique to lock the interferometer laser to the cavity. In this case we apply an effective phase-modulation spectrum to the interferometer laser light via an AOM driven with multiple frequencies. This light is injected into the cavity and we measure the reflected light via another photo-diode. This dither-lock error signal is used to apply current feedback to the diode and stabilize it with respect to the cavity.

By using this technique, we are able to replicate the master laser error signal spectrum on the feedback signal to the interferometer laser. This means that the transfer cavity lock is better than the saturation-absorption lock of the master laser to the Rb atomic line.

The cavity was originally designed to be a narrow line-width cleanup cavity (~ 5 kHz) at 850 nm, but when used at 780 nm it has a 1 MHz line-width, which is narrower than the natural line-width of Rb ($\Gamma = 2\pi \times 6.066$ MHz), which leads to the tighter lock performance.

4.4.4 Laser Amplification

All of the atomic manipulation used in this experiment requires a significant amount of laser power, and importantly, more than is available for a typical diode laser. Therefore we have to amplify our modulated light before delivering it to the atoms.

For our laser cooling light, we use 500 mW Eagleyard tapered amplifier (TA) chips in a custom housing that provide ~ 15 dB of amplification. The housing provides a strain relief for the c-mount chip and thermally couples it to a copper heat-sink using an indium pad. The heat-sink is actively cooled (or heated) by a thermo-electric cooler (TEC) that is water-cooled via an aluminum heat-sink. The TAs are actively thermally stabilized by feedback on the TEC current using a thermistor embedded in the copper heat-sink and a Wavelength Electronics temperature controller. We drive the TAs with a Wavelength Electronics current controller that provides currents up to 2 Amps. With this setup, after spatial filtering, using either a pinhole or an optical fiber, we are able to typically achieve output powers of > 300 mW from each TA.

Because the experiment is so large, we have to go through more than typical gymnastics to deliver light to the atoms. We manipulate the laser frequencies in a temperature stabilized optics lab that is located 30 m from the pit where the vacuum system is located. We transfer the light to the vacuum chamber via long fibers (30 m to the top of the pit and 40 m to the bottom). These fibers are long enough that trying to transmit > 500 m will result in significant stimulated Bruillion scattering limiting the available transmitted power. Therefore, the final TAs are located adjacent to the vacuum system and we amplify the light after transmitting through the long fibers.

For the laser cooling and trapping light, we use a single 500 mW TA located at the bottom of the pit. This provides ~ 260 mW to the 2D-MOT and ~ 10 mW to each direction of the x and y axes of the 3D-MOT. The z-axis of the 3D-MOT is provided by the interferometer TAs (described below). It is important to note that due to our frequency modulation scheme, only $75 - 80\%$ of the delivered power is the correct frequency.

We use the same basic scheme to amplify our interferometer laser light, but in this case we need more laser power because of the large detuning (~ 200 GHz). Here we use 1 Watt Eagleyard TAs in the same physical package described above. The main difference is in the current controller, which is capable of sourcing up to 10 Amps. This supply can apply a current pulse to gain larger than manufacturer specified output powers for a short period of time. We also use a larger TEC and temperature controller to handle the higher heat loads.

4.4.5 Lasers Conclusion

This system of using a single master laser and high frequency modulators proved to work well in practice. It is flexible, stable, and reliable. Modulator frequencies can be changed in 100 ns using widely available RF switches, which allows for easy application of repump and depump pulses. Requiring only a single laser lock makes for easy debugging and high up-time. The main disadvantage to this system comes from the lack of frequency cleanliness, since the NLTL modulation scheme invariably leaves spurious sidebands on the light. Thus using modulated light for absorption imaging is highly undesirable as it substantially limits the peak measurable optical depth, and we had to use a separate laser for imaging and another create another laser lock to reference it to Rb. Since this was an RF offset lock to a modulator sideband, it proved very reliable and highly tunable. Finally, the interferometer light required is too far detuned to be easily generated from the master laser, so this also required an additional laser. While it would be convenient to use only modulators to generate the required frequencies, a hybrid scheme proved to be required.

4.5 Experimental Control System

In order to control the hundreds of actions required to collect, cool, launch, and detect a cloud of atoms, we need a sophisticated computer control system that is easy to program, modify, and extend. Furthermore, because of the physical size of the apparatus, the system cannot be localized. It would be decidedly inconvenient to be required to sit at a specific computer in the office when trying to debug a signal located 60 m away and down an 8 m ladder at the bottom of the pit. Therefore we use a client-server-device architecture, with all elements communicating with each other via CORBA over the internet.

4.5.1 System Overview

The control system consists of a software network to connect, command, and control various hardware peripherals. The software network consists of a server, a client, and a series of ”devices” which are programs that control specific pieces of hardware. The primary hardware component is an FPGA-based embedded system which controls a variety of function generators, digital-to-analog converters (DACs), analog-to-digital converters (ADCs), and TTL lines. Other hardware can include any peripheral which can be connected to a computer, though GPIB-enabled devices are by far the most common.

As mentioned above, each piece of hardware is controlled by a of c-program called a "device." A "device" has "channels" and "attributes". A "channel" is something on which an "event" can be generated during a "timing sequence," which is a pre-defined set of actions that occur at set times. An "attribute" is a characteristic of a device that is true at all times after it has been set. A good example of a "channel" is a TTL line which can be changed by an "event" from high to low at a specified time. The temperature set-point of a chiller would be an "attribute," since it does not change often and typically does not have to change at a specific time to run the experiment.

A "timing sequence" is a set of "events" which occur on "channels" at specific times. Each trial run of the experiment uses a single "timing sequence" which contains all the "events" required to cool, trap, launch, and measure a cloud of atoms. These "events" are actions that cause lasers to turn on and off, detunings to change, and currents through magnetic coils to ramp up, etc... Each "event" translates to a specific action at a specific time on a "channel." For instance setting the output voltage on a DAC to 3.2 V, 1.10500100 seconds after the start of the experimental trial would be a valid "event."

The user is able to specify a sequence of "events" by writing Python code using the client, a Java console that can run on any network connected computer. The client sends the Python code for a "timing sequence" to the server, where it is converted into a set of "events" on "channels" for every "device" that is to be used. Each "device" converts the "events" into physical actions for that peripheral. These actions are executed at their specific times after all "devices" are triggered by the server to start the "timing sequence."

The central server provides communication between "devices" and the client. It parses the Python code into a list of "events" that are passed to the "devices" via CORBA over the network. The client serves as a terminal for the user to enter Python code and monitor the performance of all of the connected "devices."

After a "timing sequence" has been played, all the information related to that experimental trial is recorded. The Python code required to generate that "timing sequence" is saved. All measurements taken by devices, parameters used, "attribute" settings, etc... are saved into an XML file for easy browsing at a later time.

Finally, it is often necessary to run a series of "timing sequences" with slightly different parameters. For example, an experimenter might want to determine how laser detuning affects the atom number in the MOT. A series or "sequence" of "timing sequences" can be run in an automated fashion via the client.

4.5.2 FPGA

This section provides a brief overview of the VHDL code and hardware that constitute the FPGA component of the timing system. This is by no means an exhaustive description of the system, but it is rather meant as a primer to be used in conjunction with the code itself and the schematics.

FPGA Overview

A successful laser cooling and trapping experiment requires many events to occur within a few milliseconds. This requires a control system that is able to rapidly execute many different actions with rigorous control over the relative time between the events. We use an FPGA with a 100 MHz internal clock to provide digital commands to a selection of hardware boards. These boards can provide Direct Digital Synthesis (DDS), TTL signals (Digital Out), Digital to Analog Conversation (DAC), and Analog to Digital Conversation (ADC) with a timing resolution of 10 ns provided by the clock for the FPGA.

Each of these boards is a "device" which have "channels" that can take "events." The section below describes how this abstract concept is physically implemented in great detail.

The FPGA is directly connected to an embedded Etrax computer which provides an Ethernet interface to communicate with the rest of the control system, specifically the server via CORBA. The Etrax has a direct memory address pipeline into the FPGA in order to transfer and receive information. Through this pipeline a set of FPGA memory addresses are made available to the Etrax where it can read or write data. This allows direct control over the state machine which is programmed into the FPGA.

Each hardware board has a dedicated program, the "device," (written in c) running on the Etrax. These handler programs ("devices") communicate with the server via CORBA and receive lists of "events" for each "timing sequence". The "device" program process this list of "events" into an FPGA-readable sequence of commands and transfers them into the FPGA's memory. After this process is complete, the "device" program signals the server that it is ready to play the "timing sequence". Only when all "devices" (FPGA specific as well as all other "device" programs) have reported back to the server, can the "timing sequence" be played.

Each hardware board has a different set of requirements for what is to be loaded into the FPGA memory, but they all share a few characteristics. A typical command for an "event" consists of two 32-bit numbers which correspond to (time, value) - i.e. the time the "event" is to occur and what the "event" is supposed to do. For example, the allowed values for the digital board are "1" and "0". The low-level code on the FPGA converts the (time, value) pair into a physical action. It is imported to note that an "event" specified in Python may correspond to many (time, value) pairs to complete the desired action.

FPGA State-machine

The FPGA state-machine is the VHDL code that handles the overall behavior of the low-level board specific code and the Etrax. It consists of a trigger module, a memory bank, a timing core for each hardware board connected to it (including the Etrax), and a Wishbone memory bus connecting all of these modules together.

The state machine performs two basic functions. First, it pipelines communication between the various timing cores, the memory, and the Etrax via a Wishbone bus and an arbiter. Second, it plays back all the events in its memory when commanded to start via the trigger module. In order to understand how the system works, we will focus independently on the two aspects of its operation.

The communication pipelining is required because the FPGA state-machine has a single memory repository that can be accessed by eight different entities, one of which is the Etrax and the others are timing cores that operate a different hardware board. Each of these entities can request to use the Wishbone bus at any given time and a memory arbiter handles the requests. This allows all seven timing cores and the Etrax to all have access to the same memory addresses and not run into conflicts. The memory arbiter is a classic round-robin architecture that gives each member equal priority in order of request. Whenever a member raises its request flag, it is added to the queue and served in the order received. A Wishbone bus operation requires five clock cycles (50 ns) to complete and the Wishbone data vector is 4 bytes wide, which translates to a memory throughput of 80 megabytes per second. In the current implementation, the FPGA has 16 kilobytes of memory. Therefore this memory architecture is not a bottleneck in terms of performance for loading or for playback, as events are typically spaced by 1 μs or more.

The second function of the FPGA state machine is to provide accurate playback of the "events" loaded into its memory. As alluded to earlier, each hardware board has its own timing core inside the FPGA state machine. This timing core consists of a two "event" deep memory buffer and a comparator. The memory buffer width is hardware board dependent but typically consists of (time, value) where both are 32-bit numbers. During every clock cycle the comparator compares the current time with "time" of the current event. If the two values are equal then the timing core executes the action appropriate for "value."

