Development of new technologies for precision torsion-balance experiments

Matthew D. Turner

A dissertation submitted in partial fulfullment of the requirements for the degree of Doctor of Philosophy

University of Washington

2017

Reading Committee:

Jens H. Gundlach, Chair

Blayne R. Heckel

Boris B. Blinov

Program authorized to offer degree:

Department of Physics

University of Washington

Abstract

Development of new technologies for precision torsion-balance experiments

Matthew D. Turner

Chair of the Supervisory Committee:

Professor Jens H. Gundlach

Department of Physics

Torsion balances were first used for precision measurements of physics centuries ago and continue to be used to probe physical forces and test for new physics. Improving beyond current limits requires better sensitivity and better ability to monitor or reduce unwanted systematic effects. This dissertation describes two technologies developed to aid in these improvements. The first is an interferometric quasi-autocollimator—an optical readout device inspired by quantum weak-value amplification and capable of an angular sensitivity of 10 picoradians per root hertz. The second is a gravity gradiometer torsion balance with a mass quadrupole that can be changed *in situ*. This balance uses a wirelessly powered non-magnetic rotary actuator, also developed as part of this work, which will enable additional approaches to systematics and torsion-balance science.

All kingdoms have a law given; And there are many kingdoms; for there is no space in the which there is no kingdom; and there is no kingdom in which there is no space, either a greater or a lesser kingdom. And unto every kingdom is given a law; and unto every law there are certain bounds also and conditions.

. . .

And again, verily I say unto you, he hath given a law unto all things, by which they move in their times and their seasons; And their courses are fixed, even the courses of the heavens and the earth, which comprehend the earth and all the planets.

. . .

Unto what shall I liken these kingdoms, that ye may understand? Behold, all these are kingdoms, and any man who hath seen any or the least of these hath seen God moving in his majesty and power.

–Doctrine and Covenants 88:36-38,42-43,46-47

Acknowledgements

There are many people without whom this work would not have been possible or nearly as enjoyable. Charlie Hagedorn in particular has been a great officemate, mentor, *de facto* advisor, and friend. Krishna Venkateswara and Stephan Schlamminger have also been the best officemates one could hope for, in both professional and general terms. I've learned much about science and life from all of them. Thanks to Jens Gundlach as my advisor for enabling and encouraging me in both my scientific and creative pursuits. I also thank the rest of the Eöt-Wash group for their mentorship and support and the rest of the CENPA personnel for contributing to and maintaining such a fantastic atmosphere and facility to work and learn in.

Last I'd like to thank my parents for their support from afar; my children Callan, Marian, and Max for constantly reminding me of the wonder of science and giving me the chance to play every day; and, most importantly, my wife Michelle for her enduring love, patience, encouragement, and support. She's simply the best.

Contents

Chapter 1

Introduction

Both projects described in this thesis were motivated by considering ways to monitor the leading systematic uncertainty in our equivalence principle tests: variations in the local gravitational gradient. To explain more, it will be helpful to explain the nature of the gravitational gradient and methods of measurement.

1.1 Gravitational gradients

The gravitational gradient is a tensor field equivalent to the second spatial derivative of the gravitational potential, or the gradient of each component of the gravitational field. Put in simpler terms, the gravitational gradient is a measure of how much "down" or the acceleration due to gravity changes when you move from one point in space to another nearby point. Expressed mathematically, the gravity gradient

$$\mathbf{\Gamma} = \nabla \otimes \mathbf{g} = -\nabla \otimes \nabla \Phi,$$

where $\mathbf{g}$ is the gravitational field and Φ the gravitational potential. Expressed in Cartesian coordinates, the gravity gradient might naively be expected to have nine components, as

$$\mathbf{\Gamma} = -\begin{pmatrix} \frac{\partial^2}{\partial x \partial x} & \frac{\partial^2}{\partial x \partial y} & \frac{\partial^2}{\partial x \partial z} \\ \frac{\partial^2}{\partial y \partial x} & \frac{\partial^2}{\partial y \partial y} & \frac{\partial^2}{\partial y \partial z} \\ \frac{\partial^2}{\partial z \partial x} & \frac{\partial^2}{\partial z \partial y} & \frac{\partial^2}{\partial z \partial z} \end{pmatrix} \Phi = \begin{pmatrix} \Gamma_{xx} & \Gamma_{xy} & \Gamma_{xz} \\ \Gamma_{yx} & \Gamma_{yy} & \Gamma_{yz} \\ \Gamma_{zx} & \Gamma_{zy} & \Gamma_{zz} \end{pmatrix}.$$

However, because gravity is a conservative field that obeys Gauss's law, the tensor is symmetric and, in free space, traceless, resulting in five independent components:

$$\mathbf{\Gamma} = \begin{pmatrix} \Gamma_{xx} & \Gamma_{xy} & \Gamma_{xz} \\ \Gamma_{xy} & \Gamma_{yy} & \Gamma_{yz} \\ \Gamma_{xz} & \Gamma_{yz} & \Gamma_{zz} \end{pmatrix}$$

with $\Gamma_{xx} + \Gamma_{yy} + \Gamma_{zz} = 0$. The fact that the tensor has five independent components can also be deduced by identifying the gradient as a quadrupole field, in which case it can be described in terms of the spherical harmonics Γ_{20}, Γ_{21}, and Γ_{22}, with Γ_{21} and Γ_{22} being complex valued. We will mainly refer to the Cartesian components, but conversion can be done using these relations [1]:

$$\begin{aligned}
\Gamma_{zz} &= 2\sqrt{\frac{4\pi}{5}}\Gamma_{20} \\
\Gamma_{yz} &= 3\sqrt{\frac{8\pi}{15}}\Im\{\Gamma_{21}\} \\
\Gamma_{xz} &= -3\sqrt{\frac{8\pi}{15}}\Re\{\Gamma_{21}\} \\
\Gamma_{xy} &= -6\sqrt{\frac{2\pi}{15}}\Im\{\Gamma_{22}\} \\
\Gamma_{xx} &= 6\sqrt{\frac{2\pi}{15}}\Re\{\Gamma_{22}\} - \sqrt{\frac{4\pi}{5}}\Gamma_{20} \\
\Gamma_{yy} &= -6\sqrt{\frac{2\pi}{15}}\Re\{\Gamma_{22}\} - \sqrt{\frac{4\pi}{5}}\Gamma_{20}
\end{aligned}$$

The maximal diagonal-term gradient Γ_{ii} and cross-term gradient Γ_{ij} $(i \neq j)$ from a point mass m at distance r are

$$\Gamma_{ii} = \frac{2Gm}{r^3} \tag{1.1}$$

and

$$\Gamma_{ij} = \frac{3Gm}{2r^3}. \tag{1.2}$$

The gravitational gradient is a spatial derivative of an acceleration, and so has dimensions of acceleration per distance, which is equivalently units of per time squared or of frequency squared. Rather than using units of $1/s^2$ or Hz^2, and since surveyed values are commonly on the order of $100 \times 10^{-9}\,Hz^2$ [2], a unit called the eotvos (Eo here, also abbreviated E) has been defined to be equal to $10^{-9}\,Hz^2$. For reference, the vertical gradient

	Δa_N (fm/s^2)	Δa_W (fm/s^2)
residual gravity gradients	1.6 ± 0.2	0.3 ± 1.7
tilt	1.2 ± 0.6	-0.2 ± 0.7
magnetic	0.0 ± 0.3	0.0 ± 0.3
temperature gradients	0.0 ± 1.7	0.0 ± 1.7

Table 1.1: Systematics and uncertainties for northward and westward differential accelerations for a Be-Ti pendulum. [4].

component Γ_{zz} on the Earth's surface is about 3 kEo, and the maximal cross-term gradient component (Γ_{xy}, Γ_{xz}, or Γ_{yz}) from a mass of 100 kg at a distance of 1 m is 10.0 Eo.

The gravitational gradient is the sum of the gradient from all surrounding matter, and thus is sensitive to mass changes in the local environment. Changes in the water table or the level or nearby bodies of water are a natural source of gradient variations. Due to the masses involved and the dependence on the cube of the distance, cultural noise is less likely to be significant. A 20-ton bus would have to be brought within 12 meters to produce a gradient change of 1 Eo, but nearby excavation or building projects have the potential to result in larger changes. For example, a construction project near our laboratory has involved the excavation of $\approx 10^5$ m^3 of soil at a distance of about 120 m from the lab [3], giving an expected gradient change of about 10 Eo.

1.2 Motivation

Our group's NewWash experiment is at the limit of equivalence-principle violation tests conducted to date and has the best limits on potential new Yukawa forces for distances up to 100 Mm. The sensitivity of this experiment has been limited by statistical noise, but systematic uncertainties are close to the same magnitude, with variations in the local gravity gradient being the leading term as shown in Table 1.1. For this and earlier equivalence principle experiments we have put estimated bounds on the variation in the gravitational gradient by replacing the science pendulums with gradiometer pendulums, which are designed to have large gravitational moments, and collected runs of data before and after science runs. The result from one such run, taken during a week with significant rain, is shown in Figure 1.1.

Although the repurposing of our torsion balance setups as gradiometers has allowed for very precise measurement of the local gravitational gradient, it does not allow for simultaneous monitoring of the gradient during experimental runs. A separate standalone system is needed to do this, and this need was the initial motivation

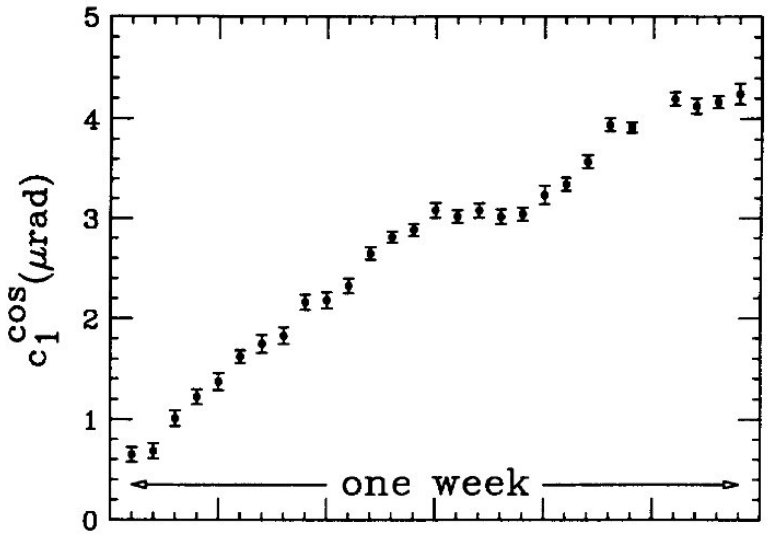

Figure 1.1: A measurement of the Γ_{21} gradient during a rainy week, from [5].

for the two parts of this work.

1.3 Gradiometry

The standard gravitational gradient unit is named after Loránd Eötvös, who pioneered the use of torsion balances in measuring gravitational gradients and testing the equivalence principle. Describing his first gradiometer, Eötvös said [6]:

> It was just a simple, straight stick that I used as instrument, specially loaded at both ends, enclosed into a metal sheath to protect it from the wind and temperature changes. Upon this stick every single mass, be it near or far, exerts a directing force; but the wire upon which it hangs resists, and whilst resisting it twists, with the degree of this twist showing us the exact magnitude of the forces acting upon the stick. This is a Coulomb balance, and that is all there is to it. It is simple, like the flute of Hamlet, you only have to know how to play on it, and just like the musician who can delight you with splendid variations, the physicist can, on this balance, with no less delight determine the finest variations of gravity. This way we can peer into such depth of the crust of the Earth, that neither our eyes, nor our longest drills could reach.

As explained, the gradiometer consisted of a rod with weights at each end, which hung from a wire within a protective casing. A horizontal gravity gradient (Γ_{xy} or $\Gamma_{xx} - \Gamma_{yy}$) produces a torque on such an arrangement, and

the magnitude of the torque could be read out optically using a telescope to view the reflection of a fixed scale seen in a mirror on the pendulum. Later verisons had one weight suspended from a wire below the pendulum rod, which enabled the gradiometer to measure additional gradient components: Γ_{xz} and Γ_{yz}. These devices demonstrated the utility of gravity gradiometry in geological surveys and were employed for oil and mineral prospecting until supplanted by gravimeters and modern gradiometers.

In the past fifty years, new gradiometers have been developed for both prospecting and submarine navigation [7]. (They are valuable for submarine navigation because they are completely passive—no signals are generated that might be detected by an enemy.) These gradiometers consist of sets of paired accelerometers. Gradient components are calculated by taking the difference between each pair and dividing by the distance between them, so each pair must be carefully matched so that common-mode signals cancel out properly. Besides being prohibitively expensive, commercial systems are not designed for long-term gradient monitoring and so are not suitable for our application. Since our expertise is in torsion balances we chose to pursue a torsion-balance based solution.

When measuring a varying value over time, naturally occuring instrument drift can be indistinguishable from a drift in the value of interest. When the measurement is not expected to vary, instrument drift is simply a source of noise. Modulation provides a way around this—instrument drift will stay at low frequency, while variations in the measured value get upconverted to a higher frequency with the modulation. Our group has employed a few methods to produce a modulated signal in torsion balance measurements. For experiments which measure the torque on a torsion pendulum due to a source mass, we have either rotated a source mass [8, 9] or repositioned it linearly [10]. For our experiments that use natural objects, such as landscape features or astronomical bodies, as source masses, controlling the position of the source mass is not feasible, so we instead modulate its effect on the torsion pendulum by rotating the entire torsion balance [4, 5].