As soon as a timing core has executed its current "event", it transfers the next "event" in its buffer (and the last one in its buffer) to be the current "event". This allows two "events" to be executed on adjacent clock cycles. Next, it makes a request on the Wishbone bus to refill its buffer. This is fulfilled by the round-robin arbiter as soon as possible. This puts a limit on how fast "events" can be played based on the Wishbone bus arbiter. Practically we have never run into this limit, and if required the internal memory buffer in each timing core could be replaced with a FIFO. This is not currently implemented due to the limited size of the current FPGA chip.

This system is coded in VHDL and is currently hardware generic. It is currently implemented using a Xilinx Spartan-3 FPGA, but the next version will use a Spartan-6 in order to gain resources to allow for more than seven timing cores and the Etrax interface.

Digital Out

The Digital Out Board provides 24 TTL output lines that can be changed independently at 100 MHz. The outputs are buffered to prevent damage to the board and they provide a high voltage level of 5 V and a low of 0.2 V which meets TTL specifications, but the user should be warned as the low voltage is not ground.

The Digital Out board is the simplest to command. It only accepts a single 24-bit number with the value of each bit corresponding to the output of that digital out channel. For example, to set "channel" 2 to "1" and all others to "0", the command value would be "00000000000000000000000000000100" in binary or "4" in hexadecimal.

Slow Analog Out

The Slow Analog Out board provides 40 channels that can provide output voltages from -10 V to 10 V, and it is able to change the output value of a single channel every 1 μs. All of these output channels are buffered with output op-amps to protect the internal logic of the digital-to-analog converter. Thirty-six of these output buffers (Ch 4 - 39) are able to drive 600 Ohm loads, while four of the output channels (Ch 0 - 3) have high power output buffers and can drive 75 Ohm loads.

From a high level programming point of view, the board is easy to command. It requires a 24-bit command to be loaded into "value" (i.e. the upper 8 bits are do-nothing bits). The first 14 bits specify the voltage output value in a continuous range from -10 V to 10 V. Bit 14 determines whether or not to update the outputs on all channels to that stored in memory. This is useful if a user wants to update all the channels at once, rather than change one at a time. The channel is selected by bits 15 to 20, with 0 corresponding to channel 0 and so forth. Bits 21 and 22 set the registry value which is typically set to "11" for our purposes. Finally bit 23 resets the board and sets all outputs to -10 V.

Fast Analog Out

The Fast Analog Out board is very similar to the Slow Analog Out board, except that it takes a 16-bit number for its output value and it only has 2 channels. It can also provide output voltages from -10 V to 10 V, and it is able to change the output value of a single channel at 10 MHz. In practice, this update rate is limited slightly by the FPGA interface to ~ 3 MHz.

In order to command this board, the first 16 bits of the command determine the output value, again in the range of -10 V to 10 V. Bit 16 commands to board to write to "Ch 0", and bit 17 commands the board to write to "Ch 1" - only one of these pins can be raised at a time. Finally bit 17 and bit 18 command values to be loaded into "Ch 0" and "Ch 1" respectively. Again only one of these pins can be raised at a time. In this parlance "write" refers to physically changing the output voltage of the channel, while "load" refers to putting a new value into memory so that it can be "written."

DDS

The DDS card provides 4 output channels that can synthesize frequencies up to 250 MHz with 32-bit precision with arbitrary phase (32-bit) and amplitude (10-bit). The board is capable of sweeping any of these parameters, while holding the others fixed with 8-bit precision on the sweep rate and sweep time.

For applications where signal cleanliness is less important, the DDS is an excellent choice. However, due to its internal method of generating frequencies, the output signal contains significant spurs. Even after low pass filtering the output at 250 MHz (which is essential), the largest spur can be only 8 dB lower than the carrier. However, when used for driving narrow-frequency range circuits such as AOMs, these spurs are effectively filtered.

The performance and flexibility of the DDS comes at the price of a complicated interface. Because the DDS requires 32-bit data values to specify its frequency, it cannot be commanded with a single 32-bit number. Rather it needs two 32-bit numbers, one to specify the output value and another to specify the options for the device. Therefore the DDS takes three 32-bit numbers in memory instead of two like all the other devices. The data format for the DDS is (time, dds command, dds data) where all three are 32-bit numbers. The DDS can perform complicated functions, such as sweeps and amplitude ramps - this is all handled by the c-code driver which runs on the Etrax. I kindly refer the reader to the c-code for documentation.

Fast Analog In

The Fast Analog In board provides 16 bit analog-to-digital conversion on two channels at up to 1 MHz total throughput. Practically this input speed is limited by the data transfer between the FPGA memory and the Etrax. The board itself is capable of running at speeds of up to 3 MHz in continuous acquisition, but for the sake of flexibility we are using it in the slower, single shot mode.

The Fast Analog In board is unique among the boards in that it reads a value on command rather than outputs a value. This poses a unique set of problems in that this value must be stored and be able to read back by the computer in a reasonable time frame. The scheme implemented writes the "measured value" into the location where the "measure command" was stored. Therefore after a timing sequence has been run, the memory bank on the FPGA will contain a sequence of (time, measured value) pairs instead of the (time, value) pairs that were there originally.

The command structure for the Fast Analog In is very simple, since the board only needs to know when to make a measurement and on which of its two channels. Therefore the board takes only two bytes of data, the first bit commanding a read on "Ch 0" and the second bit a read on "Ch 1", where these two bits are mutually exclusive.

After the "timing sequence" has been executed, the c-program must read back through the memory and record the values stored in the FPGA to another data file, otherwise they will be re-written when a new "timing-sequence" is loaded onto the FPGA. Again this is handled by the c-program driver that runs on the Etrax.

4.5.3 Peripherals

The network architecture of the control system provides an extensible interface that allows any machine which supports CORBA to become a "device" that can be controlled via the server and client. This typically refers to computers running c-programs that have peripherals attached to them via another interface (e.g. USB, GPIB, parallel-port, etc...). These "devices" are not typically referenced to the FPGA clock, unless they accept an external trigger, and therefore they are not usually used to perform timing critical tasks. However they are extremely useful for automating various other aspects of the experiment.

GPIB

Most commercial test and measurement equipment supports external control via GPIB. This allows easy control of oscilloscopes, function generators, network analyzers, voltage sources, and many other very useful devices. In order to command these devices, we use a USB-GPIB or an Ethernet-GPIB converter and a c-program layer to translate GPIB commands in language understood by the server.

Each GPIB-computer interface requires a "device" that is a GPIB-controller. All c-program GPIB "devices" communicate via the network with the GPIB-controller "device" that corresponds to the physical GPIB-controller that their instrument is attached. The end result is a set of c-programs that exactly mirrors the physical set of GPIB controllers and devices. For example, consider a computer with a USB-GPIB controller attached to it, and an oscilloscope and a function generator are connected on the GPIB bus. Then there would be a GPIB-controller "device" running on the computer with the USB-GPIB controller attached to it. There would also be two other c-programs, running on any computer (even the same computer) with network access to the computer with the USB-GPIB controller. One of these programs would contain the specific GPIB-commands required for the oscilloscope and the other would handle the function generator.

The c-program layer is transparent and easy to use. It only requires the device programmer to assign a "channel" or "attribute" for each device specific GPIB command. Then these can be accessed via the client as usual. Any "device" can access any GPIB command using low-level programming as well. This is quite useful for automating tasks, such as recording an oscilloscope trace.

Since almost all the code required to create a GPIB "device" is shared between all GPIB devices, it is stored in a c-library and does not have to be reproduced by a "device" programmer. Rather the "device" programmer need only create a few lines of code which deal with the specific GPIB commands for that piece of physical equipment.

Further information is best gathered from example code, and again I kindly refer the reader to any of the GPIB "device" c-programs for reference.

Other External Equipment

Other external devices work in much the same way as the GPIB example. For any piece of equipment which uses some external interface, the c-program "device" must include a driver that handles communication and translation code to convert desired actions into "channels" or "attributes." This is purposefully general and in practice it allows nearly anything to be converted into a "device" to use with the control system.

In this manner, we have created "devices" for our Andor iXon camera, the controller provided by Mad City Labs for the rotation compensation system, the DIO-LAN cards which communicate with our high power TA current drivers, several different AOSense laser lock boards, and a stepper motor controller. These have proved very useful for automating and trouble-shooting the experiment.

Finally, a "device" is not limited to being a single piece of equipment. Rather it can be a collection of peripherals that together perform a desired function. For example, the laser lock contains several different control loops running at different bandwidths. The first of these is a PID controller built by AOSense that has a bandwidth of 5 kHz. The computer interface is used to set gains, but the feedback is an analog circuit. There is an additional external loop which sets the piezo voltage on the master laser via a GPIB command. This loop uses a commercial ADC to measure the feedback voltage from the AOSense board via USB. It then determines if the low-passed value of that voltage is too large (i.e. close to the end of its range). If so, then it commands a piezo voltage change via GPIB. This allows for large temperature excursions in the lab and still maintains the laser lock. Even though this requires four different physical peripherals to create a single "device," the laser lock, it is relatively easy to program and is very reliable.

4.5.4 Computer Control System Conclusion

Given the complicated nature of ATEP, it requires significant automation to create a reliable and repeatable experimental apparatus. A typical day of running the experiment produces ~ 2000 trials and nearly 10 Gb of data. Each of these trails required a timing sequence and the interaction of nearly a hundred different pieces of experimental equipment. As of now, most parameters required to recreate this experiment are stored and in the future all critical pieces of data will be recorded. The significant effort expended on creating this computer control system has proven well worth it through ease of troubleshooting and taking vast quantities of data.

4.6 Evaporatively Cooled Atom Source

4.6.1 Optically Plugged Trap

In order to achieve the required tight kinematic control over the isotopes, we need to cool the atoms below 1 μK. Evaporative cooling in a magnetic trap is easiest method for doing so in our experimental apparatus. It is well-known that at the 0-field point of a quadrupole magnetic trap, trapped atoms can undergo a spin-flip to an un-trapped state, thus resulting in rapid atom loss. Many techniques have been employed to combat this problem, most notably the Time Orbiting Potential (TOP) and Optically Plugged Trap (OPT) - both of which resulted in Nobel Prizes for their inventors. We originally designed our chamber to use a TOP trap, but due to its geometry, RF radiation above 1 kHz is significantly attenuated. We need a TOP field of ~ 10 Gauss at a frequency of ~ 10 kHz so that the orbiting frequency is much faster than the harmonic trap frequency. Then the orbiting potential is effectively averaged by the atoms. Unfortunately the chamber is a very effective Faraday cage for high frequencies, and this is not feasible. Therefore we are forced to use an OPT to reach high phase space densities and low temperatures.