We initially investigated the feasibility of rotating a gradiometer torsion balance faster than the swing frequency of the pendulum. Torsion balance readout at frequencies above the torsional oscillation frequency require higher readout sensitivity due to the response drop-off above resonance. This drove an angular readout requirement of $10\,\mathrm{prad}/\sqrt{\mathrm{Hz}}$ at 10 Hz. Although we determined that this approach was not feasible due to instability of a rotating pendulum when surrounded by rotating objects [11], we were able to meet this readout requirement with the development of a new device described in Part I. Part II describes a second approach with a new method of modulation: repositioning components of the torsion pendulum *in situ*.

Part I

An interferometric quasi-autocollimator

Chapter 2

Angular sensing

The ability to measure small angular deflections quickly, precisely, and accurately is important in many fields of experimental physics. In the Eöt-Wash experimental gravity group at the University of Washington, our torsion balance experiments depend upon our ability to measure minuscule angular deflections [4, 12].

The simplest form of light-based angular readout might be an optical lever, as shown in Figure 2.1 [13]. Collimated light is reflected off of a target mirror and onto a position-sensitive detector (PSD). Gain is proportional to the distance between the target and the detector ($\Delta x = 2\theta l$). The optical lever is sensitive to displacements of the target: if the incident light is not normal to the mirror in the plane of interest, then the beam spot will move on the detector if the target moves perpendicular to its surface, with $\Delta x = 2\theta\Delta y$. For lever length l, this corresponds to an effective angular measurement error of $\Delta\theta = \theta\Delta y/l$.

Most of our group's balances make use of an optical autocollimator, which collimates the light from a point source, reflects it one or more times off of a target mirror, and then focuses the light onto a position-sensitive photodetector using the original collimating lens as shown in Figure 2.2. By placing the light source and detector in the focal plane of the lens (with the assistance of a beamsplitter), a conversion between transverse position-

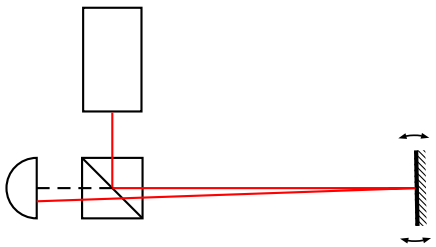

Figure 2.1: Schematic diagram of an optical lever. A collimated beam of light is reflected off of the target and onto a position-sensitive detector (PSD).

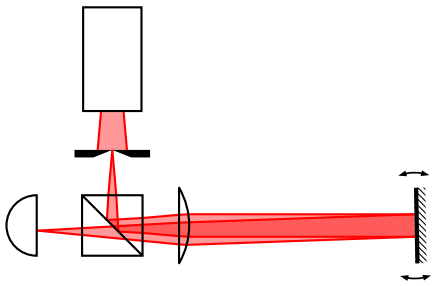

Figure 2.2: Schematic diagram of an autocollimator. Light from a pinhole is collimated, reflected off of the target, and focused onto a position-sensitive detector (PSD).

and momentum-space is made. An angular deflection of the target mirror results in a change in the transverse momentum of the returning light, which is converted to a transverse displacement of the beam spot on the detector. A longitudinal displacement of the target mirror could result in a change in the position of the returning light on the lens, but this would be converted to a change in the transverse momentum of the light incident on the detector and not on the position of the focused beam spot, making the autocollimator insensitive to displacements of the target. Using the same lens for collimation and focusing (hence the "auto-") reduces the effect of optical aberrations from the lens, and because the beam is spread out for reflection off of the target, imperfections in its surface can be effectively averaged away. Our group's best standard autocollimators have a sensitivity of $\sim 4\,\mathrm{nrad}/\sqrt{\mathrm{Hz}}$. Amplification of the signal can be increased by using longer focal length lenses, at the cost of a larger system size, or by repeated reflections off of the target.

Angular deflection measurements can also be made by interferometric setups, which compare the length of two different optical paths by recombining coherent light that has been sent along each path and measuring the intensity of the resulting interference. These are inherently differential measurements, requiring a comparison either between two points on the target or between the target and a fixed reference. Using two points on the target cancels out common-mode displacements. Instability in the light source can introduce noise, as the difference in path length is measured in wavelengths. Although the range over which interference will occur is limited only by the coherence length of the light, the absolute path length difference is only known modulo a half-wavelength. To achieve a dynamic range larger than this, peak/trough crossings can be tracked or a second system can be used to provide extra information.

For any of these devices, systems limited by instrument noise can be improved by using a number of measurement devices in parallel, with the noise dropping with the square root of the number of devices. In the case of an autocollimator this can be done using a one- or two-dimensional array of light sources that are imaged onto a detector capable of resolving each beam spot, as reported by Cowsik *et al.* [14]. A multi-slit

autocollimator developed by our group has demonstrated sensitivities below 1 nrad/$\sqrt{\text{Hz}}$ [15].

An investigation into the feasiblity of rotating a torsion pendulum about its fiber axis at a frequency much higher than its natural frequency led to interest in developing a device with a sensitivity better than 10 prad/$\sqrt{\text{Hz}}$. Although other limitations precluded the development of this rotating torsion balance, we were able to develop a device that did meet this sensitivity requirement. The device is based upon the work of John Howell's group at Rochester [16–18], which was in turn inspired by the concept of weak-value amplification.

Chapter 3

Weak-value amplification

Weak-value amplification (WVA) is a quantum effect first posited by Aharanov *et al*. in 1988 [19]. Using this phenomenon the value of a measurement can be effectively amplified by proper preselection and postselection of the ensemble of particles used to make the measurement.

As an introductory example, consider a beam of spin-1/2 particles traveling in y that have been preselected to have spin up (= +1/2) in the z basis. The particles are later separated according to their spin in the x basis and those with spin up (= +1/2) are kept or postselected. For this remaining ensemble, their initial spin was up in z and their final spin was up in x. If nothing interacted with the particles between the preselection and the postselection, then the spin during that time will not have changed. Thus, for the intermediate time, one might say that the value of the particles' spin is both +1/2 when projected into z and +1/2 when projected into x. This would suggest that the spin has a value of $\sqrt{2}/2$ directed at 45° between the x- and z-axes. Obviously this is a non-physical conclusion, and in fact was made with the assumption that no interaction with (and hence no observation of) the particles' spin was made during this intermediate time.

Now suppose that an interaction with an outside system is added between the preselection and postselection, causing a transverse displacement (Δr) of the particles proportional to their spin as measured 45 degrees between x and z. The initial uncertainty in the x-z position of the electrons (e.g. the beam width σ_r) sets a scale for which this interaction may be considered small or weak: $\Delta r \ll \sigma_r$. If the interaction is weak, then the effect of the interaction on the transverse position is much smaller than the initial uncertainty and the interaction does not disturb the spin of the particles. This also means that no information about the interaction is available by looking at a single particle. However, when the position of a sufficiently large number of particles ($N_f > \sigma_r^2/\Delta r^2$) is averaged, the shift in position can be resolved and will be found to correspond to a spin of $\sqrt{2}/2$ (surprisingly,

given that the spin of a spin-1/2 particle was being measured). Given that the actual magnitude of the spin is known, this could also be considered as an amplification of the interaction.

In the given example the gain is quite small—only $\sqrt{2}$ larger than what would be gained by measuring without the pre- and postselection. The more orthogonal the initial and final states are made to be, however, the greater the gain. If the postselected state was chosen to be much closer to spin down ($=-1/2$) in the z basis, the magnitude of the intermediate backprojected spin must have a value much larger than $1/2$, giving an amplification factor much larger than 1. In fact, as the postselected spin approaches $-1/2$ in the z basis, the amplification will approach infinity. This amplification comes at a cost, however, as the number of particles that make it through the postselection will drop. Given some number of incoming particles N_0 and an amplification factor A, where $A = 1$ for the case where the initial and final selection states are the same, the number making it through the postselection $N_f = N_0/A^2$, a general result for WVA. With the resulting displacement $\Delta r = A\Delta r_0$, the number of postselected particles necessary to resolve the displacement is $N_f = \sigma_r^2/(A\Delta r_0)^2$, giving a required incoming number of $N_0 = A^2 N_f = \sigma_r^2/\Delta r_0^2$. In other words, the quantum limit of the measurement is independent of the amplification factor—quantum noise cannot be overcome. As will be shown later, however, this amplification does have benefits in systems that are not quantum limited.

In general terms, a weak-value measurement consists of a preselection of a certain quantum state; an interaction with the system to be measured, which entangles the state of the particles with one of their continuous degrees of freedom; and a postselection on a different state. The measurement is made "weak" by requiring that the direct result of each interaction, as measured in the continuous degree of freedom, be smaller than the uncertainty in that degree of freedom. By making the preselection and postselection states nearly orthogonal, the effect of the interaction, as observed in the selected degree of freedom after postselection, can be greatly amplified. As already stated, this amplification comes at the cost of a decreased sample size: the number of postselected particles is inversely proportional to the square of the amplification factor.

Since its proposal, WVA has been highly controversial, and continued work by Aharanov *et al*. still leads to argument [20]. However, experimental realization of weak-value amplification was demonstrated by Ritchie *et al*. in 1991 [21], and has since been used in a number of demonstrations and experiments, including the first measurement of the photonic spin Hall effect [22, 23]. It was also the basis for the development of an optical setup highly sensitive to mirror rotations, which will be discussed in the next chapter.

Chapter 4

The Rochester device

A paper by Dixon *et al.* from John Howell's group at the University of Rochester describes a device that uses weak-value amplification to measure angular deflection [16]. Their device, subsequently referred to as the Rochester design, is based on a Sagnac interferometer.

In a free-space Sagnac interferometer, collimated light is directed into a beamsplitter. The beamsplitter splits the light into two beams traveling the same circuit, which is formed by a set of mirrors, but in opposite directions. After completing their circuits the two beams are recombined at the original beamsplitter. (A three-mirror example design is shown schematically in Figure 4.1.) For a perfectly aligned setup this would result in total constructive interference in the direction from which the beam entered (the "light port") and total destructive interference in the other exit direction (the "dark port"). Sagnac interferometers can also be constructed using fiber optics or a ring-shaped lasing medium.

The Rochester design, shown schematically in Figure 4.2, uses a free-space Sagnac geometry with an intentional phase offset introduced between the two paths. This phase offset is produced either by the combination of a half-wave plate and a Soliel-Babinet compensator or by a small out-of-plane deflection of one of the mirrors (see [17]) and allows some light to exit via the dark port. Interference between the beams causes sufficiently small in-plane deflections of the target (or, in fact, any of the mirrors) to result in movement of the beam spot at the dark port, with the movement of the spot being proportional to the deflection of the target and inversely proportional to the square root of the fraction of light exiting the dark port.

Put in our earlier WVA terms, photons are the particles used to make the measurement. The photons are directed into the Sagnac interferometer, where the selection basis consists of the two paths available to the photons. The preselected state is a superposition of the two paths, and the observed degree of freedom

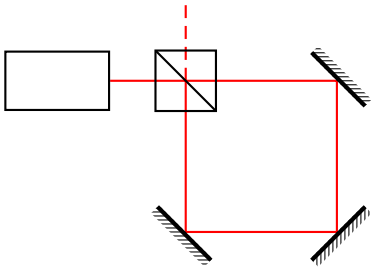

Figure 4.1: Schematic diagram of a Sagnac autocollimator. A beam of light is sent in opposite directions around the same path by means of a beamsplitter. Perfect alignment would result in total constructive interference of the light exiting in the direction of the light source (the "light port") and total destructive interference at the remaining side of the beamsplitter (the "dark port").

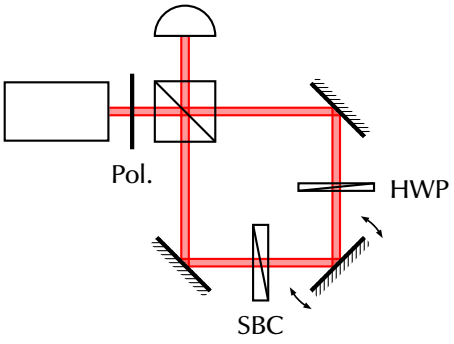

Figure 4.2: Schematic diagram of the Rochester device. A phase offset between the two paths allows some light to exit the dark port. In-plane rotations of any mirror causes a shift of the position of the beam spot at the dark port.

is the transverse momentum of the photons in the plane of the interferometer. The postselection is done by recombining the beams at the dark port of the beamsplitter, with the postselected state determined by the phase difference between the light traveling the two paths. For in-plane mirror deflections resulting in a change in transverse momentum smaller than the transverse momentum uncertainty, weak-value amplification results in proportional displacements of the beam spot at the dark port. As in the case of an autocollimator, a conversion between momentum-space and position-space occurs, but here comes about because the weak-value amplification factor is imaginary. The magnitude and sign of the amplification factor can be adjusted by varying the phase offset of the two paths.

The Rochester design offers the same intrinsic quantum noise (*i.e.* photon shot noise) limit as an autocollimator, as explained by Starling *et al.* [17]. However, it has certain advantages when compared to an autocollimator. Both an autocollimator and the Rochester design have an enhanced signal-to-quantum-noise ratio when compared with a simple beam-deflection setup with the same collimated beam diameter. The ratio improves in an autocollimator because focusing the beam onto the detector reduces the position uncertainity or quantum noise. In the the case of the Rochester device, the ratio improves because the signal is larger, which also improves the ratio of the signal to technical (*e.g.* electronic or digitization) noise. Another advantage is that the design is size independent—the amplification factor and sensitivity do not depend on the size of the setup or the focal length of a focusing lens. Additionally, the number of photons incident on the detector for a given signal-to-quantum-noise ratio is lower compared to that of an autocollimator, allowing for the use of low-saturation-intensity detectors. It should also be noted that a Michelson interferometer has the same intrinsic quantum noise limit as an autocollimator and the Rochester design [24]. Compared to a Michelson interferometer, the Rochester design and autocollimators are less sensitive to optical-path-length fluctuation and have the possible advantage of making a position-based, rather than intensity-based, measurement.

Besides being free to rotate about the axis of the torsion fiber, torsion balances also swing. If the Rochester setup were used to monitor a torsion balance, displacements normal to the mirror surface would result in equal displacements (multiplied by $\sqrt{2}$) of the laser spot on the detector, as shown in Figure 4.3. Such displacements would be indistinguishable from rotation of the pendulum. Despite the implementation of swing-damping techniques, our torsion pendulums have displacement noise amplitudes of about 5 μm, which would limit angular detection using the Rochester setup to sensitivities well above those available with our existing autocollimators.

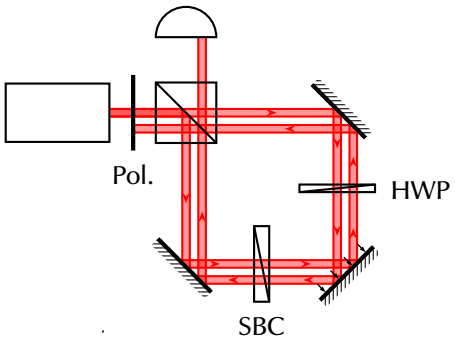

Figure 4.3: Schematic diagram of the effect of longitudinal displacement of the target mirror in the Rochester device.

Chapter 5

The iQuAC

To make the benefits of the Rochester device available for use in an autocollimator replacement, we have developed an angle-measuring device designed to be insensitive to translations and out-of-plane displacements of the target mirror. We refer to this device as an interferometric quasi-autocollimator (iQuAC). In order for the device to be insensitive to target displacements, the laser beams incident on the target must be approximately normal to the reflecting surface. This requires an extra degree of freedom by which the two paths of the Rochester design can be distinguished, which is provided by manipulating the polarization of the beams.

In one form of our design, shown in Figure 5.1, the light passes through a polarizer and into a 50/50 beam splitter. The reflected light then reflects off two mirrors and passes through a half-wave plate, which rotates the polarization by 90 degrees, allowing it to pass through a polarizing beam splitter. It then passes through a quarter-wave plate, reflects off the target mirror at normal incidence, and returns through the quarter-wave plate. The two passes through the quarter-wave plate rotate the polarization by another 90 degrees, causing the light to reflect off the polarizing beam splitter. The light then is reflected by three mirrors and is returned to the 50/50 beam splitter. Light that initially passes through the 50/50 beam splitter follows the same path in reverse. An intentional slight out-of-plane misalignment of one of the non-target mirrors allows for some light to exit the nominally dark port of the interferometer due to the resulting difference in path lengths. An in-plane rotation of the target results in a displacement of the beam spot at the dark port, and the magnitude of the out-of-plane misalignment controls the weak-value amplification factor of the displacement. The mirror arrangement in the interferometer is such that the two paths between the two beam splitters are of equal length, fixing the target in the center of both paths between the light and dark ports and mitigating the effects of out-of-plane angular displacements of the target. For the weak-value amplification to occur, the number of reflections in each path

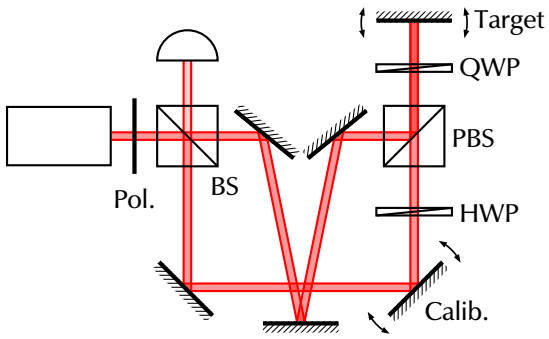

Figure 5.1: Schematic diagram of one form of the interferometric quasi-autocollimator (iQuAC).

must be such that one path has an even number of reflections after the target and the other has an odd number, resulting in the light from the two paths having opposite in-plane transverse momentum due to reflections off the target when exiting the dark port. This requirement precludes the use of a more symmetric arrangement, such as removing the dogleg and placing the polarizing beam splitter and non-polarizing beam splitter at opposite corners of a simple rectangle.

We constructed an iQuAC setup using a 10-mW, 660-nm diode laser (QPhotonics QFLD-660-10S) connected via a single-mode fiber to a 3.4-mm diameter collimator (Thorlabs F280FC-B), standard optics, and a 5-mm position-sensitive photodetector (OSI Optoelectronics SL5-1), shown in Figure 5.2. The signal from the photodetector was amplified, low-passed, and then digitized and read using a data-acquisition board. A piezoelectric disk (Piezo Kinetics ND 0.205-0.000-0.080-502) connected to a signal generator was used to generate in-plane angular displacements of one of the mirrors. It was calibrated by removing the 50/50 beam splitter and making a simple beam-deflection measurement. A 2-Hz sine-wave voltage signal was applied to the piezoelectric disk and used as a calibration signal to determine the weak-value amplification factor.

A noise plot from the setup is shown in Figure 5.3. We observe a noise floor of less than 10 prad/$\sqrt{\text{Hz}}$ in the 10–200 Hz band. Below 10 Hz $1/f$ noise dominates and above 200 Hz acoustic pickup increases. This noise floor was only achieved by enclosing the setup in a polystyrene box and suspending the optical breadboard from three points by approximately 150 rubber bands (Alliance 117B), as shown in Figure 5.4.

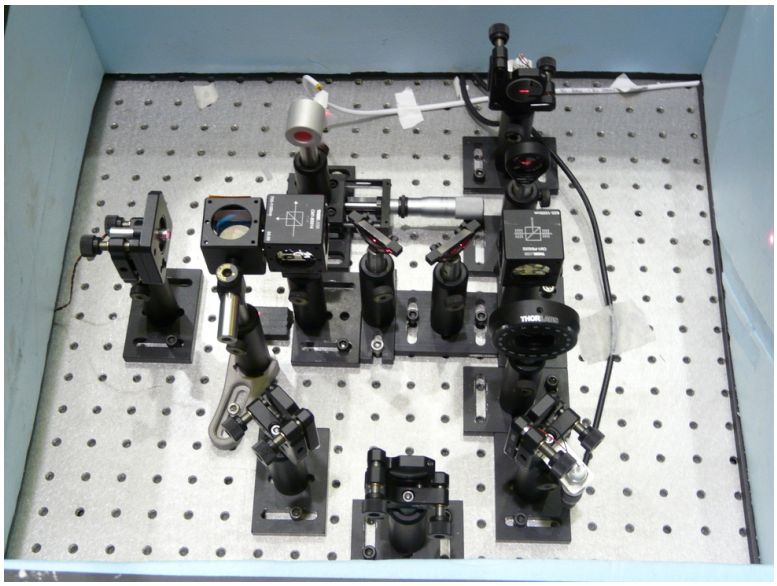

Figure 5.2: A photograph of the iQuAC prototype.

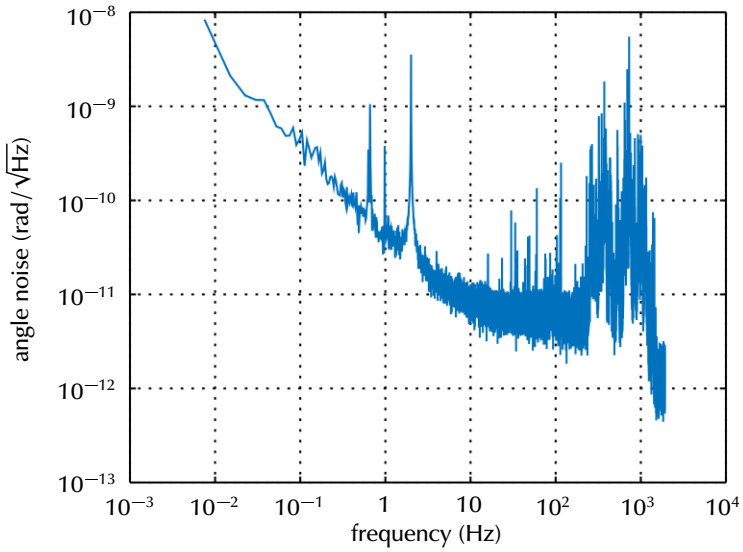

Figure 5.3: A plot of the noise floor of an implemented iQuAC setup. A 620-prad calibration signal is visible at 2 Hz. A weak-value amplification factor of 60, as compared to a beam-deflection measurement, was measured for this data.

Figure 5.4: The setup used to isolate the iQuAC demonstrator from seismic and acoustic noise sources.

Chapter 6

Conceptual derivation

Howell *et al.* have provided the mathematical quantum-mechanical and classical derivation of the weak-value amplification scheme used by both the Rochester design and our design [18]. The classical description can also be explained in conceptual terms.

Suppose a spatially-uniform coherent source is directed into the interferometer. If the system were to be exactly aligned with zero phase offset between the two paths, complete destructive interference would occur at the dark port. If the target were to be rotated in plane, a series of equally spaced fringes, oriented perpendicular to the plane of the apparatus, would appear due to the opposite transverse momenta of the two beams. The center point would have zero intensity, and the fringe spacing would be inversely proportional to the angle of the target. The introduction of a phase offset between the two paths would result in the fringe pattern being displaced by the same phase, regardless of the fringe spacing (or angle of the target), so that the intensity of the center point is the same for any angular displacement of the target. A representation of these fringes is shown in Figure 6.1.

Now suppose that a source with a Gaussian profile is used. The fringe pattern would still be present but would be multiplied by the beam profile. For sufficiently small angular deflections, the fringe spacing would be large compared to the width of the beam, and the fringe pattern in the region of the beam spot could be approximated by a linear function. A linear function of a small slope multiplied by a zero-centered Gaussian is approximately equal to a Gaussian of equal width that has been translated by some distance, *i.e.* for $b\sigma \ll a$,

$$(a + bx) \exp\left[-x^2/\left(2\sigma^2\right)\right] \approx a \exp\left[\frac{-\left(x - b\sigma^2/a\right)^2}{2\sigma^2}\right]. \tag{6.1}$$

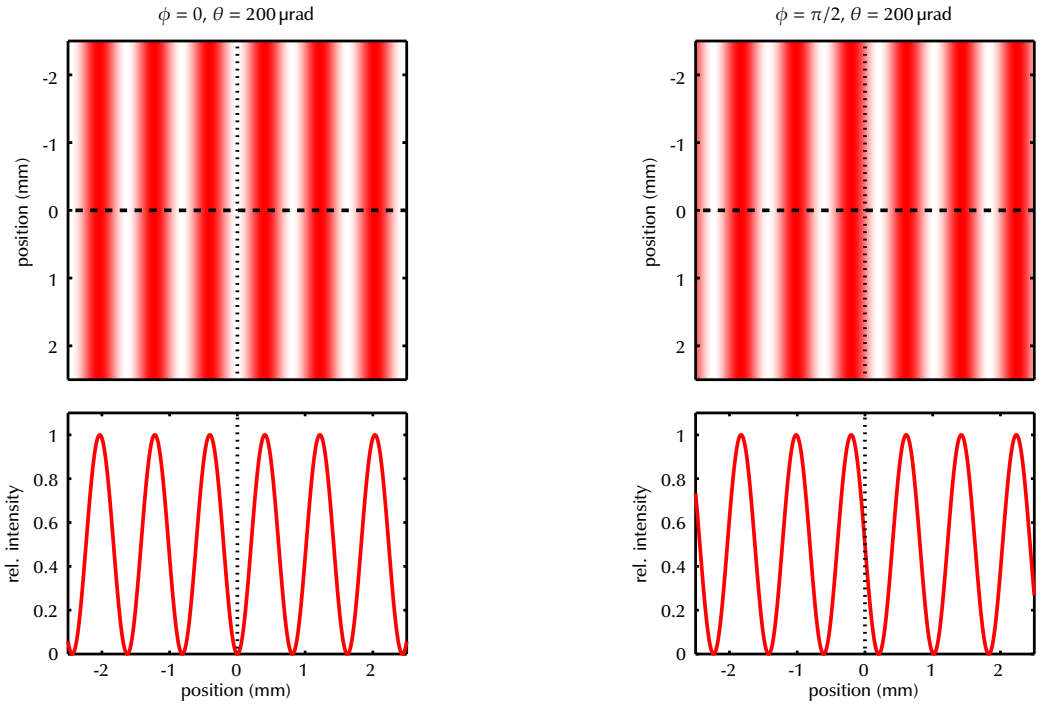

Figure 6.1: Two simulated images and horizontal intensity profiles for plane-wave interference at the dark port with a deflection angle of 200 µrad and phase offsets of 0 and $\pi/2$. The interference pattern shifts by twice the phase offset.