The physics behind the OPT is simple - use a repulsive optical potential to create a repulsive barrier at the 0-field region of the quadrupole trap. In practice this requires a far detuned laser with lots of power in order to create a tall enough optical potential without significant losses due to spontaneous emission. The energy shift from a far detuned laser is

$$\Delta E = \frac{\hbar}{2}\left(\sqrt{\Omega^2 + \delta^2} - \sqrt{\delta^2}\right) \tag{4.1}$$

where δ is the laser detuning and Ω is the Rabi frequency. This scales as $\frac{power}{detuning}$ for a fixed focus in the far detuned limit. Therefore, for a fixed plug height, we need to linearly increase power as we increase detuning. Given this scaling, the best choices we found for high power, far blue detuned light, were at 532 nm. Therefore we use a 6 W Lighthouse Photonics laser at 532 nm to provide the blue detuned laser light. At a 30 μm focus, we are able to achieve a ~ 10 μK trap depth, as measured by imaging the plug laser on an Rb molasses. This is sufficient to prevent significant Majorana loss at low temperatures.

4.6.2 Microwave Evaporation

In order to increase the phase-space density of our magnetically trapped sample, we use an energy dependent microwave knife to selectively remove hot atoms from the sample [129]. Since the atom energy shift in a quadrupole trap is linearly dependent on distance from the trap center, we can selectively address different kinetic energy classes by tuning the frequency of the microwave knife. This allows for selective removal of hotter atoms and the evaporative cooling of the sample.

To load atoms into the magnetic trap, we follow a slightly peculiar sequence. We have to deviate from the standard procedure in the literature because the large quantity of steel and aluminum present near our quadrupole coils allows large eddy currents to persist for up to 50 ms after a rapid change in the magnetic field. Therefore we load atoms into the magnetic trap by first detuning the laser light red by an additional 70 MHz, then ramping up the magnetic field gradient relatively slowly (in ~ 20 ms typically) to to 40 Gauss/cm while keeping the detuned MOT light on. Before

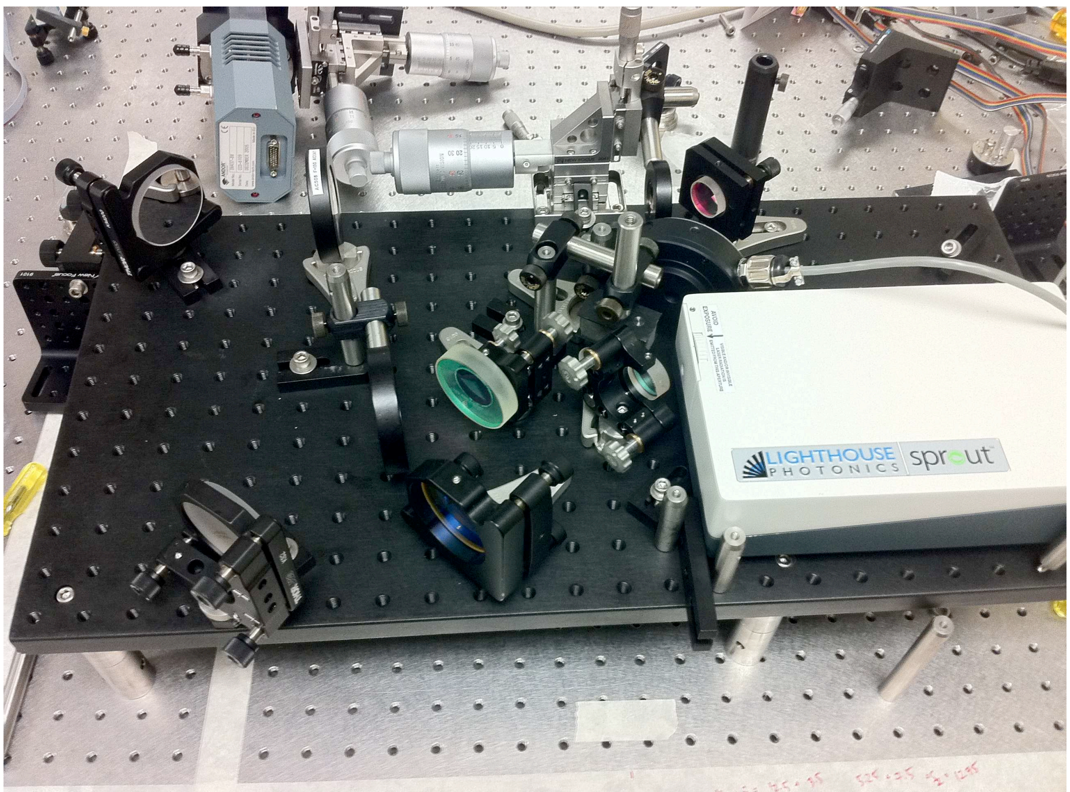

Figure 4.4: The optics system for the blue detuned optical plug.

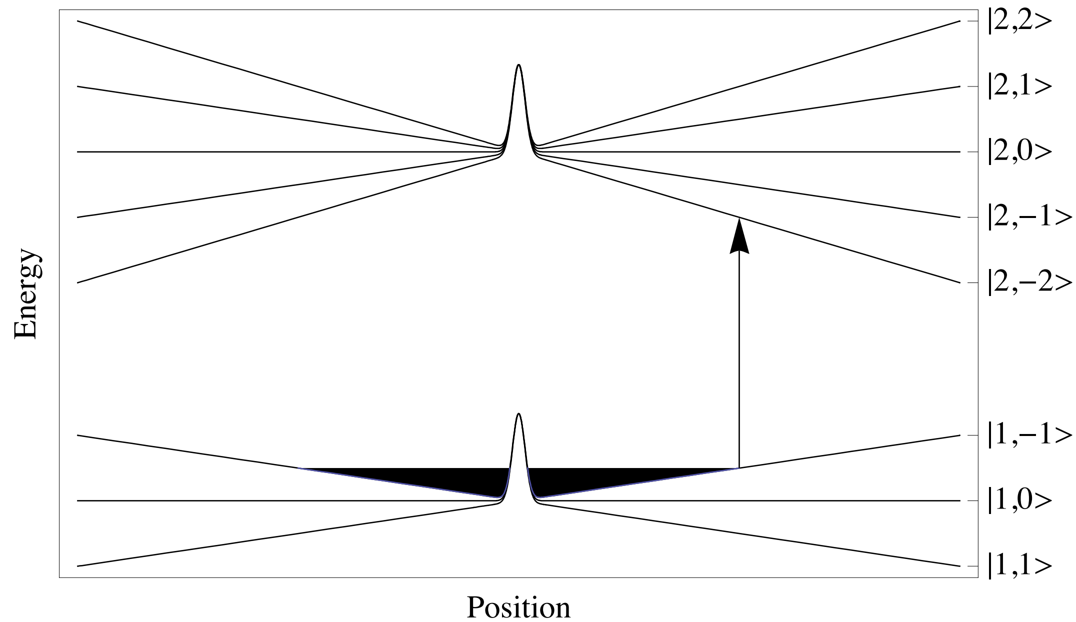

Figure 4.5: Microwave Evaporation Scheme for an Optical Plug Trap. By driving microwave transitions for atoms at a specific energy, we can selectively remove hot atoms from a trapped sample. This has been shown to lead to runaway evaporation and eventually BEC.

the field gradient has reached its full value, we turn off the re-pump light, which allows atoms to occasionally fall down from the $|F = 2 >$ state into the $|F = 1 >$ state, where the $|1, -1 >$ atoms are trapped in the 40 Gauss/cm quadrupole magnetic trap. Finally, an additional de-pump pulse, resonant with the $F = 2 \rightarrow 2$ transition is applied and $> 99\%$ of the atoms are transfered into the $F = 1$ state where 1/3 of them are trapped.

In order to remove atoms from the magnetic trap, we need to drive a transition from a trapped state to an un-trapped state. Since the trapped atoms are in the $|1, -1 >$ state, it is possible to drive microwave transitions to the $|2, -2 >$, $|2, -1 >$, and $|2, 0 >$ states while satisfying selection rules. Of these three states, none are trapped, since the g-factor for $F = 1$ and $F = 2$ states have opposite signs. The relative energy shift is the smallest between $|1, -1 >$ and $|2, -2 >$ since these states differ in energy by $3g_F\mu_B B$ since $\Delta m_F = 3$ in this case. The microwave frequency source is a $\sim 10 - 200$ MHz DDS channel mixed with a 6.8 GHz base frequency. In order to avoid the effects of spurious sidebands due to the imperfection of the DDS, we tune the microwave source to the lowest frequency transition and increase the frequency to reduce the energy of the evaporated atoms. This evaporation scheme is shown schematically in Fig. 4.5, with the energies of the trapped atoms shown as the filled-in region and the microwave knife transition represented by an arrow between the two states.

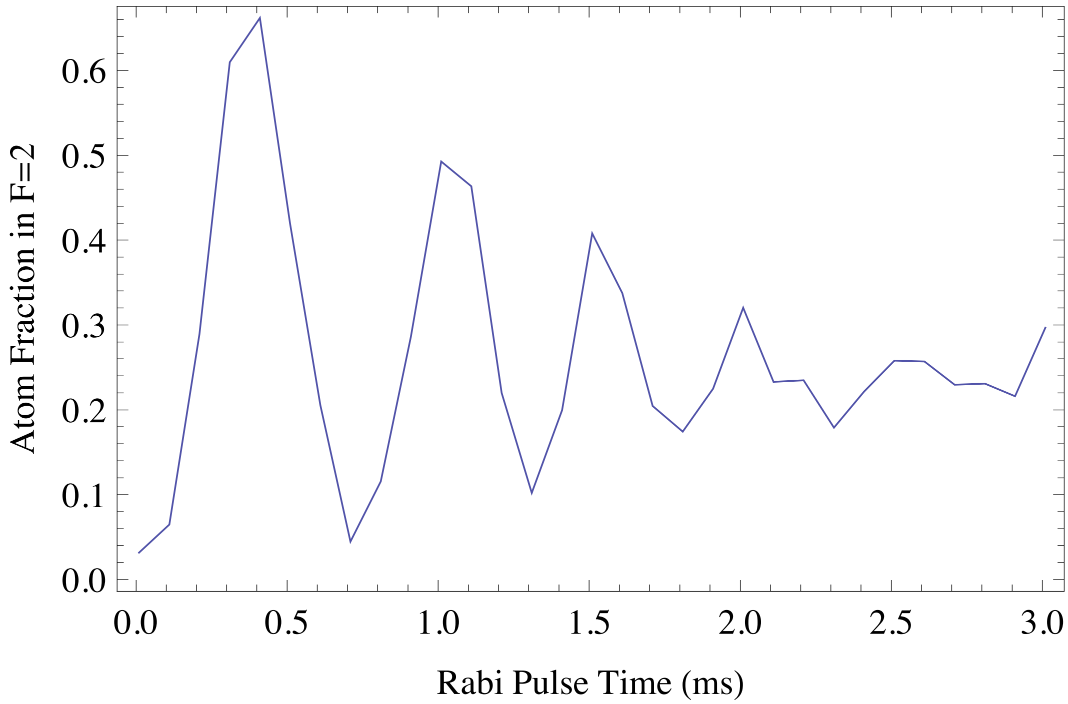

Figure 4.6: Typical Rabi Flop Data. Population in the $F = 2$ state after applying a boxcar microwave pulse of a given duration. By removing the horn and only using a bare antenna, we achieved a factor of 10 increase in Rabi frequency.

In order to efficiently remove atoms from the sample, a sufficiently fast enough Rabi frequency is required, otherwise evaporation proceeds too slowly. This proved difficult given our chamber geometry and insufficient access for the microwaves to enter our stainless steel chamber (essentially a large Faraday cage). By using a microwave antennae that is impedance matched to free space and no horn, we were able to move the horn sufficiently close to the chamber to have a > 50 kHz Rabi frequency (see Fig. 4.6).

4.7 Conclusion

The system is proceeding rapidly towards condensation and eventually atom interferometry. We have achieved phase-space densities of 10^{-4} without the optical plug. Similarly, we were able to perform Raman spectroscopy using a conventional 2-photon transition. In the very near future, we expect condensation and large scale atom interferometry.

Chapter 5

Conclusion

After early demonstrations of light pulse atom interferometers, it became immediately desirable to take advantage of the linear increase in sensitivity with length. Several atom interferometers with baseline $L \sim 1$ m have been built, and the work described here aims to push L to 10 m and eventually 30 km! These proposed, and under construction, interferometers have sensitivities that offer tantalizing opportunities to test fundamental gravitational physics in a controlled laboratory environment.

The current effort at Stanford to build an 8.8 m atom interferometer is nearing completion. Currently, the interferometer region is shielded to have magnetic fields below 100 μGauss, the rotation compensation system is built and can steer the interferometer beam through the required 200 μrad, and atoms have been cooled to $\sim 5\mu$K at the base of the fountain. A lot of work remains before a test of the WEP at $\eta \sim 10^{-15}$, but the future looks bright.

Finally, the proposed AGIS-LEO gravitational wave detector offers a chance at detection of a predicted phenomenon of general relativity that has yet to be observed. The future of precision accelerometers promises to be quite exciting as these experiments begin to produce results.

Appendix A

NLTL-Driven EOMs

These appendixes discuss the use of Non-Linear Transmission Lines to the modulate the frequency and amplitude of light. The work presented here was done in collaboration with Jason Hogan, Sheng-wey Chow, and Mark Kasevich.

A.1 Serrodyne Modulation

A version of this work can be found online at http://arxiv.org/abs/0909.1834.

A.1.1 Introduction

Optical frequency shifting has broad commercial and scientific applications. For example, optical telecommunication FSK protocols can be implemented using single-sideband (SSB) frequency shifters[131]. In spectroscopy and laser cooling experiments, agile frequency shifts from the MHz through the GHz range are often required[132, 133].

Acousto-optic modulators (AOMs) are commonly used for MHz-level shifts, but a given AOM has a limited tuning range which is only a small fraction of its fixed center frequency. AOMs that can deliver GHz-level shifts typically offer this increased bandwidth by sacrificing diffraction efficiency. Broadband electro-optic phase modulators (EOMs) avoid many of the bandwidth and tuning range limitations of AOMs. Commercially available fiber phase modulators can have modulation bandwidths > 10 GHz. However, typical sinusoidal phase modulation is an inherently inefficient method of frequency shifting. At best, the fraction of the power in the first-order sideband is $\eta = (J_1(\beta_{\max}))^2 \approx 0.34$, where $J_n(\beta)$ are the Bessel functions of the first kind. Also, the presence of large undesirable frequency spurs at other harmonics can cause problems in some applications.

In this Letter, we use a serrodyne phase modulation signal[134] to demonstrate broadband electro-optic frequency shifting with high efficiency into the desired sideband and correspondingly small undesirable spurs.

We briefly review the theory of serrodyne phase modulation[135]. The electric field for the light exiting the EOM is $E(t) = E_0 \cos(\omega t + \phi(t))$ where ω is the optical frequency and $\phi(t)$ is the phase imprinted by the modulator. A direct frequency shift can be imparted by applying a linear phase ramp,

$$\phi(t) = (m\delta) \cdot t \bmod 2\pi m \tag{A.1}$$

where $\dot{\phi} = m\delta$ is the desired shift and m is an integer. This phase ramp need only extend from 0 to $2\pi m$ before resetting since cosine is 2π-periodic. The resulting waveform is a sawtooth with angular frequency δ and amplitude $2\pi m$. The first order ($m = 1$) condition results in a frequency shift δ, but higher order shifts ($m > 1$) are also possible. Although a serrodyne signal can theoretically mimic a linear phase ramp within a finite tuning range, it requires a high bandwidth to faithfully reproduce the discontinuities at the end of each period.

The usefulness of the serrodyne technique is limited by the quality of the sawtooth waveform that can be produced within an experimentally accessible bandwidth, and this gets increasingly difficult for higher frequency shifts. Early work was done with MHz-level frequency shifts[134, 136], and more recently improved spur suppression has been demonstrated at these frequencies[135, 137, 138, 139]. Higher frequency serrodyne shifts were achieved with photonic arbitrary waveform generation (PAWG) using mode-locked lasers[140, 141], but these techniques are complicated and have limited tunability. In addition, PAWG based on spatial-spectral holography[142, 143] as well as commercially available conventional arbitrary wave generators with > 10 GS/s sampling rates could conceivably be used to produce high-quality GHz sawtooth waves.

A.1.2 Serrodyne Experimental Setup

Here we generate high frequency (200 MHz - 1.6 GHz) sawtooth waveforms with good fidelity by using a Non-Linear Transmission Line (NLTL). In an NLTL, an electronic signal experiences an amplitude-dependent propagation speed[144]. This effect results in a steepening of the input waveform as the higher amplitude components catch up with the lower amplitude components. A sinusoidally driven NLTL therefore outputs an approximate sawtooth waveform at the drive frequency. The NLTLs that we use are commercially available, passive components and generate harmonic content out to > 20 GHz.

We directly drive a $LiNbO_3$ fiber phase modulator (Photline Technologies NIR-MPX850-LN08, > 8 GHz bandwidth) with the serrodyne signal generated by an NLTL and analyze the resulting spectrum with a Coherent model 240 Fabry-Perot spectrometer (Fig. A.6a). We drive the NLTLs by amplifying the output of a signal generator (HP83712A) with a broadband RF amplifier (Mini-Circuits ZHL-42W). The NLTLs produce high output power serrodyne signals, and since the fiber phase modulators have a low $V_\pi \simeq 8$ V, we can directly drive the modulator without post-amplifying, and thus further bandwidth limiting, the NLTL output.

We generate 780 nm narrow-linewidth (~ 1 MHz) laser light using a MOPA laser setup (New Focus Vortex injecting an Eagleyard Tapered Amplifier). A small fraction of the light (30 mW) is coupled into the fiber phase modulator, providing up to 14 mW of output light. We couple a fraction of this power into the Fabry-Perot cavity (FSR 7.5 GHz, resolution ~ 25 MHz).

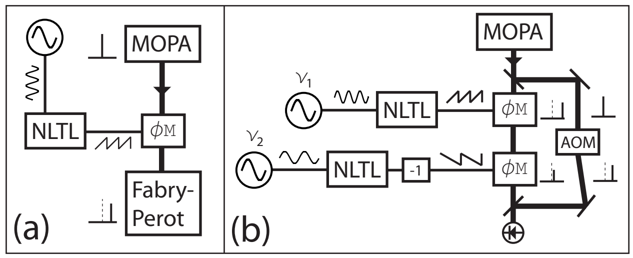

Figure A.1: Experimental setup. (a) A single phase modulator (ϕM) measured with a Fabry-Perot cavity. (b) Two phase modulators in a push-pull configuration measured via a heterodyne beat with the reference beam shifted 194 MHz by an AOM.

A.1.3 Serrodyne Performance

The serrodyne output of the NLTL varies with input frequency and amplitude. We characterized three NLTLs manufactured by Picosecond Pulse Labs with different frequency ranges and rated input powers: 7112-110 (300 – 700 MHz @ 29 dBm), 7113-110 (600 – 1600 MHz @ 29 dBm) and 7102-110 (300 – 700 MHz @ 24 dBm, requiring a ZHL-42W post-amplifier). Each NLTL was swept through its frequency range, and the input amplitude was varied at each point to determine the optimum serrodyne signal. We captured a data spectrum and a reference spectrum (no RF drive applied to the NLTL) at each drive frequency with a digital oscilloscope. Example traces are shown in Fig. A.2. Note that since the fiber phase modulator is electrically floating, we can reverse the signal and ground connections to apply a minus sign to the serrodyne signal (a balun could also be used).

The efficiency η of the serrodyne is the fraction of the unmodulated carrier power shifted into the desired sideband. Our serrodyne frequency shift has $\eta > 0.6$ from 200 MHz to 1.2 GHz as shown in Fig. A.3. Of particular interest are the regions from 400 – 500 MHz and from 1 – 1.1 GHz which maintain $\eta \sim 0.8$. This is comparable to a well-aligned single pass AOM. Similarly, we also characterize the cleanliness of the resulting spectrum by its spurious sideband fraction SF, defined as the ratio of the largest spurious frequency component to the desired signal. We measure SF < 0.2 from 200 MHz to 1.2 GHz as shown in Fig. A.3. In the region from 700 MHz to 1.1 GHz we find SF < 0.1. Figure A.2 (left trace) shows a shift with $\eta = 0.82$ and SF $= -16$ dB.

As a consistency check of the observed spectra, we made a time domain measurement of the applied sawtooth waveform $V(t)$ using a 63 GHz Agilent Infiniium DCA 86100B oscilloscope (Fig.