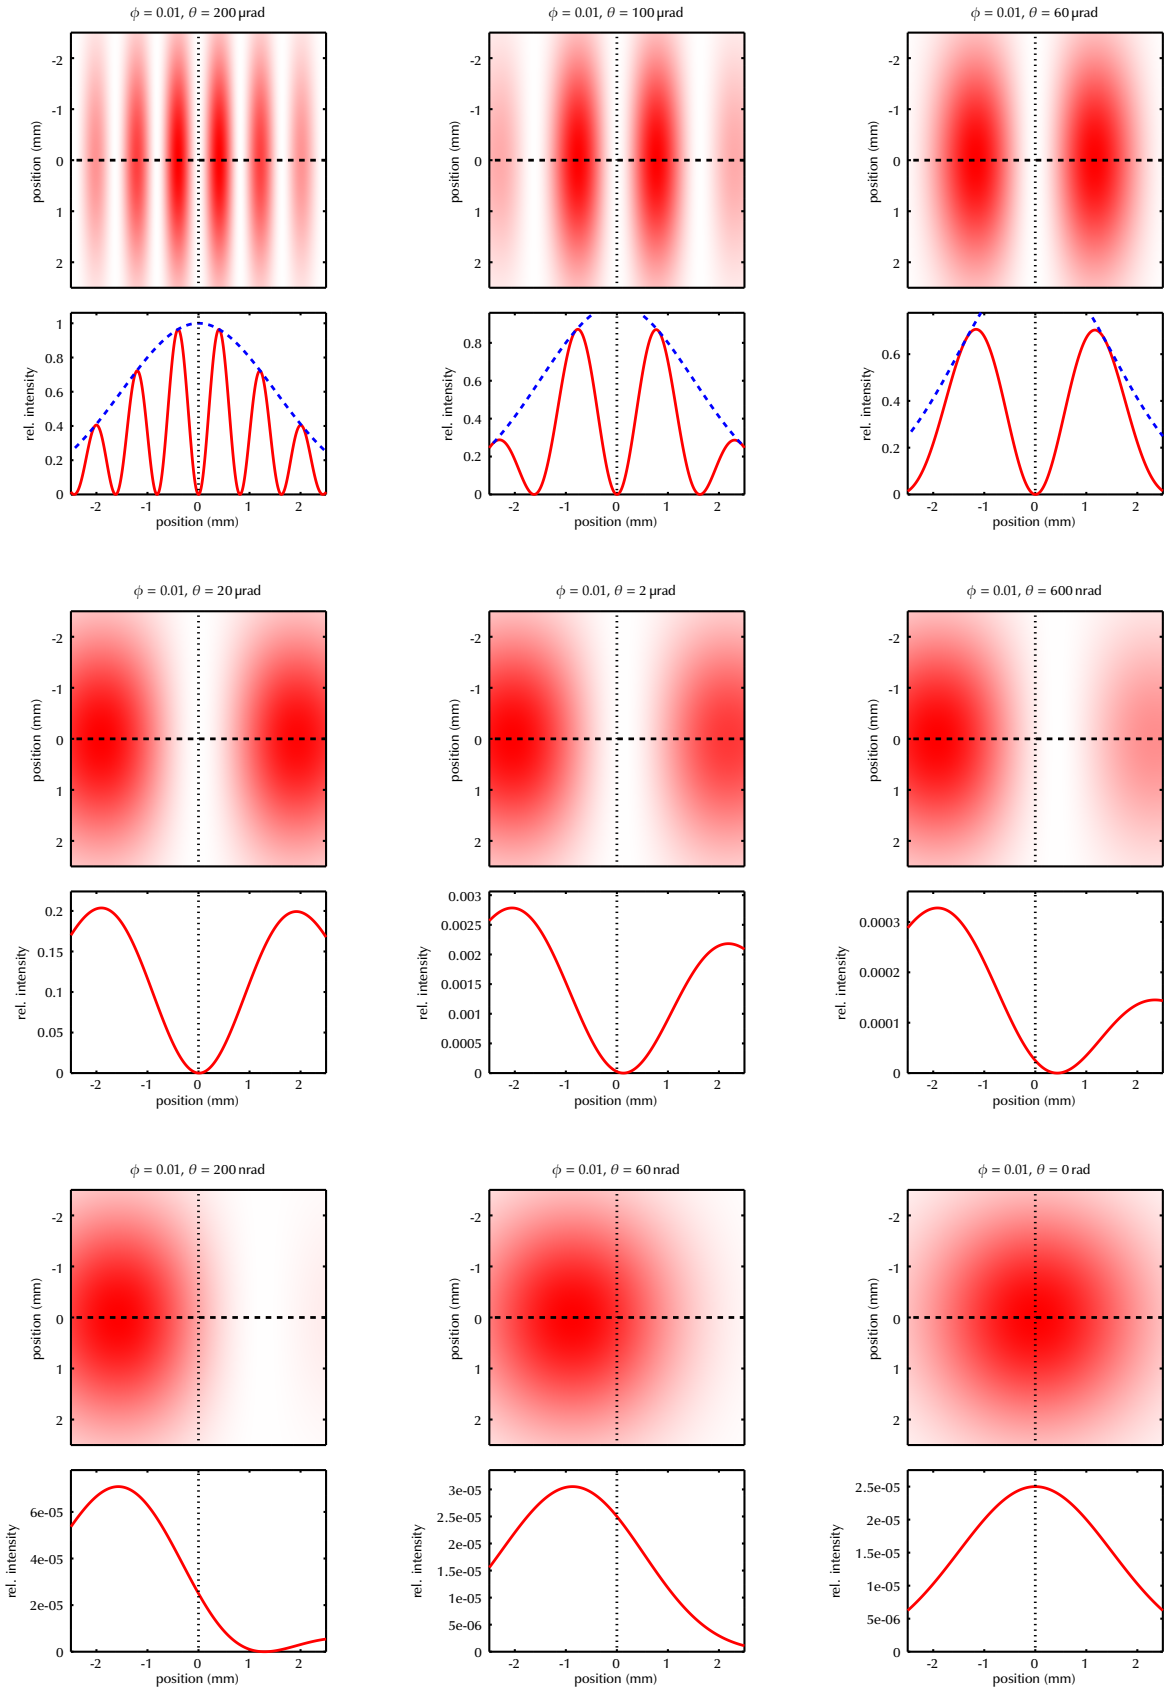

Figure 6.2: A series of simulated images and horizontal intensity profiles of the dark port output for a constant phase offset of $\phi = 0.01$ and decreasing deflection angle θ (left to right, top to bottom). The intensity axis is rescaled to accomodate the 40,000-fold decrease in peak intensity from the first plot to the last.

For sufficiently wide fringe spacing (or small angular deflections), the fringe pattern in the region of the Gaussian can be approximated by the linear function

$$\sin^2(\phi/2) - 2k_0\theta(\sin\phi)x,$$

where ϕ is the phase offset, k_0 the wave number, θ the angular deflection of the target, and x the transverse position. Replacing the values in Equation 6.1 and approximating for small values of ϕ gives a translation of $8k_0\sigma^2\theta/\phi$. This result corresponds to an amplification, as compared to a simple beam-deflection setup with target-detector separation l_{td}, of $4k_0\sigma^2/(\phi\, l_{td})$, and matches the result of Howell *et al*. A representation of the multiplication of a changing fringe pattern with a Gaussian profile is shown in Figure 6.2.

Appendix A

Exact derivation

Dixon *et al.* have introduced a method to make interferometric weak-value measurements of optical beam deflections [16]. The quantum mechanical descriptions given by both their original and later papers make use of approximations to derive the final observable [18]. It is also possible to obtain the value of interest without the use of approximations while maintaining a quantum mechanical description.

The description presented by Dixon *et al.* begins with the state

$$|\Psi\rangle = \int dx\ \psi(x)|x\rangle \exp(-ix\hat{A}k)|\psi_i\rangle \tag{A.1}$$

where x is the transverse position, $\psi(x)$ is the initial position wave function, k is the momentum shift given by the mirror, $\hat{A}$ the system operator $|\circlearrowleft\rangle\langle\circlearrowleft| - |\circlearrowright\rangle\langle\circlearrowright|$, and $|\psi_i\rangle$ the system state $|\psi_i\rangle = (ie^{i\phi/2}|\circlearrowleft\rangle + e^{-i\phi/2}|\circlearrowright\rangle)/\sqrt{2}$. The states $|\circlearrowleft\rangle$ and $|\circlearrowright\rangle$ represent the which-path information. From this point the referenced papers expand the exponential to first order in k, which allows the system operator to act directly on the state $|\psi_i\rangle$. The final system state $|\psi_f\rangle = (|\circlearrowleft\rangle + i|\circlearrowright\rangle)/\sqrt{2}$ is then applied, and the result is factored and reexponentiated to obtain

$$\langle\psi_f|\Psi\rangle = \langle\psi_f|\psi_i\rangle \int dx\ \psi(x)|x\rangle \exp(-ixA_w k), \tag{A.2}$$

where the weak value $A_w = \langle\psi_f|\hat{A}|\psi_i\rangle/\langle\psi_f|\psi_i\rangle$. This technique does have the benefit of clearly displaying the weak-value amplification effect by replacing $\hat{A}$ with the weak value, but the result is not the exact solution, as confirmed by comparison with the classical derivation in [18].

The expansion of the exponential was likely prompted by the existence of an operator in the exponent. Let

us now consider the full expansion of the exponential,

$$\exp(-ix\hat{A}k) = \sum_{n=0}^{\infty} \frac{(-ix\hat{A}k)^n}{n!} = 1 - ix\hat{A}k - \frac{x^2\hat{A}^2k^2}{2} + \frac{ix^3\hat{A}^3k^3}{6} \ldots . \tag{A.3}$$

For notational simplicity, we choose to now represent the which-path information in matrix form. We choose the representation

$$|\circlearrowleft\rangle = \begin{pmatrix} 1 \\ 0 \end{pmatrix} \text{ and } |\circlearrowright\rangle = \begin{pmatrix} 0 \\ 1 \end{pmatrix}, \tag{A.4}$$

which results in

$$\hat{A} = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}. \tag{A.5}$$

One may now recognize that $\hat{A}^2 = \hat{I}$, where $\hat{I}$ is the identity matrix (as is the case with the Pauli matrices). By replacing even powers of $\hat{A}$ with $\hat{I}$ and odd powers with $\hat{A}$, the above expansion is simplified:

$$\begin{aligned}
\exp(-ix\hat{A}k) =& \hat{I} - ikx\hat{A} - \frac{x^2k^2}{2}\hat{I} + \frac{ix^3k^3}{6}\hat{A} \ldots \\
=& \sum_{n=0,\ \text{even}}^{\infty} \frac{(-ikx)^n}{n!}\hat{I} + \sum_{n=1,\ \text{odd}}^{\infty} \frac{(-ikx)^n}{n!}\hat{A} \\
=& \hat{I}\cos(-kx) + i\hat{A}\sin(-kx) \\
=& \hat{I}\cos(kx) - i\hat{A}\sin(kx) \\
=& \begin{pmatrix} \cos(kx) + i\sin(kx) & 0 \\ 0 & \cos(kx) - i\sin(kx) \end{pmatrix} \\
=& \begin{pmatrix} e^{ikx} & 0 \\ 0 & e^{-ikx} \end{pmatrix}.
\end{aligned} \tag{A.6}$$

Note that the series expansion used here, which serves to elucidate the mathematics, is not truncated: no approximations have been made.

The final state $\langle\psi_f|\Psi\rangle$ can be now be found by inserting the simplified exponential, as well as the vector forms of the initial and final system states,

$$|\psi_i\rangle = \frac{1}{\sqrt{2}}\begin{pmatrix} ie^{i\phi/2} \\ e^{-i\phi/2} \end{pmatrix} \text{ and } |\psi_f\rangle = \frac{1}{\sqrt{2}}\begin{pmatrix} 1 \\ i \end{pmatrix}, \tag{A.7}$$

into Equation A.1, giving

$$\begin{aligned}
\langle\psi_f|\Psi\rangle =& \langle\psi_f| \int dx\ \psi(x)|x\rangle \exp(-ix\hat{A}k)|\psi_i\rangle \\
=& \frac{1}{\sqrt{2}}\begin{pmatrix}1 & -i\end{pmatrix} \int dx\ \psi(x)|x\rangle \begin{pmatrix} e^{ikx} & 0 \\ 0 & e^{-ikx}\end{pmatrix} \frac{1}{\sqrt{2}}\begin{pmatrix} ie^{i\phi/2} \\ e^{-i\phi/2}\end{pmatrix} \\
=& \frac{1}{2}\int dx\ \psi(x)|x\rangle \begin{pmatrix}1 & -i\end{pmatrix}\begin{pmatrix} ie^{i(\phi/2+kx)} \\ e^{-i(\phi/2+kx)}\end{pmatrix} \\
=& \frac{1}{2}\int dx\ \psi(x)|x\rangle \left(ie^{i(\phi/2+kx)} - ie^{-i(\phi/2+kx)}\right) \\
=& -\int dx\ \psi(x)|x\rangle \sin(\phi/2 + kx).
\end{aligned} \tag{A.8}$$

It should be noted that this result is equivalent to that obtained by the classical treatment in [18].

Given this simplification, the position expectation

$$\langle x\rangle = \frac{\int dx\langle\Psi|\psi_f\rangle x\langle\psi_f|\Psi\rangle}{\int dx\langle\Psi|\psi_f\rangle\langle\psi_f|\Psi\rangle} \tag{A.9}$$

can be calculated exactly for certain beam profiles. For the case of a Gaussian profile, where $\psi(x) = \exp(-x^2/(4\sigma^2))$, the integrals are known, and the expectation value is found to be

$$\langle x_{\text{Gaussian}}\rangle = \frac{2k\sigma^2 \sin\phi}{\exp(2k^2\sigma^2) - \cos\phi}. \tag{A.10}$$

To compare this solution with the approximate result of the original papers, we expand to first order in k and second order in ϕ, giving

$$\langle x_{\text{Gaussian}}\rangle \approx \frac{2k\sigma^2\phi}{1 - 1 + \phi^2/2} = \frac{4k\sigma^2}{\phi}. \tag{A.11}$$

Part II

A moving-mass gravity-gradiometer torsion balance

Chapter 7

Pendulum and actuator development

The objective of this work was to develop a low-drift gravity gradiometer. Although a gravity gradiometer was the initial goal, undertaking this work also led to the first known demonstration of a torsion pendulum with actuated components in which a property of the pendulum could be adjusted *in situ*. Included in this was the development of a non-magnetic wirelessly powered rotary actuator

The basic idea of this project was to create a torsion pendulum with a quadrupole mass distribution that can be rotated about an axis perpendicular to the torsion fiber without mechanical interaction with the pendulum. Allowing for the mass arrangement to be put in at least 8 different positions, as shown in Figure 7.1, would allow for measurements of both the xy and xz components of the gravity gradient and cancellation of any potential effects linked to the full rotation cycle, such as a changing center of mass or magnetic torques.

The only component of the proposed system that presented any new challenges was the actuator. Beyond the basic capability of rotating the mass arrangement, two criteria for the actuator needed to be met. First, the actuator must be able to be wirelessly powered or driven for a long (>1 month) period of time by an on-board battery. Opening a torsion balance's vacuum chamber and interacting with the pendulum comes with time and material costs as well as risks, such as the possibilities of breaking the fiber or damaging the vacuum seals. Second, the mass configuration must be precisely positionable.

7.1 Early iterations

The initial approach to the problem of rotating weights on the pendulum involved tests with two different DC motors. The first motor was a 6-V DC gearmotor with a 297.92:1 gear ratio (Pololu item #994). This gearbox

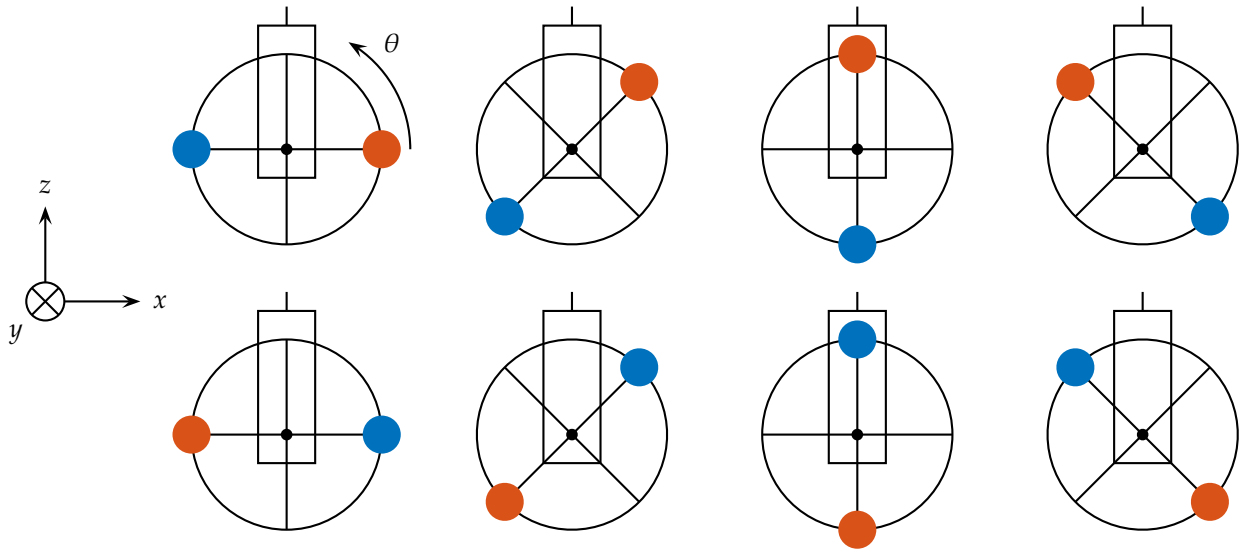

Figure 7.1: The initial concept for an active-mass gravity gradiometer pendulum showing the rotation of two masses on a wheel attached to a torsion pendulum. The masses are shown in different color simply to aid in tracking their positions for a full 360° cycle.

was directly connected to a wheel cut from aluminum sheet. We powered the motor by a capacitor recharged by an on-board button-cell battery and triggered motion by shining a light on a photodiode. The pendulum is shown in Figure 7.2. Even before testing its suitability as a torsion balance, the system exhibited two problems: large disturbances in the "wobble" mode due to the change in angular momentum of the wheel and difficulties in driving the wheel to the desired location.

The problem of the disturbances was resolved by switching to a slightly different motor unit that has an extended shaft on the motor extending out the opposite side of the gear box, presumably for use with a rotational encoder (Pololu item #2218). Fortuitously, the gear box had an odd number of gears, so this shaft turned in the opposite direction as the gearbox output. We designed a small aluminum reaction wheel such that the ratio of its moment of inertia to that of the primary wheel matched the gear reduction ratio and attached it to the extended shaft as shown in Figure 7.3. This arrangement mostly canceled out the change in angular momentum of the pendulum inherent in starting and stopping the wheel, reducing the disturbances that were previously observed by an order of magnitude. Backlash in the gearbox prevented a total elimination of the disturbances.

The difficulty in positioning the wheel was somewhat resolved by cutting semicircular notches into the wheel and attaching a roller-lever snap-action switch to the pendulum body, with the roller pressing against the outside

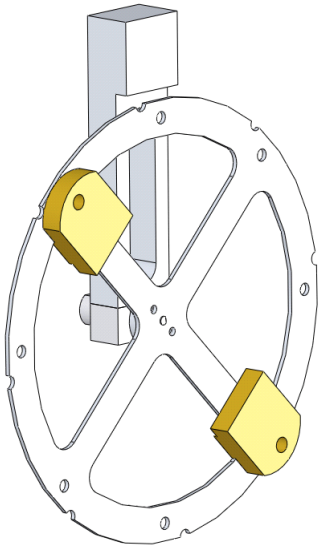

Figure 7.2: A CAD rendering of an early prototype using a DC gear motor.

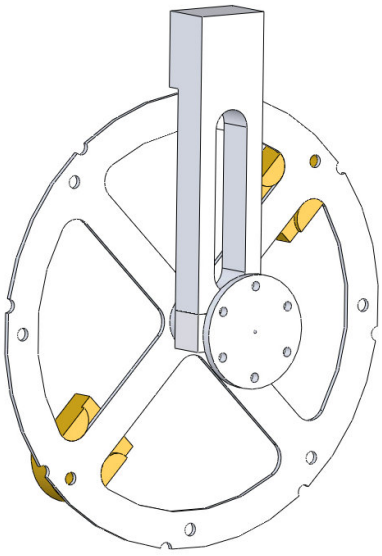

Figure 7.3: A CAD rendering of an early prototype with a reaction wheel attached the the motor's shaft. The reaction wheel, which turns in the opposite direction of the main wheel attached to the gearbox output, serves to cancel out the change in angular momentum when starting and stopping the main wheel.

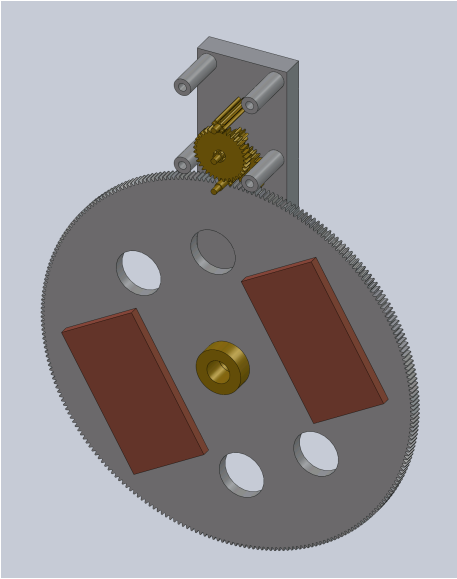

Figure 7.4: A CAD rendering of the second-phase prototype with a custom-built gearbox driven by a ⌀3.2 mm DC motor. The gearbox top plate and wheel shield are not shown.

edge of the wheel. The compression of the spring was adjusted such that the motor's torque was sufficient to drive the roller out of the notch. After an initial pulse, the motor would be turned off and the wheel allowed to coast to the next position, where the roller was pressed into the notch by its spring, stopping the wheel. Once this system was put into a vacuum chamber and its behavior observed, however, it became apparent that the pendulum's changing magnetic dipole, due to the rotation of the motor's rotor, would preclude its use for gradiometry. (Although the magnetic nature of the motor made it unsuitable for this application, it did make feedback quite simple: a large spool of wire was placed outside the chamber, which when driven by the DAQ produced a magnetic field capable of controlling the pendulum's position.)

The magnetic dipole on the pendulum was reduced by replacing the first motor and off-the-shelf gearbox with a tiny (⌀3.2 mm × 8 mm) coreless DC electric motor (Shicoh F7) and a custom-built 668.25:1 gearbox, with the last gear being the pendulum wheel, as shown in Figure 7.4. This arrangement was wirelessly powered by an induction coil arrangement, with a rectifying diode in series with the coil and a smoothing capacitor in parallel with the motor. (Without the capacitor the motor would not turn.) However, despite the motor being coreless and the permanent magnets being very small, variations in the local magnetic field were still strong enough to cause noise above acceptable levels. We then imposed a third criteria for the actuator: that it be entirely non-magnetic. Since standard batteries have steel or nickel components, this narrowed our earlier criteria to require that the mechanism be wirelessly powered.

Other actuator options were considered but not tested in depth. Various types of piezoelectric-based actuators are available. The most promising might have been the rotary piezo stages made by PCBMotor, but the challenges associated with providing the high-voltage power required by piezos to the pendulum discouraged any useful investigation. Other types of piezo motors are available but also require high-voltage power and are typically very expensive or have steel components. Another option considered was a coreless universal electric motor, which uses coils instead of permanent magnets on both the rotor and the stator. No suitably small or coreless off-the-shelf models were found. Even if one had been available, the efficiency may have been too low to make it practical for this application, as heating of the pendulum later turned out to be a possible limiting factor.

7.2 New rotary actuator

A brief investigation of the possibility of driving a linear-to-rotary conversion mechanism with a dual-coil solenoid actuator (analogous to a universal motor) led to the identification and development of such a mechanism, the discovery of a more suitable linear actuator, and a final demonstrator pendulum.

7.2.1 Basic mechanism and wheel

The idea of using a solenoid as the prime mover led to a search for a suitable linear-to-rotary conversion mechanism. A crude mechanism, developed for use in a hobbyist 3D printer and pictured in Figure 7.5, was discovered and adapted for use in this application by increasing the step resolution and making other minor modifications described later. Initial prototyping was done in laser-cut acrylic. Later prototypes and the final version were milled from aluminum.

The functional part of the wheel is an irregularly-shaped circular groove, shown in Figure 7.6, which has offset valleys in the inner and outer walls. A pin is seated into the groove, oriented perpendicularly to the wheel surface. Movement of the pin towards the center of the wheel while contacting the inner wall of the groove causes the wheel to rotate until the pin reaches an inner valley, and movement of the pin away from the center moves the wheel until the pin is in an outer valley. The valleys are offset so that movement of the pin from an outer valley to an inner, or *vice versa*, always causes the wheel to rotate in the same direction. Figure 7.7 shows the movement of the pin from the wheel's frame of reference.

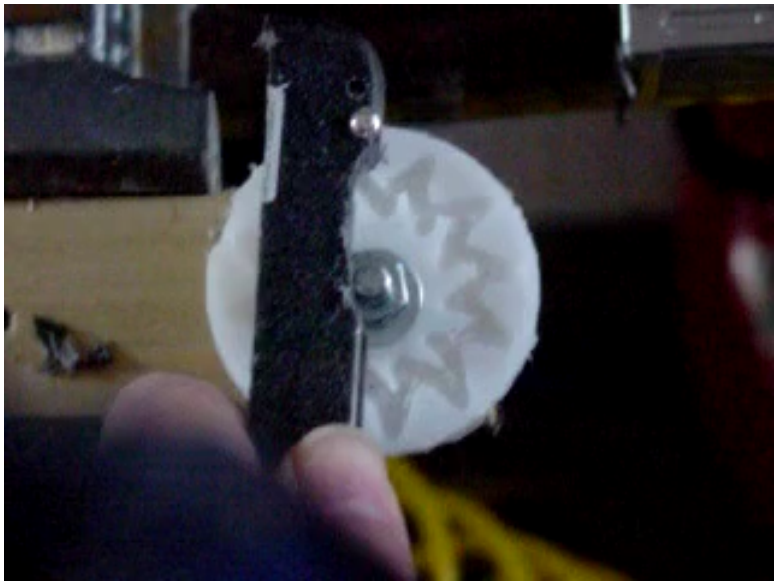

Figure 7.5: A frame from the video providing the inspiration for the rotary actuator [25].

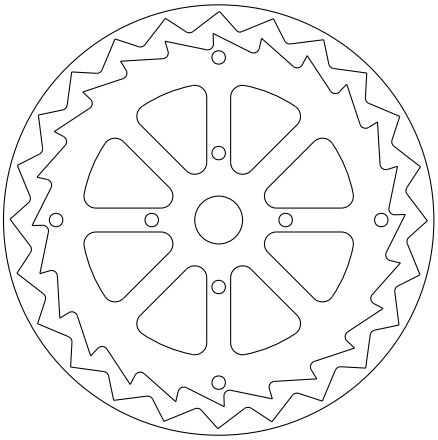

Figure 7.6: A 1:1-scale sketch of the final wheel design showing the groove used for drive and indexing.

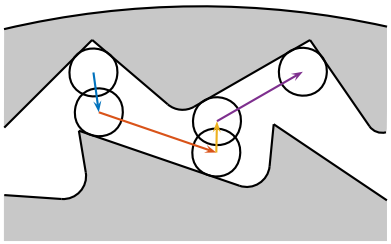

Figure 7.7: 4:1-scale diagram of the movement of the actuator pin from the wheel frame of reference. Starting in an outer valley, the pin moves inward until it contacts the inner wall of the groove. It continues to move inward, rotating the wheel. After a set distance, the pin moves outward until it contacts the outer wall and forces further rotation, completing a single cycle.