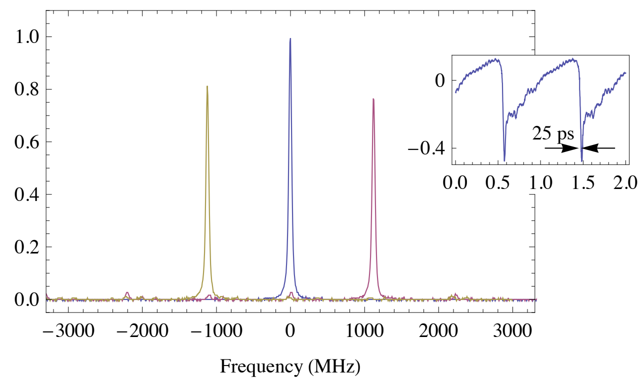

Figure A.2: Optical spectrum with and without serrodyne modulation normalized to the unmodulated carrier. The center (blue) curve shows the optical spectrum with no RF modulation. The right (red) and left (yellow) curves are the spectrum with a serrodyne modulation of +1.1 GHz and -1.1 GHz, respectively. The inset shows the applied 1.1 GHz serrodyne waveform in volts versus time in nanoseconds as measured by a 63 GHz oscilloscope.

A.2 inset). We then calculated the implied serrodyne spectrum using $\phi(t) = \alpha V(t)$ and fit the result to the observed spectrum shown in Fig. A.2 with a single free parameter α. The predicted spectrum matches the observations very well, with an RMS peak height difference of $< 2\%$ relative to the carrier.

In addition to the first order serrodyne results, we demonstrated second order shifts ($m = 2$ in Eq. A.1) by increasing the amplitude of the sawtooth waveform using the post-amplified 7102-110. This allows for a larger serrodyne frequency shift of 2δ without exceeding the frequency range of the NLTL. We observed second order shifts with $\eta_2 > 0.5$ and $\mathrm{SF}_2 < 0.25$ for a frequency shift range of 800 MHz to 1 GHz. These results were likely limited by the 4.2 GHz bandwidth of the post-amplifier.

A.1.4 Push-Pull Serrodyne

To obtain smaller frequency shifts, we arranged two identical fiber phase modulators in a push-pull configuration (Fig. A.6b). The two modulators are driven by NLTLs supplied by independently tunable amplified function generators running at frequencies ν_1 and ν_2, respectively. With this flexible setup we can continuously scan the serrodyne output at $\Delta\nu = \nu_1 - \nu_2$ from positive to negative frequencies out to the maximum frequencies of the NLTLs. We used a heterodyne measurement to characterize the push-pull spectrum for $\Delta\nu$ smaller than the linewidth of the Fabry-Perot (Fig. A.4). An AOM in the reference arm of the interferometer shifted the carrier by a fixed frequency of 194 MHz so that we could distinguish positive frequency beat notes from negative frequency beat notes with respect to the unshifted carrier. We find that higher-order harmonics of $\Delta\nu$ are suppressed by at least 25 dB compared to the desired shifted signal.

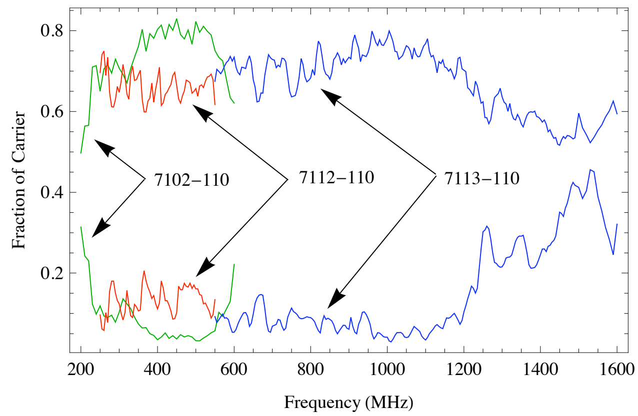

Figure A.3: Serrodyne modulation efficiency. The upper and lower curves show η and SF, respectively.

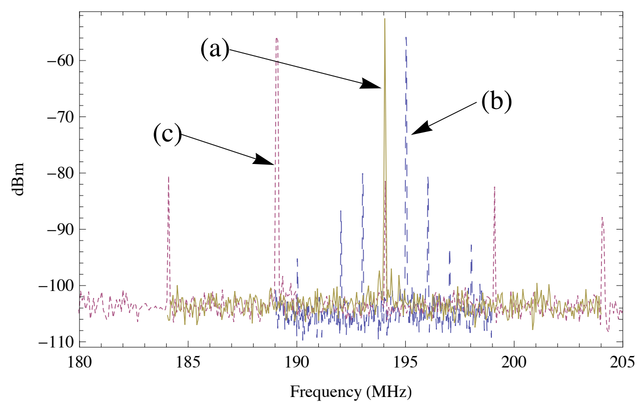

Figure A.4: Heterodyne spectra of push-pull frequency shift. (a) No modulation (b) $\Delta\nu = +1$ MHz (c) $\Delta\nu = -5$ MHz.

We investigated the phase noise of the serrodyne-shifted light by comparing the unmodulated heterodyne beat note signal to that of the push-pull signal out to 1 MHz with a resolution bandwidth of 1 kHz. The phase noise of the unmodulated signal was −95 dBc/Hz at 150 kHz offset. During serrodyne modulation we observe no increase in phase noise above this noise floor.

A.1.5 Conclusion

Optical SSB modulation using a dual Mach-Zehnder geometry[145] is another well-known solution for frequency shifting in the GHz range. However the maximum efficiency is $\eta_{\mathrm{SSB}} = (J_1(\beta_{\max}))^2 \approx 0.34$, at which point the -3δ spur is only suppressed by $\mathrm{SF}_{\mathrm{SSB}} = \left(\frac{J_3(\beta_{\max})}{J_1(\beta_{\max})}\right)^2 \approx -15$ dB. Our measured efficiency is better than this limit, and in some frequency ranges the spur suppression is better as well. Additionally, the NLTL serrodyne scheme needs only a single phase modulator compared to the four that must be integrated into an SSB modulator. Finally, the multiple path Mach-Zehnder geometry is a potential source of low-frequency drift.

The serrodyne technique can potentially be improved to offer cleaner frequency shifts and a larger tuning range. The fall time is the principle performance driver[135], and NLTLs can generate sub-picosecond transients[146]. This is more than 50 times faster than our measured fall time. Additionally, the serrodyne could be improved by adding a passive network to adjust the phase and amplitude of the existing frequency components.

Although the dual modulator push-pull configuration allows for a much larger tuning range, it does so at the expense of an additional modulator. In principle, the two serrodyne drive signals could be subtracted using a high-frequency balun and then applied to a single modulator.

It is interesting to compare these NLTL results to what could be achieved with state-of-the-art commercially available AWGs. An AWG with a 24 GS/s sampling rate (e.g. Tektronix AWG7122B) can in principle generate sawtooth waveforms with fall times as short as 35 ps which would allow for GHz serrodyne shifts with spurious sideband fraction and phase noise similar to what we observe with the NLTLs. However, an AWG-based scheme would most likely require high-power, broadband ($\sim$1 W, BW $\approx$ 10 GHz) post-amplification to achieve efficient phase modulation, and the cost and overhead associated with high frequency arbitrary synthesis is a disadvantage.

Serrodyne frequency shifting using NLTLs offers several advantages over traditional optical frequency manipulation. Unlike an AOM, the serrodyne phase modulator does not spatially shift the beam and is thus immune to temperature induced alignment changes. An NLTL-driven phase modulator offers a greater than three octave dynamic tuning range, and using a push-pull configuration we have demonstrated continuous tuning from -1.6 GHz to $+1.6$ GHz. This is more flexible than previous frequency shifting techniques in this band.

A similar implementation of the serrodyne technique recently appeared in [147]. Many thanks to the Fejer group for lending us the Infiniium oscilloscope.

Bibliography

[1] Mark Kasevich and Steven Chu. Atomic interferometry using stimulated raman transitions. *Phys. Rev. Lett.*, 67(2):181–184, Jul 1991.

[2] P. Berman, editor. *Atom Interferometry.* Academic Press, San Diego, 1997.

[3] David W. Keith, Christopher R. Ekstrom, Quentin A. Turchette, and David E. Pritchard. An interferometer for atoms. *Phys. Rev. Lett.*, 66(21):2693–2696, May 1991.

[4] Pippa Storey and Claude Cohen-Tannoudji. The Feynman path integral approach to atomic interferometry. a tutorial. *J. Phys. II France*, 4(11):1999–2027, November 1994.

[5] K. Bongs, R. Launay, and M. A. Kasevich. High-order inertial phase shifts for time-domain atom interferometers. *Applied Physics B: Lasers and Optics*, 84:599–602, September 2006.

[6] B. Dubetsky and M. A. Kasevich. Atom interferometer as a selective sensor of rotation or gravity. *Phys. Rev. A*, 74(2):023615, Aug 2006.

[7] C. J. Bordé. Quantum Theory of Atom-Wave Beam Splitters and Application to Multidimensional Atomic Gravito-Inertial Sensors. *General Relativity and Gravitation*, 36:475–502, March 2004.

[8] Jae-Hoo Gweon and Jeong-Ryeol Choi. Propagator and geometric phase of a general time-dependent harmonic oscillator. *J. Korean Phys. Soc.*, 42(3), 2003.

[9] C. Antoine. Matter wave beam splitters in gravito-inertial and trapping potentials: generalized ttt scheme for atom interferometry. *Applied Physics B: Lasers and Optics*, 84:585–597, September 2006.

[10] M. Jansen. *Atom interferometry with cold metastable helium.* PhD thesis, Technische Universiteit Eindhoven, 2007.

[11] Kathryn Moler, David S. Weiss, Mark Kasevich, and Steven Chu. Theoretical analysis of velocity-selective Raman transitions. *Phys. Rev. A*, 45(1):342–348, Jan 1992.

[12] A. F. Bernhardt and B. W. Shore. Coherent atomic deflection by resonant standing waves. *Phys. Rev. A*, 23(3):1290–1301, March 1981.

[13] L. Allen and J.H. Eberly. *Optical Resonance and Two-Level Atoms*. Dover Publications, Inc., New York, 1987.

[14] D. S. Durfee, Y. K. Shaham, and M. A. Kasevich. Long-term stability of an area-reversible atom-interferometer sagnac gyroscope. *Phys. Rev. Lett.*, 97(24):240801, Dec 2006.

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

[16] J. B. Fixler, G. T. Foster, J. M. McGuirk, and M. A. Kasevich. Atom Interferometer Measurement of the Newtonian Constant of Gravity. *Science*, 315(5808):74–77, 2007.

[17] J. B. Fixler. *Atom Interferometer-Based Gravity Gradiometer Measurements*. PhD thesis, Yale University, 2003.

[18] A. Bertoldi, et al. Atom interferometry gravity-gradiometer for the determination of the newtonian gravitational constant g. *Eur. Phys. J. D*, 40(2):271–279, 2006.

[19] G. Biedermann. *Gravity Tests, Differential Accelerometry and Interleaved Clocks with Cold Atom Interferometers*. PhD thesis, Stanford University, 2007.