Property	Value
product name	Flexinol 150LT Muscle Wire
supplier	Dynalloy Inc.
diameter	150 µm
recommended loading force	3.15 N
lifetime	100,000 cycles
nominal transition temperature	70 °C
contraction against bias spring	3%

Table 7.1: Specifications of the Nitinol wire used in the pendulum nitispec

7.2.2 Nitinol

Although the development of our linear-to-rotary mechanism was prompted by the idea of using a two-coil solenoid, we quickly identified a more suitable linear actuator. Nitinol, a nickel-titanium alloy, is a common shape-memory alloy (SMA) that is used in a number of industrial and consumer applications [26]. SMAs are materials that can be deformed and then restored to an original shape by heating through a crystalline phase transition. In some applications, such as eyeglass frames, the transition temperature is set through the alloy composition to be below standard ambient temperatures, so that the material returns to the original state immediately. For actuators the transition temperature is chosen to be higher than expected ambient temperatures so that the action can be initiated by heating the material. In addition to recovering from deformations, the material can also recover from small elongations, which is the mode used in our actuator. Off-the-shelf linear actuators using Nitinol are available (such as the Miga NanoMuscle), which employ a mechanical advantage to amplify small contractions, but in our pendulum these features are not necessary: a single length of Nitinol wire and a bias spring are sufficient to convert the rotary mechanism into an electrically powered actuator. The specifications of the wire used in the actuator are given in Table 7.1.

7.3 Final assembly

The final pendulum assembly, as shown in Figure 7.8, is based around the aluminum chassis, which is a long aluminum bar with a bored channel passing lengthwise through roughly half its length, a narrower cross-section for the other half, and mounting points for the Nitinol wire, the wheel, and two end plates. The Nitinol wire is held by nuts on two screws at the narrower end of the chassis. These screws pass through oversized holes in the

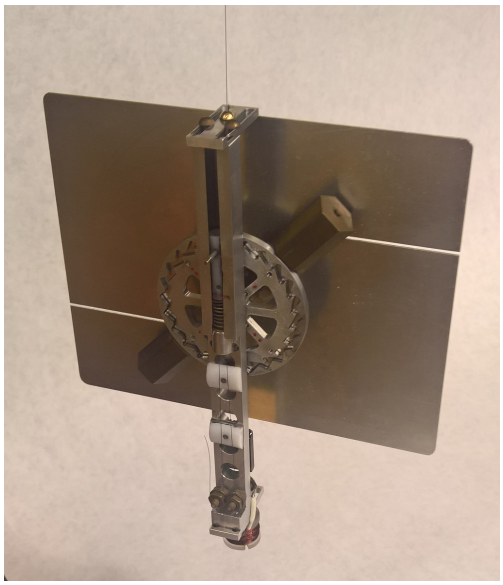

Figure 7.8: A photograph of the final pendulum.

chassis, are wrapped with polyimide sheet within the holes, and have nylon washers immediately outside the holes to electrically isolate their heads and attached nuts from the chassis. The screws also serve as connection points between the terminals of the inductive coil and the Nitinol wire. A short acetal tube (the slider) sits in the bored channel, where it is free to slide. A titanium pin passes through a pair of oppositely placed holes in the tube. The holes were drilled undersized and slit by a small distance parallel to the cylinder axis to give a tight fit on the pin. One side of the pin passes through a slot in the body, which prevents the rotation of the slider within the channel and allows the pin to seat into the groove of the wheel. The axle of the wheel is formed by a brass screw protruding perpendicularly from the chassis, which holds a brass bushing about which the wheel turns. The bias force for the Nitinol is provided by a brass compression spring held within the bored channel between the slider and a spacer at the end of the bore. The Nitinol wire passes through the spring and bore of the slider and is folded in half about the pin.

The changing gravitational quadrupole of the pendulum is provided by two weights made from brass hex stock and weighing 62 g each. Pan-headed screws passing through countersunk holes in the weights and into tapped holes in the wheel fix the weights into well-defined positions 180° apart. The effective radius of gyration of the weights about the wheel axle is 34 mm.

Two end plates are attached to the top and bottom of the pendulum body by brass screws passing through slots in the plates into tapped and vented holes. The torsion fiber is attached to the top plate and the secondary coil of the inductive power supply is attached to the bottom. These plates allow the fiber position and coil

position to be adjusted to accommodate changes in the wheel and weights during development. Adjustment is necessary to allow for the pendulum to hang vertically and for the coil to be positioned along the fiber axis.

Other pendulum features include a gold-coated glass mirror epoxied to the pendulum body for autocollimator readout and an aluminum shield plate, which is held by the wheel axle and serves as the surface on which the potential difference of the feedback electrode exerts a force.

Chapter 8

Complete apparatus

This chapter will cover the hardware and software involved in measuring and controlling the behavior of the pendulum, including readout and feedback on its torsional mode and rotation of the actuator. We identify a number of modifications that could be made to improve the performance and reduce the cost of this instrument in chapter 11.

8.1 Hardware

The hardware consists of both the mechanical structure surrounding and supporting the pendulum and the control system, which includes an autocollimator and camera for input and the feedback electrode, power coil, and "anticoil" for output. Data acquisition and instrument control is done primarily through a National Instruments USB-6211 card and managed by a custom LabVIEW program documented in section 8.2.

8.1.1 Structure

The system is housed in a stainless-steel bell jar. A custom aluminum breadboard was fabricated to mount on existing points and provide access to the necessary electrical and mechanical feedthroughs. A post supports a plate holding three enmeshed gears with axles that seat into holes in the plate. The first gear is securely connected via a rod to a rotary feedthrough in the baseplate of the vacuum vessel. The torsion fiber is connected to the last gear. This setup, a "phi-top" in our group's vernacular and shown in Figure 8.1, allows for manual adjustment of the pretorque angle of the pendulum.

Figure 8.1: A photograph of the top of the pendulum stand showing the phi-top assembly, swing damper, pre-hanger, upper slug, and the upper clamping screw of the main fiber.

8.1.2 Torsion fiber

The torsion fiber consists of three separate tungsten fibers connected in series. As with our other experiments, the initial connection to the phi-top is made with a "pre-hanger", a short fiber of thicker diameter that serves to isolate the main fiber from tilts. Our group has previously found that using a larger diameter fiber at the uppermost connection point reduces the magnitude of tilt-twist coupling due to fiber eccentricity [27]. In addition, since this is our first pendulum with active masses, it is also our first torsion balance with a "post-hanger", which serves to isolate the main fiber from changes in the center of mass of the pendulum. The relevant properties of these fibers are given in Table 8.1.

The pre- and post-hangers are connected to the main fiber by copper and aluminum slugs into which the fiber screws connect. The upper slug sits within a magnet arrangement, shown in Figure 8.1, to provide eddy-current damping of swing. The lower slug serves as a thermal sink to smooth the spread of heat from the pendulum into the main fiber.

The pre- and post-hanger use our group's standard fiber screws and crimp tubes, but the grip provided by the crimp tubes on the smaller diameter main fiber was not reliable in supporting the weight of the pendulum. The main fiber instead uses the clamping screws shown in Figure 8.2, which consist of a small aluminum block

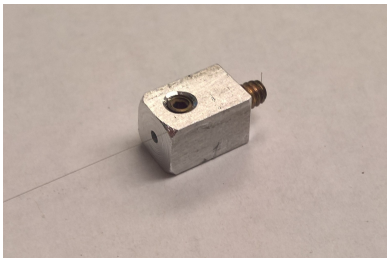

Figure 8.2: A photograph of one of the clamping screws for the main fiber.

Fiber	Diameter	Length	Spring constant
Pre-hanger	76 µm	10.9 cm	4.7×10^{-6} N m/rad
Main fiber	41 µm	24.8 cm	1.7×10^{-7} N m/rad
Post-hanger	76 µm	3.7 cm	1.4×10^{-5} N m/rad
Total	–	48 cm (including slugs and screws)	1.6×10^{-7} N m/rad

Table 8.1: Torsion fiber properties

with two set screws—one extending from the block for connection to the slugs and one for clamping the fiber. The fiber is sheathed and centered inside the clamp by a short length of small-diameter copper tubing. (The tubing is available at low cost as EDM electrode material.)

8.1.3 Control system

Autocollimator

Angular readout of the torsional mode of the pendulum is done using a simple autocollimator (see chapter 2). We constructed the autocollimator from off-the-shelf optical and optomechanical components, including a convex lens, three beam tubes, a cube-mounted beamsplitter, and optical mounting posts. The light source is a 532-nm laser diode powered by a current-regulated circuit and located behind a pinhole. After reflecting off a mirror on the pendulum the light is refocused onto a position-sensitive detector (Edmund Scientific 58-283). Although this particular detector is designed for 2D operation, we ganged together the top and bottom leads on each side to reduce it to a 1D detector and simplify vertical alignment. The current from each channel is amplified by transimpedence amplifiers and then differenced and summed by in-vacuum op-amps before being passed to the DAQ card.

Using a focal length of 200 mm and a detector width of 4 mm, the nominal difference-over-sum conversion

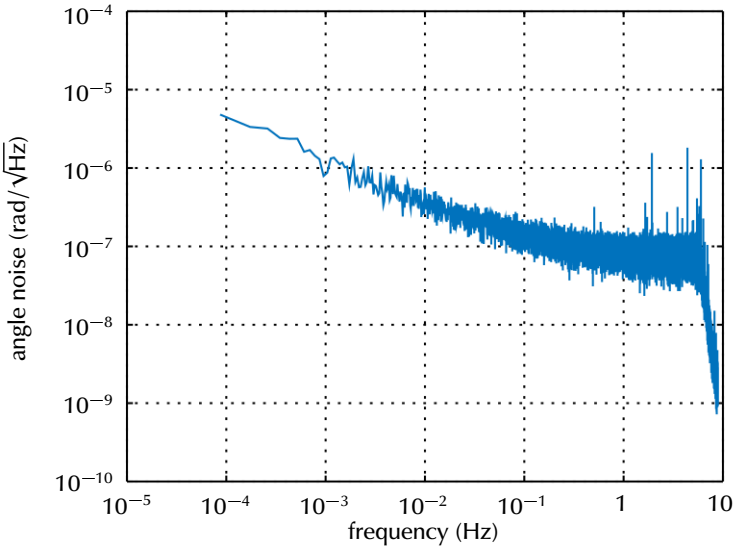

Figure 8.3: An autocollimator noise plot, taken using a stationary mirror, showing a white-noise floor of less than 100 nrad/$\sqrt{\text{Hz}}$ and $1/\sqrt{f}$ noise below 0.5 Hz.

to radians should be 1/200. This was confirmed based on a measurement referencing the known deflection per turn of a Thorlabs mirror mount. A noise curve taken from a reading against a fixed mirror is shown in Figure 8.3.

Camera

The position of the pendulum wheel is monitored using a generic USB webcam, shown in Figure 8.4, that was converted into a vacuum-compatible camera. The gooseneck cable was removed and the four wires were connected to an electrical feedthrough made by epoxying (Varian Torr Seal) copper pins into four holes drilled through a vacuum flange. The camera fit inside a vacuum nipple and a window was attached at one end and the feedthrough at the other. Putting the camera within this vacuum-tight enclosure allowed it to remain in atmospheric-pressure air, reducing the chance of arcing or overheating.

The feedthrough and electrical connections limit communications with the camera to USB1 protocol, which, along with the lack of light in the chamber, only allowed for snapshots of the pendulum rather than continuous video monitoring. USB1 protocol is enforced by connecting the camera to the computer through a USB1 hub. Two white LEDs were positioned to shine against the chamber wall, which reflects and diffuses their light onto the pendulum. The LEDs are connected to a DC power supply with an inline MOSFET controlled by the DAQ card.

Because of space constraints in the chamber, a mirror is used to allow the camera to see the pendulum. In

Figure 8.4: A photograph of the USB camera that was placed inside a vacuum nipple to create a low-cost vacuum compatible camera. An AAA battery is shown for scale.

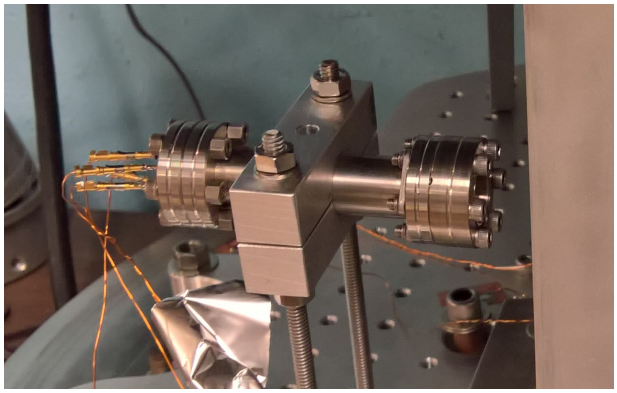

Figure 8.5: A photograph of the installed vacuum-compatible camera.

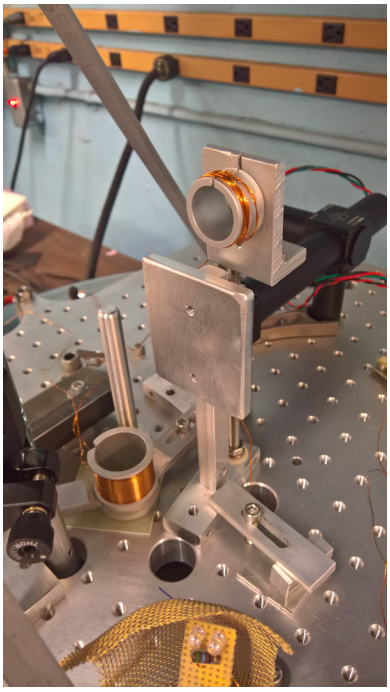

Figure 8.6: A photograph showing the electrode (center), anticoil (above electrode), and power coil (on base-plate). The post to the left of the electrode is a bumper that limits the rotation of the pendulum away from the electrode.

addition to its primary purpose of monitoring the wheel rotation, the camera can also be used to identify the position of the pendulum when it goes off the detector.

Feedback electrode

As with some other torsion balances built by our group [10, 28], the WashCycle pendulum is held in electrostatic feedback. This is especially necessary due to the unexplained and relatively large torque experienced by the pendulum when it is powered. We installed a single feedback electrode with a large active area, shown in Figure 8.6, which acts against the pendulum shield. The electrode is an aluminum plate securely fastened to and electrically isolated from a post. The electrode's potential difference from the grounded structure and pendulum is driven by an amplifier, which amplifies a voltage signal from the DAQ card by a factor of 14. A 10-kΩ resistor placed in series with the electrode limits the current for safety and for when shorts to the pendulum occur.

To limit the rotation of the pendulum away from the electrode, we placed an aluminum post a significant distance away from the shield side of the pendulum on the opposite end as the electrode. The pendulum does not come into contact with this post during normal operation, but it reduces the time to damp and catch the pendulum after large disturbances.

Power coil

As noted in chapter 7, the pendulum is powered by an inductive-coil arrangement. A larger drive coil affixed to the baseplate surrounds a smaller coil on the pendulum. The drive coil consists of ~100 turns of wire on an 32-mm diameter aluminum spool. The spool is slotted to reduce the formation of eddy currents within the spool. The secondary coil has ~100 turns of wire on a 13-mm diameter polycarbonate spool.

The AC current for the coil is provided by an amplifier driven by an amplitude-modulated function generator, with the amplitude set by an analog output from the DAQ card. In the final setup the function generator is set to produce a 5-kHz 6.8-V peak-to-peak sine wave at full amplitude, which the amplifier amplifies by a factor of 3.6.

Anticoil

As mentioned, an unexplained torque is exerted on the pendulum when it is powered. This torque exceeds the torque available from the feedback system, so a second coil was installed above the feedback electrode, with its axis perpendicular to the pendulum's shield. Passing an AC current through this coil, which consists of 10 turns on a 25-mm diameter slotted aluminum spool, exerts a torque on the pendulum via eddy-current forces, countering the unwanted torque from the power coil. Because of its role this coil is referred to as the "anticoil". The anticoil is powered by an amplifier driven by a PWM output from the DAQ card, which provides a current of 6 A. The torque on the pendulum is set by changing the frequency of the PWM signal. The anticoil only exerts a torque on the pendulum when the current is changing, which, for a square wave, only happens at the edges. The angular impulse from one edge does not depend on the PWM frequency, but the frequency does set the number of edges per second, which makes the time-averaged torque go linearly with the frequency.

8.2 Software

Only the features involved in normal operation of the current system are documented here. The program has a number of additional features, many being vestigal additions created during development and earlier prototypes.

8.2.1 Acquisition and preprocessing

The DAQ card collects the difference and sum signals from the autocollimator at 200 Hz. The data is low-passed and averaged down to 20 Hz before passing through a notch filter at 690 mHz and a low-pass filter at 1 Hz to remove the swing and wobble modes from the signal. These secondary filters are disabled for better feedback

response when catching the pendulum. The data is then passed to the feedback loop and further downsampled for recording.

8.2.2 PID feedback

The potential for the feedback electrode is determined by a manually tuned PID (proportional-integral-derivative) loop. Because the derivative term (D-term) acts to reduce the angular momentum of the pendulum, to be properly tuned it must have some relationship to the moment of inertia of the pendulum. As the weights move from their vertical to horizontal positions, the moment of inertia of the pendulum roughly doubles, necessitating a corresponding change in the D-term. We addressed this by first tuning and recording the D-terms for these two positions, D_V and D_H. We then added code to generate the D-term $D(\theta)$ for the intermediate positions using the function

$$D(\theta) = D_V + (D_H - D_V)\cos^2\theta. \tag{8.1}$$

Because the electrostatic torque on the pendulum goes with the square of the potential difference between the pendulum and the electrode, the square root of the PID output is used to set the voltage of the electrode after applying a gain factor. For ease of operation the PID loop has other features such as a bumpless integral term (I-term), which rescales the integral accumulator when the loop gain is adjusted. A description of the calibration of the feedback system is given in section 9.5.

8.2.3 Switching

Moving the pendulum to a new position ("switching") consists of a number of steps:

1. A logged switching flag is set to 1 to identify in the data that a switch routine is in progress.
2. The I-term of the PID loop is pinned to its current value to limit the effect of the switching torque and the anticoil on the science signal.
3. The LEDs are switched on and a series of ten exposures is taken by the camera, which is necessary to allow the camera to adjust to the low light levels. The last exposure is displayed and stored. USB signal issues have sometimes caused the camera to stop working. If this happens the software will attempt to restart the USB1 hub. If this fails to re-enable the camera, the software will simply proceed with the rest of the sequence.
4. The D-term of the PID loop is changed to the appropriate value for the next position, as described above.

5. The power coil and anticoil are switched on for 15 s.
6. The power coil signal amplitude is ramped down over 300 ms to provide a degaussing effect.
7. After a wait time of 20 s corresponding to the Nitinol relaxation time, the I-term is unpinned.
8. The switching flag is reset to 0.

8.2.4 Data recording

When recording is enabled the user is prompted for a filename. A header section is written, documenting static parameters such as sample rate and PID coefficients (these parameters are locked during recording), followed by the data stream for each channel. Data is written at 10 Hz.

Chapter 9

Data analysis and test signal

The data analysis software used to validate the hardware prototype is written in Octave and has been written to be used within a single Octave session. The goal of the data analysis is to extract, fit, and quantify the sinusoidal signals corresponding to multipole components of the local gravitational field. This is done by importing the data, cutting and averaging switch cycles, flattening and segmenting, fitting the segments, and performing statistical analysis of the fit results.

9.1 Data import

An import command prompts for the dataset filename and then generates a data structure within the current Octave workspace, reading data from the header and the following time series into corresponding fields within the structure. In addition to importing the time series given within the data file, new time series based on the imported data are created, *e.g.*, the difference-over-sum of the autocollimator signals and the square of the feedback voltage. The filename is also stored within the structure, and it or a comment field input by the user are used to refer to the run in plot legends, *etc.* Importing additional runs converts the data structure into an array of data structures, allowing each run to be referenced by its index in the analysis commands. Subsequent analysis commands store their output in the respective structures.

9.2 Cuts and averaging

For the gradient measurement the integral term is used as the science signal. At low frequencies the I-term is equivalent to the feedback signal, and the inherent low-pass filtering from the integral makes visual inspection of the signal straightforward, whereas inspection of the feedback signal would be hindered by high-frequency noise. For the analysis the information from each step is reduced to a single value. To do this the steps are separated ("chopped") according to the switching flag. Only full steps, *i.e.*, steps preceded and followed by a switch, are kept, and only data between switching routines is used. Because switching times are sometimes irregular due to camera resets, *etc.*, the length of chops could potentially vary. To prevent this the chops are limited to the length of the shortest period between switches, imposed by measuring backward from the start of the subsequent switch routine. The length of the chops is further limited to a user-specified value, again measured backward from the following switch, to exclude post-switch settling time from the data analysis. The result of this process is a number of chops one less than the number of switches in the run, with each chop consisting of an equal length of data immediately preceding a switch.

With earlier wheel designs, occasional mechanical disturbances in the switching produced some steps in which sufficient settling had obviously not occurred by the time the next switch occurs. To exclude these steps a cut is performed based on the standard deviation within each chop. The threshold is set very conservatively to only exclude chops in which settling has been significantly delayed. There is also a general cut on the autocollimator sum signal to ensure that the pendulum is on the detector for the entire chop. The average values of the integral term for the retained chops, along with their indices and the wheel angle, are passed along to the next analysis steps.

9.3 Drift subtraction and segmentation

In order to extract the sinusoidal signals from the data, the long-term drift is first removed. This is done by defining an "anchor angle" within each full rotation. The specific angle chosen has no significant effect on the results—the distribution of the resulting fits from all possible anchor angles are within the statistical uncertainty of the fits. A cubic-spline interpolation is made between subsequent anchor points and subtracted from the data. If an anchor point is missing from any cycle, *i.e.*, it was discarded due to failing one of the cuts, then the data between the preceding and following anchor points is discarded.

Following the drift subtraction the data is segmented into subsequent lengths corresponding to two full rotations, each beginning at an anchor point. This data, along with the wheel positions, is used by the fitting

routine.

9.4 Fitting and averaging

Each 48-point segment is fit for sine and cosine components at the fundamental rotation frequency and the second to fourth harmonics, as well as a DC offset, for a total of nine parameters. The fit routine is based on the `ols2` function from Charlie Hagedorn's thesis work [29], which is a modification of the `ols` function included in Octave and returns values and errors for each component of the fit.

After performing the fitting for each segment, the fits for each segment can be averaged, weighted by their respective errors, to produce an overall average fit for the run. The weighted mean β_w is produced by the formula

$$\beta_w = \frac{\sum \frac{\beta_n}{\sigma_n^2}}{\sum \sigma_n^{-2}}, \tag{9.1}$$

where β_n are the fit parameters and σ_n their errors. The weighted error σ_w is given by

$$\sigma_w = \sqrt{\frac{1}{\sum \sigma_n^{-2}} \frac{1}{N-1} \sum \frac{(\beta_n - \beta_w)^2}{\sigma_n^2}}. \tag{9.2}$$

The resulting values for the second harmonic are multiplied by a scaling factor to convert the feedback signal into a gradient measurement in SI units.

9.5 Calibration

Calibration is necessary to convert the values returned by the system into meaningful measurements. Our group's previous work has used different methods for calibrating systems either running with or without feedback. For systems without feedback the calibration can be done with two factors: one relating the signal returned by the autocollimator to the twist angle of the pendulum and another relating the twist angle to the torque on the pendulum. The autocollimator calibration can be obtained by using a calibrated adjustable mirror, and the twist-to-torque calibration factor (for signals at frequencies well below the pendulum resonance) is simply the torsional spring constant of the fiber. Given a known moment of inertia of the pendulum, this can be calculated by measuring the resonant frequency of the pendulum. The response function of an internally damped harmonic

oscillator,

$$\frac{\phi(\omega)}{\tau(\omega)} = \frac{1/\kappa}{(1-\omega^2/\omega_0^2) + i/Q}, \tag{9.3}$$

where $\phi(\omega)$ is the angular deflection, $\tau(\omega)$ the torque, κ the spring constant, ω and ω_0 the observed and resonant angular frequencies, and Q the mechanical quality factor of the fiber, can be used to back out the torque for a given deflection and frequency. The quality factor can be determined by measuring the decay of oscillation after a disturbance larger than the background noise level.

For our balances that operate in feedback, a different calibration relating the feedback signal to torque is required. This has been obtained for these balances by rotating a known mass arrangement around the pendulum at a fixed frequency and measuring the resulting feedback signal [10, 28]. Using the gravitational moments of the pendulum and the external masses the expected gravitational torque on the pendulum can be calculated. Together, the measured feedback signal and expected torque give the necessary conversion factor.

In the case of the WashCycle prototype, the location did not lend itself to installing a rotating mass arrangement. An alternative method sufficient for a proof-of-concept measurement was used instead. Extra code was added to the control software to allow the option of sinusoidally varying the feedback set point. This allowed the pendulum to be driven at an oscillation of set amplitude and frequency and was intended to provide a direct scaling between the feedback signal and torque.

Given that the nominal electrostatic force F_E between two parallel plates at a potential difference V is

$$F_E = \frac{\varepsilon_0 A V^2}{2d^2}, \tag{9.4}$$

where A is the overlapping area of the plates and d the distance between the plates, the torque exerted on the pendulum would be

$$\tau_E = -\frac{\varepsilon_0 A \ell V^2}{2d^2}, \tag{9.5}$$

where ℓ is the distance between the center of the feedback electrode and the pendulum's torsion axis, with the sign convention of movement toward the electrode being negative. The opposing torque τ_T from the torsion fiber as a function of displacement angle and using the same sign convention is

$$\tau_T(\phi) = \kappa(\phi_0 - \phi), \tag{9.6}$$

where κ is the torsional spring constant of the fiber and ϕ_0 the pretorque angle. Given the quadratic dependence on d and the difficulty of measuring it *in situ*, we meant to determine a scaling factor B between the feedback

signal and torque by determining the relationship between the observed twist angle and the feedback.

As mentioned above, we used the integral term as the science signal; this is valid because the feedback signal and the integral term are equivalent at low frequencies (frequencies below the ratio of the integral gain to the proportional gain). Hereafter feedback signal will refer to the integral feedback term, and the square of the feedback voltage will be replaced using $V^2 = K_S(S_0 + \Delta S)$, where $S = S_0 + \Delta S$ is the feedback signal, with S_0 equal to the feedback signal in steady state when $\phi = 0$, as measured by the autocollimator, and K_S the feedback gain with units of V^2. Hence

$$\kappa(\phi_0 - \phi) = B(S_0 + \Delta S), \tag{9.7}$$

which given the definition of S_0 reduces to

$$B = \frac{\kappa\phi}{\Delta S}. \tag{9.8}$$

Fitting to the calibration sinusoids in angle and feedback signal and using the known κ gives B, with the intention of using this to scale the feedback signal to torque and, using the pendulum's gravitational moment, the incident gravitation gradient.