[20] E. Fischbach, G. T. Gillies, D. E. Krause, J. G. Schwan, and C. Talmadge. Non-Newtonian Gravity and New Weak Forces: an Index of Measurements and Theory. *Metrologia*, 29:213–260, January 1992.

[21] Clifford M. Will. The confrontation between general relativity and experiment. *Living Rev. Rel.*, 9:3, 2005.

[22] James G. Williams, Slava G. Turyshev, and Dale H. Boggs. Progress in lunar laser ranging tests of relativistic gravity. *Phys. Rev. Lett.*, 93(26):261101, Dec 2004.

[23] S. Schlamminger, K.-Y. Choi, T. A. Wagner, J. H. Gundlach, and E. G. Adelberger. Test of the equivalence principle using a rotating torsion balance. *Phys. Rev. Lett.*, 100(4):041101, Jan 2008.

[24] J. Mester, R. Torii, P. Worden, N. Lockerbie, S. Vitale, and C. W. F. Everitt. The STEP mission: principles and baseline design. *Classical and Quantum Gravity*, 18:2475–2486, July 2001.

[25] P. Touboul, B. Foulon, L. Lafargue, and G. Metris. The microscope mission. *Acta Astronautica*, 50:433–443, April 2002.

[26] Wolfgang Petrich, Michael H. Anderson, Jason R. Ensher, and Eric A. Cornell. Stable, tightly confining magnetic trap for evaporative cooling of neutral atoms. *Phys. Rev. Lett.*, 74(17):3352–3355, Apr 1995.

[27] J. H. Denschlag, J. E. Simsarian, H. Häffner, C. McKenzie, A. Browaeys, D. Cho, K. Helmerson, S. L. Rolston, and W. D. Phillips. A Bose-Einstein condensate in an optical lattice. *Journal of Physics B: Atomic, Molecular and Optical Physics*, 35(14):3095, 2002.

[28] Holger Müller, Sheng-wey Chiow, Sven Herrmann, and Steven Chu. Atom interferometers with scalable enclosed area. *Phys. Rev. Lett.*, 102(24):240403, Jun 2009.

[29] J. M. McGuirk, M. J. Snadden, and M. A. Kasevich. Large area light-pulse atom interferometry. *Phys. Rev. Lett.*, 85(21):4498–4501, Nov 2000.

[30] Martin Weitz, Brenton C. Young, and Steven Chu. Atomic interferometer based on adiabatic population transfer. *Phys. Rev. Lett.*, 73(19):2563–2566, Nov 1994.

[31] G. Santarelli, Ph. Laurent, P. Lemonde, A. Clairon, A. G. Mann, S. Chang, A. N. Luiten, and C. Salomon. Quantum projection noise in an atomic fountain: A high stability cesium frequency standard. *Phys. Rev. Lett.*, 82(23):4619–4622, Jun 1999.

[32] J. M. Obrecht, R. J. Wild, M. Antezza, L. P. Pitaevskii, S. Stringari, and E. A. Cornell. Measurement of the temperature dependence of the Casimir-Polder force. *Phys. Rev. Lett.*, 98(6):063201, Feb 2007.

[33] J. M. McGuirk, D. M. Harber, J. M. Obrecht, and E. A. Cornell. Alkali-metal adsorbate polarization on conducting and insulating surfaces probed with bose-einstein condensates. *Phys. Rev. A*, 69(6):062905, Jun 2004.

[34] Savas Dimopoulos, Peter W. Graham, Jason M. Hogan, and Mark A. Kasevich. General relativistic effects in atom interferometry. *Phys. Rev. D*, 78(4):042003, Aug 2008.

[35] Marco Pisani and Milena Astrua. Angle amplification for nanoradian measurements. *Appl. Opt.*, 45(8):1725–1729, 2006.

[36] J. H. Müller, O. Morsch, D. Ciampini, M. Anderlini, R. Mannella, and E. Arimondo. Atomic micromotion and geometric forces in a triaxial magnetic trap. *Phys. Rev. Lett.*, 85(21):4454–4457, November 2000.

[37] I. Bloch, M. Greiner, O. Mandel, T. W. Hänsch, and T. Esslinger. Sympathetic cooling of ^{85}Rb and ^{87}Rb. *Phys. Rev. A*, 64(2):021402, August 2001.

[38] T. Bitter, et al. A large volume magnetic shielding system for the ILL neutron-antineutron oscillation experiment. *Nuclear Instruments and Methods in Physics Research A*, 309:521–529, November 1991.

[39] T. Kurosu, Y. Fukuyama, K. Abe, and Y. Koga. Measurement of a Weak Magnetic Field Using Cold Atoms. *Japanese Journal of Applied Physics*, 41:L586–L588, May 2002.

[40] Asimina Arvanitaki, Savas Dimopoulos, Andrew A. Geraci, Jason Hogan, and Mark Kasevich. How to test atom and neutron neutrality with atom interferometry. *Phys. Rev. Lett.*, 100(12):120407, Mar 2008.

[41] Savas Dimopoulos, Peter W. Graham, Jason M. Hogan, and Mark A. Kasevich. Testing general relativity with atom interferometry. *Phys. Rev. Lett.*, 98(11):111102, Mar 2007.

[42] S. Dimopoulos, P. W. Graham, J. M. Hogan, M. A. Kasevich, and S. Rajendran. Gravitational wave detection with atom interferometry. *Physics Letters B*, 678:37–40, July 2009.

[43] Savas Dimopoulos, Peter W. Graham, Jason M. Hogan, Mark A. Kasevich, and Surjeet Rajendran. Atomic gravitational wave interferometric sensor. *Phys. Rev. D*, 78(12):122002, Dec 2008.

[44] B F Schutz. Gravitational wave astronomy. *Classical and Quantum Gravity*, 16(12A):A131, 1999.

[45] C. Cutler and K. S. Thorne. An Overview of Gravitational-Wave Sources. *arXiv:gr-qc/0204090*, April 2002.

[46] Michele Maggiore. Gravitational wave experiments and early universe cosmology. *Physics Reports*, 331(6):283 – 367, 2000.

[47] M. Hohensee, S.-Y. Lan, R. Houtz, C. Chan, B. Estey, G. Kim, P.-C. Kuan, and H. Mueller. Sources and technology for an atomic gravitational wave interferometric sensor. *arXiv:1001.4821 [gr-qc]*, January 2010.

[48] B. F. Schutz. Determining the Hubble constant from gravitational wave observations. *Nature*, 323:310, September 1986.

[49] Juan García-Bellido, Daniel G. Figueroa, and Alfonso Sastre. Gravitational wave background from reheating after hybrid inflation. *Phys. Rev. D*, 77(4):043517, Feb 2008.

[50] Richard Easther, John T. Giblin, and Eugene A. Lim. Gravitational wave production at the end of inflation. *Phys. Rev. Lett.*, 99(22):221301, Nov 2007.

[51] Jean-François Dufaux, Amanda Bergman, Gary Felder, Lev Kofman, and Jean-Philippe Uzan. Theory and numerics of gravitational waves from preheating after inflation. *Phys. Rev. D*, 76(12):123517, Dec 2007.

[52] Marc Kamionkowski, Arthur Kosowsky, and Michael S. Turner. Gravitational radiation from first-order phase transitions. *Phys. Rev. D*, 49(6):2837–2851, Mar 1994.

[53] Chiara Caprini, Ruth Durrer, and Géraldine Servant. Gravitational wave generation from bubble collisions in first-order phase transitions: An analytic approach. *Phys. Rev. D*, 77(12):124015, Jun 2008.

[54] Christophe Grojean and Geraldine Servant. Gravitational Waves from Phase Transitions at the Electroweak Scale and Beyond. *Phys. Rev. D*, 75:043507, 2007.

[55] R. Apreda, M. Maggiore, A. Nicolis, and A. Riotto. Gravitational waves from electroweak phase transitions. *Nuclear Physics B*, 631:342–368, June 2002.

[56] Lisa Randall and Graldine Servant. Gravitational waves from warped spacetime. *Journal of High Energy Physics*, 2007(05):054, 2007.

[57] J. Polchinski. Cosmic String Loops and Gravitational Radiation. *arXiv:0707.0888 [astro-ph]*, July 2007.

[58] Matthew R. DePies and Craig J. Hogan. Stochastic gravitational wave background from light cosmic strings. *Phys. Rev. D*, 75(12):125006, Jun 2007.

[59] Craig J. Hogan. Cosmological gravitational wave backgrounds. *arXiv:astro-ph/9809364*, 1998.

[60] Craig J. Hogan. Gravitational waves from mesoscopic dynamics of the extra dimensions. *Phys. Rev. Lett.*, 85:2044–2047, 2000.

[61] Craig J. Hogan. Scales of the extra dimensions and their gravitational wave backgrounds. *Phys. Rev. D*, 62:121302, 2000.

[62] M. Gasperini and G. Veneziano. The pre-big bang scenario in string cosmology. *Phys. Rept.*, 373:1–212, January 2003.

[63] Michael S. Turner and Frank Wilczek. Relic gravitational waves and extended inflation. *Phys. Rev. Lett.*, 65(25):3080–3083, Dec 1990.

[64] A. Bresson, Y. Bidel, P. Bouyer, B. Leone, E. Murphy, and P. Silvestrin. Quantum mechanics for space applications. *Applied Physics B: Lasers and Optics*, 84:545–550, 2006. 10.1007/s00340-006-2412-x.

[65] G. Amelino-Camelia, et al. Gauge: the grand unification and gravity explorer. *Experimental Astronomy*, 23:549–572, 2009. 10.1007/s10686-008-9086-9.

[66] P. Wolf, et al. Quantum physics exploring gravity in the outer solar system: the sagas project. *Experimental Astronomy*, 23:651–687, 2009. 10.1007/s10686-008-9118-5.

[67] L. Cacciapuoti, et al. Atomic clock ensemble in space: Scientific objectives and mission status. *Nuclear Physics B - Proceedings Supplements*, 166:303 – 306, 2007. Proceedings of the Third International Conference on Particle and Fundamental Physics in Space, Proceedings of the Third International Conference on Particle and Fundamental Physics in Space.

[68] L. Cacciapuoti and Ch. Salomon. Space clocks and fundamental tests: The aces experiment. *The European Physical Journal - Special Topics*, 172:57–68, 2009. 10.1140/epjst/e2009-01041-7.

[69] Alexander D. Cronin, Jörg Schmiedmayer, and David E. Pritchard. Optics and interferometry with atoms and molecules. *Rev. Mod. Phys.*, 81(3):1051–1129, Jul 2009.

[70] J. M. Hogan, D. M. S. Johnson, and M. A. Kasevich. Light-pulse atom interferometry. In E. Arimondo, W. Ertmer, and W. P. Schleich, editors, *Proceedings of the International School of Physics "Enrico Fermi" on Atom Optics and Space Physics*, pages 411–447, Amsterdam, 2009. IOS Press.

[71] Brahim Lamine, Rémy Hervé, Astrid Lambrecht, and Serge Reynaud. Ultimate decoherence border for matter-wave interferometry. *Phys. Rev. Lett.*, 96(5):050405, Feb 2006.

[72] B. Lamine, M.T. Jaekel, and S. Reynaud. Gravitational decoherence of atomic interferometers. *The European Physical Journal D-Atomic, Molecular, Optical and Plasma Physics*, 20(2):165–176, 2002.

[73] P. Delva, M.C. Angonin, and P. Tourrenc. A comparison between matter wave and light wave interferometers for the detection of gravitational waves. *Physics Letters A*, 357(4-5):249–254, 2006.

[74] GM Tino and F. Vetrano. Is it possible to detect gravitational waves with atom interferometers? *Classical and Quantum Gravity*, 24:2167, 2007.

[75] S. Foffa, A. Gasparini, M. Papucci, and R. Sturani. Sensitivity of a small matter-wave interferometer to gravitational waves. *Physical Review D*, 73(2):022001, 2006.

[76] E. DAmbrosio, L. Maleki, and N. Yu. Translational invariance of gravitational wave atom interferometers. *Physical Review D*, 76(12):122001, 2007.

[77] R.Y. Chiao and A.D. Speliotopoulos. Towards MIGO, the matter-wave interferometric gravitational-wave observatory, and the intersection of quantum mechanics with general relativity. *Journal of Modern Optics, 51*, 6(7):861–899.

[78] A. Roura, D.R. Brill, BL Hu, C.W. Misner, and W.D. Phillips. Gravitational wave detectors based on matter wave interferometers (MIGO) are no better than laser interferometers (LIGO). *Physical Review D*, 73(8):084018, 2006.

[79] D. Lorek, C. Lämmerzahl, and A. Wicht. Orientational atom interferometers sensitive to gravitational waves. *Physical Review A*, 81(2):023621, 2010.

[80] David M. Giltner, Roger W. McGowan, and Siu Au Lee. Atom interferometer based on bragg scattering from standing light waves. *Phys. Rev. Lett.*, 75(14):2638–2641, Oct 1995.

[81] Ernst M. Rasel, Markus K. Oberthaler, Herman Batelaan, Jörg Schmiedmayer, and Anton Zeilinger. Atom wave interferometry with diffraction gratings of light. *Phys. Rev. Lett.*, 75(14):2633–2637, Oct 1995.

[82] T. Lévèque, A. Gauguet, F. Michaud, F. Pereira Dos Santos, and A. Landragin. Enhancing the area of a raman atom interferometer using a versatile double-diffraction technique. *Phys. Rev. Lett.*, 103(8):080405, Aug 2009.

[83] K Kawabe and the LIGO Collaboration. Status of ligo. *Journal of Physics: Conference Series*, 120(3):032003, 2008.

[84] J R Smith and the LIGO Collaboration. The path to the enhanced and advanced ligo gravitational-wave detectors. *Classical and Quantum Gravity*, 26(11):114013, 2009.

[85] Shane L. Larson, William A. Hiscock, and Ronald W. Hellings. Sensitivity curves for spaceborne gravitational wave interferometers. *Phys. Rev. D*, 62(6):062001, Aug 2000.

[86] Nico Sneeuw and Hanspeter Schaub. Satellite clusters for future gravity field missions. In Fernando Sans, Christopher Jekeli, Luisa Bastos, and Joana Fernandes, editors, *Gravity, Geoid and Space Missions*, volume 129 of *International Association of Geodesy Symposia*, pages 12–17. Springer Berlin Heidelberg, 2005. 10.1007/3-540-26932-0_3.

[87] Nelson Christensen. Measuring the stochastic gravitational-radiation background with laser-interferometric antennas. *Phys. Rev. D*, 46(12):5250–5266, Dec 1992.

[88] B. Abbott, et al. Searching for a stochastic background of gravitational waves with the laser interferometer gravitational-wave observatory. *The Astrophysical Journal*, 659(2):918, 2007.

[89] Bruce Allen. The stochastic gravity-wave background: Sources and detection. *arXiv:gr-qc/9604033*, 1996.

[90] Tristan L. Smith, Elena Pierpaoli, and Marc Kamionkowski. New cosmic microwave background constraint to primordial gravitational waves. *Phys. Rev. Lett.*, 97(2):021301, Jul 2006.

[91] Bruce Allen and Scott Koranda. Cbr anisotropy from primordial gravitational waves in inflationary cosmologies. *Phys. Rev. D*, 50(6):3713–3737, Sep 1994.

[92] A. J. Farmer and E. S. Phinney. The gravitational wave background from cosmological compact binaries. *Mon. Not. Roy. Astron. Soc.*, 346:1197–1214, December 2003.

[93] H.J. Metcalf and P. Straten. *Laser Cooling and Trapping*. Springer, New York, 1999.

[94] T. van Zoest, et al. Bose-Einstein Condensation in Microgravity. *Science*, 328(5985):1540–1543, 2010.

[95] G. Stern, et al. Light-pulse atom interferometry in microgravity. *Eur. Phys. J. D*, 53(3):353–357, 2009.

[96] R. S. Conroy, Y. Xiao, M. Vengalattore, W. Rooijakkers, and M. Prentiss. Compact, robust source of cold atoms for efficient loading of a magnetic guide. *Optics Communications*, 226(1-6):259 – 266, 2003.

[97] Mark Kasevich, David S. Weiss, Erling Riis, Kathryn Moler, Steven Kasapi, and Steven Chu. Atomic velocity selection using stimulated raman transitions. *Phys. Rev. Lett.*, 66(18):2297–2300, May 1991.

[98] Hubert Ammann and Nelson Christensen. Delta kick cooling: A new method for cooling atoms. *Phys. Rev. Lett.*, 78(11):2088–2091, Mar 1997.

[99] Jean-Sébastien Bernier, Corinna Kollath, Antoine Georges, Lorenzo De Leo, Fabrice Gerbier, Christophe Salomon, and Michael Köhl. Cooling fermionic atoms in optical lattices by shaping the confinement. *Phys. Rev. A*, 79(6):061601, Jun 2009.

[100] J.-F. Clément, J.-P. Brantut, M. Robert-de Saint-Vincent, R. A. Nyman, A. Aspect, T. Bourdel, and P. Bouyer. All-optical runaway evaporation to bose-einstein condensation. *Phys. Rev. A*, 79(6):061406, Jun 2009.

[101] Tim Kovachy, Jason M. Hogan, David M. S. Johnson, and Mark A. Kasevich. Optical lattices as waveguides and beam splitters for atom interferometry: An analytical treatment and proposal of applications. *Phys. Rev. A*, 82(1):013638, Jul 2010.

[102] Pierre Cladé, Saïda Guellati-Khélifa, Fran çois Nez, and Fran çois Biraben. Large momentum beam splitter using bloch oscillations. *Phys. Rev. Lett.*, 102(24):240402, Jun 2009.

[103] Malo Cadoret, Estefania de Mirandes, Pierre Cladé, Saïda Guellati-Khélifa, Catherine Schwob, Fran çois Nez, Lucile Julien, and Fran çois Biraben. Combination of bloch oscillations with a ramsey-bordé interferometer: New determination of the fine structure constant. *Phys. Rev. Lett.*, 101(23):230801, Dec 2008.

[104] Laser interferometer space antenna (LISA) mission concept. LISA Project internal report number LISA-PRJ-RP-0001, May 2009.

[105] M. G. Hauser, T. Kelsall, D. Leisawitz, and J. Weiland, editors. *COBE Diffuse Infrared Background Experiment (DIRBE) Explanatory Supplement*. NASA/GSFC, Greenbelt, MD, COBE ref. pub. no. 98 A edition, 1998. available in electronic form from the NSSDC.

[106] Kenji Numata, Masaki Ando, Kazuhiro Yamamoto, Shigemi Otsuka, and Kimio Tsubono. Wide-band direct measurement of thermal fluctuations in an interferometer. *Phys. Rev. Lett.*, 91(26):260602, Dec 2003.

[107] D. A. Vallado and D. Finkelman. In *Astrodynamics Specialist Conference and Exhibit*, Honolulu, Hawaii, 2008.

[108] A. Gillespie and F. Raab. Thermally excited vibrations of the mirrors of laser interferometer gravitational-wave detectors. *Phys. Rev. D*, 52(2):577–585, Jul 1995.

[109] Yu. Levin. Internal thermal noise in the ligo test masses: A direct approach. *Phys. Rev. D*, 57(2):659–663, Jan 1998.

[110] M. Cerdonio, L. Conti, A. Heidmann, and M. Pinard. Thermoelastic effects at low temperatures and quantum limits in displacement measurements. *Phys. Rev. D*, 63(8):082003, Mar 2001.

[111] Derek J. Edinger and Alison A. Nordt. Selection of i-220h beryllium for nircam optical bench. volume 5868, page 58680X. SPIE, 2005.

[112] For example, see the 4D Technology FizCam 2000 Fizeau interferometer.

[113] P. Cheinet, B. Canuel, F. Pereira Dos Santos, A. Gauguet, F. Leduc, and A. Landragin. Measurement of the sensitivity function in time-domain atomic interferometer. *arXiv:physics/0510197*, October 2005.

[114] J. Le Gouet, J. Kim, C. Bourassin-Bouchet, M. Lours, A. Landragin, and F. Pereira Dos Santos. Wide bandwidth phase-locked diode laser with an intra-cavity electro-optic modulator. *Optics Communications*, 282(5):977 – 980, 2009.

[115] A. D. Ludlow, X. Huang, M. Notcutt, T. Zanon-Willette, S. M. Foreman, M. M. Boyd, S. Blatt, and J. Ye. Compact, thermal-noise-limited optical cavity for diode laser stabilization at 1×10^{-15}. *Opt. Lett.*, 32(6):641–643, 2007.

[116] A. E. Hedin. Extension of the MSIS thermosphere model into the middle and lower atmosphere. *J. Geophys. Res.*, 96:1159–1172, February 1991.

[117] S. Maus, S. Macmillan, S. McLean, B. Hamilton, A. Thomson, M. Nair, and C. Rollins, 2010. The US/UK World Magnetic Model for 2010-2015, NOAA Technical Report NESDIS/NGDC.

[118] P. Storey and C. Cohen-Tannoudji. The Feynman path integral approach to atomic interferometry. A tutorial. *Journal de Physique II*, 4(11):1999–2027, 1994.

[119] N. K. Pavlis, S. A. Holmes, S. C. Kenyon, and J. K. Factor. An Earth Gravitational Model to Degree 2160: EGM2008, presented at the 2008 General Assembly of the European Geosciences Union, Vienna, Austria, April 13-18, 2008.

[120] Jeffrey Park, et al. Earth's Free Oscillations Excited by the 26 December 2004 Sumatra-Andaman Earthquake. *Science*, 308(5725):1139–1144, 2005.

[121] F. A. Dahlen and J. Tromp. *Theoretical Global Seismology*. Princeton University Press, Princeton, NJ, 1998.

[122] Guy Masters, Misha Barmine, Susan Kientz, J. F. Gilbert, and John Woodhouse. the MINEOS Computational Infrastructure for Geodynamics.

[123] For example, see ABLE Engineering, Inc., http://www.aec-able.com/Booms/coilboom.html.

[124] B. D. Tapley, D. P. Chambers, S. Bettadpur, and J. C. Ries. Large scale ocean circulation from the GRACE GGM01 Geoid. *Geophys. Res. Lett.*, 30(22):220000–1, November 2003.

[125] J. L. Chen, C. R. Wilson, D. Blankenship, and B. D. Tapley. Accelerated Antarctic ice loss from satellite gravity measurements. *Nature Geoscience*, 2:859–862, December 2009.

[126] B. D. Tapley, S. Bettadpur, J. C. Ries, P. F. Thompson, and M. M. Watkins. GRACE Measurements of Mass Variability in the Earth System. *Science*, 305:503–506, July 2004.

[127] N. Olsen, H. Lühr, T. J. Sabaka, M. Mandea, M. Rother, L. Tøffner-Clausen, and S. Choi. CHAOS-a model of the Earth's magnetic field derived from CHAMP, Ørsted, and SAC-C magnetic satellite data. *Geophysical Journal International*, 166:67–75, July 2006.

[128] A. Kuvshinov and N. Olsen. A global model of mantle conductivity derived from 5 years of CHAMP, Ørsted, and SAC-C magnetic data. *Geophys. Res. Lett.*, 33:18301–+, September 2006.

[129] W. Ketterle, D. S. Durfee, and D. M. Stamper-Kurn. Making, probing and understanding bose-einstein condensates. 1999.

[130] H. J. Lewandowski, D. M. Harber, D. L. Whitaker, and E. A. Cornell. Simplied system for creating a bose-einstein condensate. *Cornell Group Website at JILA*, 2003.

[131] Tetsuya Kawanishi, Takahide Sakamoto, and Masayuki Izutsu. High-speed control of lightwave amplitude, phase, and frequency by use of electrooptic effect. *Selected Topics in Quantum Electronics, IEEE Journal of*, 13(1):79–91, Jan-Feb 2007.

[132] J. L. Hall, L. Hollberg, T. Baer, and H. G. Robinson. Optical heterodyne saturation spectroscopy. *Applied Physics Letters*, 39(9):680–682, 1981.

[133] Steven Chu, L. Hollberg, J. E. Bjorkholm, Alex Cable, and A. Ashkin. Three-dimensional viscous confinement and cooling of atoms by resonance radiation pressure. *Phys. Rev. Lett.*, 55(1):48–51, Jul 1985.

[134] K. K. Wong, R. M. De La Rue, and S. Wright. Electro-optic-waveguide frequency translator in $LiNbO_3$ fabricated by proton exchange. *Opt. Lett.*, 7(11):546–548, 1982.

[135] L.M. Johnson and III Cox, C.H. Serrodyne optical frequency translation with high sideband suppression. *Journal of Lightwave Technology*, 6(1):109–112, Jan 1988.

[136] L. Thylen, P. Sjoberg, and G.E. Lindqvist. Electro-optical serrodyne frequency translator for $\lambda = 1.3$ μm. *Optoelectronics, IEE Proceedings J*, 132(2):119–121, April 1985.

[137] C. Laskoskie, H. Hung, T. El-Wailly, and C.L. Chang. Ti-$LiNbO_3$ waveguide serrodyne modulator with ultrahigh sideband suppression for fiber optic gyroscopes. *Lightwave Technology, Journal of*, 7(4):600–606, Apr 1989.

[138] Sarper Ozharar, Sangyoun Gee, Frank Quinlan, and Peter J. Delfyett. Heterodyne serrodyne with high sideband suppression via time division multiplexing for arbitrary waveform generation. volume 5814, pages 79–83. SPIE, 2005.

[139] S. Ozharar, S. Gee, F. Quinlan, I. Ozdur, and P.J. Delfyett. Ultra-high spur-free dynamic range RF synthesis using optical homodyne serrodyne technique. *Electronics Letters*, 43(5):65 –66, March 2007.

[140] M. Shen and R.A. Minasian. Serrodyne optical frequency translation using photonics-based waveforms. *Electronics Letters*, 40(24):1545 – 1547, 2004.

[141] I.Y. Poberezhskiy, B. Bortnik, J. Chou, B. Jalali, and H.R. Fetterman. Serrodyne frequency translation of continuous optical signals using ultrawide-band electrical sawtooth waveforms. *IEEE Journal of Quantum Electronics*, 41(12):1533–1539, Dec. 2005.

[142] R. Reibel, T. Chang, M. Tian, and W. R. Babbitt. Optical linear sideband chirp compression for microwave arbitrary waveform generation. In *IEEE Microwave Photonics Conference Proceedings*, pages 197–200. IEEE, 2004.

[143] R. R. Reibel, Z. W. Barber, J. A. Fischer, M. Tian, and W. R. Babbitt. Broadband demonstrations of true-time delay using linear sideband chirped programming and optical coherent transients. *Journal of Luminescence*, 107(1-4):103 – 113, 2004.

[144] M. J. W. Rodwell, et al. Active and nonlinear wave propagation devices in ultrafast electronics and optoelectronics [and prolog]. *Proceedings of the IEEE*, 82(7):1037–1059, July 1994.

[145] S. Shimotsu, S. Oikawa, T. Saitou, N. Mitsugi, K. Kubodera, T. Kawanishi, and M. Izutsu. Single side-band modulation performance of a $LiNbO_3$ integrated modulator consisting of four-phase modulator waveguides. *Photonics Technology Letters, IEEE*, 13(4):364 –366, apr 2001.

[146] D. W. van der Weide. Delta-doped schottky diode nonlinear transmission lines for 480-fs, 3.5-v transients. *Applied Physics Letters*, 65(7):881–883, 1994.

[147] Rachel Houtz, Cheong Chan, and Holger Müller. Wideband, efficient optical serrodyne frequency shifting with a phase modulator and a nonlinear transmission line. *Opt. Express*, 17(21):19235–19240, 2009.

[148] Masatoshi Saruwatari. High-speed all-optical technologies for photonics. In Nadir Dagli, editor, *High-Speed Photonic Devices*, pages 217–248. Taylor & Francis, 2007.

[149] Lothar Moeller, Yikai Su, Chongjin Xie, Jurgen Gripp, Xiang Liu, and Roland Ryf. 640-Gb/s OTDM RZ-DQPSK signal enabling 2.4-Bit/s/Hz spectral efficiency and its detection with an EAM-based receiver. In *Asia Optical Fiber Communication and Optoelectronic Exposition and Conference*, page SaA2. Optical Society of America, 2008.

[150] Hsu-Feng Chou, John E. Bowers, and Daniel J. Blumenthal. Compact 160-Gb/s add-drop multiplexer with a 40-Gb/s base rate using electroabsorption modulators. *IEEE Photon. Technol. Lett.*, 16:1564–1566, 2004.

[151] T. Ohara, et al. 160-Gb/s OTDM Transmission Using Integrated All-Optical MUX/DEMUX With All-Channel Modulation and Demultiplexing. *IEEE Photonics Technology Letters*, 16:650–652, February 2004.

[152] I. Glesk and P.R. Prucnal. 250-Gb/s self-clocked optical TDM with a polarization-multiplexed clock. *Fiber and Integrated Optics*, 14:71–82, 1995.

[153] Andrei I. Siahlo, Anders T. Clausen, Leif K. Oxenløwe, Jorge Seoane, and Palle Jeppesen. 640 Gb/s OTDM transmission and demultiplexing using a NOLM with commercially available highly nonlinear fiber. In *Conference on Lasers and Electro-Optics/Quantum Electronics and Laser Science and Photonic Applications Systems Technologies*, page CTuO1. Optical Society of America, 2005.

[154] C. J. Madden, M. J. W. Rodwell, R. A. Marsland, D. M. Bloom, and Y. C. Pao. Generation of 3.5-ps fall-time shock waves on a monolithic GaAs nonlinear transmission line. *IEEE Electron Device Letters*, 9:303–305, June 1988.

[155] D. M. S. Johnson, J. M. Hogan, S. w. Chiow, and M. A. Kasevich. Broadband optical serrodyne frequency shifting. *Opt. Lett.*, 35(5):745–747, 2010.

[156] J. Breitbarth. *Design and Characterization of Low Phase Noise Microwave Circuits*. PhD thesis, University of Colorado, 2006.

[157] Picosecond Pulse Labs. A new breed of comb generators featuring low phase noise and low input power. *Microwave Journal*, 49:5, 2006.

[158] M. R. et al. Tan. A 12 psec gaas double heterostructure step recovery diode. *HP Labs Technical Report*, HPL-91-187, 1991.

[159] M. Birk, L.M. Lunardi, A.H. Gnauck, H. Schumacher, and D. Behammer. 10-Gbit/s RZ pulses using an all-silicon nonlinear transmission line integrated circuit. In *Optical Fiber Communication Conference, 2000*, volume 1, pages 248–250 vol.1, 2000.

[160] M. Lee, H. E. Katz, C. Erben, D. M. Gill, P. Gopalan, J. D. Heber, and D. J. McGee. Broadband Modulation of Light by Using an Electro-Optic Polymer. *Science*, 298:1401–1404, November 2002.

[161] M. Lee. Dielectric constant and loss tangent in $LiNbO_3$ crystals from 90 to 147 GHz. *Applied Physics Letters*, 79:1342, August 2001.

[162] J. Shin, S. Wu, and N. Dagli. 35-GHz Bandwidth, 5-V-cm Drive Voltage, Bulk GaAs Substrate Removed Electrooptic Modulators. *IEEE Photonics Technology Letters*, 19:1362–1364, September 2007.

[163] A. Kuver, J. M. Heaton, A. Miller, G. Murdoch, Y. Zhou, P. Jiang, and D. R. Wight. Development of GaAs electro-optic modulators with improved linearity. In *Proceedings of the 2nd EMRS DTC Technical Conference-Edinburgh*, 2005.

[164] Y. Cui and P. Berini. Modeling and Design of GaAs Traveling-Wave Electrooptic Modulators Based on the Planar Microstrip Structure. *Journal of Lightwave Technology*, 24:2368, June 2006.

[165] A. Yariv. *Optical Electronics in Modern Communications*. Oxford University Press, New York, 1997.