However, after the science runs had concluded and the experiment decommissioned, it was determined that the calibration was incorrect. This error is due to the non-negligible dependence of the electrostatic torque on the displacement angle of the pendulum. A more accurate equation for the electrostatic torque is

$$\tau_E = -\frac{\varepsilon_0 A\ell V^2}{2(d + \ell\phi)^2} = -\frac{\varepsilon_0 A\ell K_S(S_0 + \Delta S)}{2(d + \ell\phi)^2}. \tag{9.9}$$

This breaks the linear relationship between ΔS and ϕ given above. Instead we solve for two different scaling factors: D, the scaling between ΔS and equilibrium angle, and H, the feedback signal to torque conversion ratio when $\phi = 0$.

We can solve for the first scaling factor by equating the torque from the torsion fiber with the electrostatic torque, as

$$\kappa\,(\phi_0 - \phi) = \frac{\varepsilon_0 A\ell K(S_0 + \Delta S)}{2(d + \ell\phi)^2}, \tag{9.10}$$

expanding to

$$2\kappa\,(\phi_0 d^2 + 2\phi_0 d\ell\phi + \phi_0\ell^2\phi^2 - d^2\phi - 2d\ell\phi^2 - \ell^2\phi^3) = \varepsilon_0 A\ell K_S(S_0 + \Delta S). \tag{9.11}$$

Given that S_0 is defined to be the necessary feedback signal in steady state when $\phi = 0$, the terms on each side

not containing ϕ and ΔS are equal, reducing the equation to

$$2\kappa\left(2\phi_0 d\ell\phi + \phi_0\ell^2\phi^2 - d^2\phi - 2d\ell\phi^2 - \ell^2\phi^3\right) = \varepsilon_0 A\ell K_S \Delta S. \tag{9.12}$$

With the assumption that $\phi \ll 1$, higher order terms in ϕ can be ignored, leaving

$$2d\kappa\phi\left(2\phi_0\ell - d\right) \approx \varepsilon_0 A\ell K_S \Delta S. \tag{9.13}$$

We can further simplify the approximation by recognizing that, given that the pretorque angle is known empirically to be of order 1, $2\phi_0\ell \gg d$. This gives the scaling factor

$$D = \frac{\phi}{\Delta S} = \frac{\varepsilon_0 A K_S}{4d\kappa\phi_0}. \tag{9.14}$$

This scaling can be extracted from the calibration routine, which also allows for a determination of d, as

$$d = \frac{\varepsilon_0 A K_S}{4\kappa\phi_0 D}. \tag{9.15}$$

The pretorque angle ϕ_0 is not known, but by solving the equation for the pretorque equilibrium

$$\kappa\phi_0 = \frac{\varepsilon_0 A\ell K_S S_0}{2d^2}, \tag{9.16}$$

it can be obtained in terms of the measured parameters and d, giving

$$d = 2D\ell S_0. \tag{9.17}$$

The desired conversion ratio H can then be obtained by substituting the value for d into the expected electrostatic torque from ΔS when $\phi = 0$, giving

$$H = \frac{\Delta\tau_E(\phi = 0)}{\Delta S} = -\frac{\varepsilon_0 A K_S}{8D^2\ell S_0^2}. \tag{9.18}$$

This factor H, obtained from the calibration routine result and other measured parameters, can be used along with the relevant quadrupole moment components of the weights to obtain the conversion ratio between the feedback signal and gravitational gradient. The quadrupole moment components, taking $\hat{\mathbf{y}}$ to be parallel to the

weight rotation axis as shown in Figure 7.1, can be approximated as

$$q_{xy} = \frac{1}{2} I_0 (1 + \cos 2\theta) \text{ and} \tag{9.19a}$$

$$q_{yz} = \frac{1}{2} I_0 \sin 2\theta, \tag{9.19b}$$

where I_0 is the moment of inertia of the weights about the fiber axis when the weight axis is horizontal and θ is the angle of the weight axis with respect to the horizontal. The amplitude of the θ-dependent variation in these quadrupoles is the same: $I_0/2$.

The torque from a gravitational gradient Γ_{ij} is simply $\tau_G = -q_{ij}\Gamma_{ij}$, and so the feedback signal to gradient conversion factor M_{ij} when the pendulum is locked in feedback at $\phi = 0$ is given by

$$M_{ij} = \frac{\Gamma_{ij}}{\Delta S} = \frac{H}{q_{ij}}. \tag{9.20}$$

The torque from the two gradient components of interest are orthogonal in θ, having the same angular periods of 180° and being out of phase by 45°. They correspond with the fitted $\sin 2\theta$ and $\cos 2\theta$ terms resulting from the fit described in the previous section. Using the amplitude of the θ-dependence for the moments given in Equation 9.19a and Equation 9.19b gives the final scaling from fitted amplitudes to gravitational gradient:

$$M_{xy} = M_{yz} = M = \frac{2H}{I_0}. \tag{9.21}$$

The values for all measured and derived parameters used in this section, based on a calibration run taken after the science campaign, are given in Table 9.1.

Parameter	Value	Description	Source
κ	150 nN m/rad	torsional spring constant of fiber	free-run measurement
I_0	140 mg m^2	moment of inertia of weights at $\theta = 0$	CAD model
A	35 cm^2	area of feedback electrode	direct measurement
ℓ	44 mm	radial distance to center of feedback electrode	direct measurement
K_S	2350 V^2	feedback signal gain	adjustable factor
B	6.9 nN m	*incorrect* feedback signal to torque scaling factor	calibration fit
D	46.0 mrad	feedback signal to equilibrium angle scaling factor	calibration fit
S_0	0.83	steady-state feedback signal at $\phi = 0$	measurement
d	3.4 mm	electrode-shield separation	derived
ϕ_0	790 mrad	pretorque angle	derived
H	-140 nN m	feedback signal to torque at $\phi = 0$	derived
M	-2.0 MEo	2θ signal to gravitational gradient	derived

Table 9.1: Feedback calibration parameters.

Chapter 10

Results

Although the goal of this project was the development of a low-cost gravity gradiometer, an intermediate and necessary goal for the chosen solution was the identification or development of a rotary actuator that could be used on the pendulum. A few key successes along with the results of a final measurement run with the gradiometer are discussed here.

10.1 Successes

10.1.1 Actuator reliability

After producing the initial vacuum-compatible pendulum and operating it for a number of days, the actuator would get stuck. Three modifications were needed to improve the actuator's reliability to acceptable levels.

Since it could be observed in post-failure analysis that the throw of the slider had decreased and that the tension in the Nitinol wire was reduced, the first modification was to improve the attachment of the wire to the pendulum body. This was done by replacing the nylon screws used to clamp the wire with brass screws, electrically isolated as described in section 7.3, due to a suspicion that the nylon screws were not capable of providing sufficient clamping force on the wire.

After improving the clamping of the wire, failures still occurred, and we suspected that the wire was irrecoverably stretching itself when the slider pin reached the valley of the inner wall. We first addressed this by cutting relief slots into the wheel , but we found this to cause too much of a disturbance to the pendulum when the pin slipped over the corner into the slot. This design was replaced by one which extended the "downward" slope farther than was necessary for cycling, as shown in Figure 7.7.

A final change to improve actuator reliability involved the slider and an education in vacuum-compatible lubricants. The original brass slider appeared to be binding in its bore, despite the application of powdered graphite to the sliding surfaces. After some discussion and research, we discovered that graphite is known to not be a vacuum-compatible lubricant, as it requires water to be present to allow the the crystal planes to slide against each other [30]. (We later discovered that graphite is also known to cause corrosion in aluminum in the presence of water [31].) Molybdenum disulfide ("dry moly") lubricant was substituted but appeared to exhibit the same behavior, although to a lesser degree. Finally, the brass slider was replaced with an acetal slider, which has not exhibited similar issues.

After these modifications, the rotary actuator has been able to operate continuously for multiple weeks without missing a step, making thousands of steps and over one hundred full rotations. The run of successful rotations appeared to be curtailed only by the rising gas pressure in the valved-off vacuum chamber (see section 10.1.3), which increased the thermal conduction away from the Nitinol and prevented it from making full contractions at the preset power level.

10.1.2 Feedback lock

Just as important as the reliability of the actuator is the reliability of the feedback system in holding the pendulum in lock for a sufficient amount of time to take a measurement in each position. As mentioned in section 8.1.3, powering the pendulum produces a torque on the pendulum greater than the maximum torque of the feedback electrode. This was a problem because it could push the pendulum out of feedback and off the detector, causing delays as the pendulum was caught and potentially causing collisions with the electrode or the bumper installed opposite the electrode, possibly disturbing the fiber and changing the equilibrium torque. We attempted various schemes to compensate for this torque such as releasing the pendulum from feedback before powering the coil, to allow it to build some momentum opposite to the coil torque, or pinning the pendulum to the electrode with a large feedback potential before powering the coil, to prevent it from building momentum that would result in a collision. In the end, Jens Gundlach suggested adding a coil that could push back against the pendulum. This resulted in the anticoil described in the earlier chapter. The anticoil is powered at a set frequency and amplitude while the coil is powered. Although it should be possible to integrate the anticoil into the feedback loop, simply tuning the drive signal to match the torque from the coil was sufficient (and a lack of available hardware at the time precluded such integration).

Using the anticoil, the pendulum stays on the detector for the whole switching cycle. As shown in Figure 10.1, it is not fully locked in place while the coil is powered and the wheel is moving, but the excursions are smaller

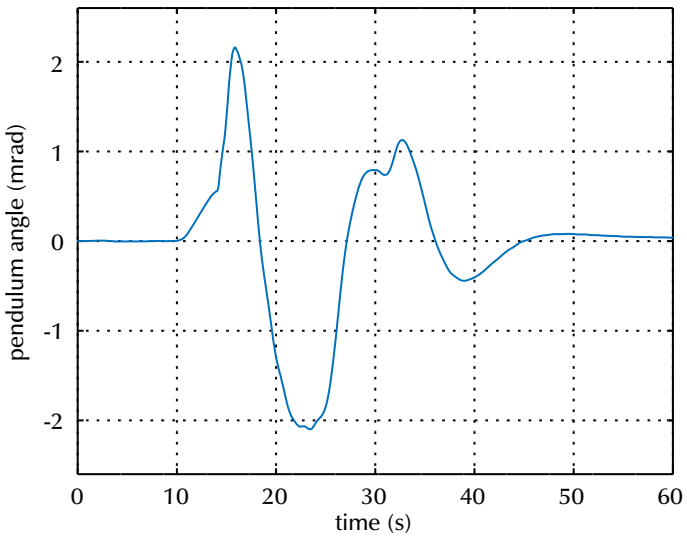

Figure 10.1: A plot of the pendulum angle during a switching cycle. The excursions cover just over half of the autocollimator range. Features of the trace correspond with coil activation and pin and wheel movement.

than the range of the detector and the feedback system quickly recovers after the coil is turned off. The magnitude of the excursions depend on the wheel position, probably due to the changing moment of inertia and possibly due to a changing center of mass. In a few instances (about 30 in the 3300 switches used in the results) the excursions did go outside the range of the detector but most were small enough that the system still recovered. The others resulted in the effective end of that run. Many of the excursions happened in a single run, all at the same wheel positions, which points to the anticoil being mistuned during that run.

10.1.3 Reduced cycle time

Once we were able to make the switching and feedback reliable enough for continuous data taking, the issue became one of time. We are fighting against $1/f$ noise, which results in the measurement sensitivity becoming worse as the signal frequency is decreased [32]. Measuring at a higher signal frequency, corresponding to faster rotation of the wheel, should give better sensitivity.

The coil is only powered for 15 s for each switch, but the cycle must be longer than that to allow for the Nitinol to cool and relax. At lower vacuum pressures cooling is dominated by conduction into the pendulum body and with a turbopump running relaxation of the Nitinol took about 200 s. Relaxation time was decreased to about 60 s by adding PTFE bumpers that press against the wire at two points along its length. The larger surface area of the pendulum allows heat to escape at a reasonable rate through the residual gas. Even so, with

the turbopump on, cycle times below a certain threshold allow the temperature of the pendulum to rise above the transition point of the Nitinol and prevent the Nitinol from relaxing. A cycle time of 400 s allows for enough cooling of the pendulum for continuous cycling to occur.

We observed a consistent drift in the feedback signal after the pendulum is switched. This appears to be a thermal effect, possibly due to changing properties of the post-hanger. In addition to allowing sufficient time for the Nitinol to relax and the pendulum to cool, a longer cycle time also allows for a reduction in the drift prior to the next switch, as the drift rate drops exponentially along with the drift.

To reduce the cooldown time and allow for faster settling of the feedback we allowed the pressure in the vacuum chamber to increase, initially by turning off the turbopump and only using the roughing pump and finally by valving off the chamber from the roughing pump and not pumping on it at all. Given the mass of the pendulum, thickness of the torsion fiber, and magnitude of the signals of interest, the noise due to gas damping of the pendulum is not significant, which was shown by comparing the noise spectrum at different pressures. Running at increased pressures reduced the Nitinol relaxation time to 15 s, but the drift in the signal is still substantial. The average change in the feedback signal for the final 15 s (360 s after the power input is turned off) corresponds to a gravitational gradient drift of 3 Eo/s. Given this large residual drift, variations between cycles in total heat delivered via the coil or in the cooldown rate could potentially have a significant effect on the measurement, either in systematic effects or decreased sensitivity.

10.2 Gradient measurement

10.2.1 Test signal

To test the system's ability to measure a gravitational gradient, we built a support system with block-and-tackle for a 25-kg weight, which allowed for easy positioning of the weight in two positions along the same vertical axis and at roughly equal distances above and below the horizontal plane passing through the pendulum wheel . The weight was suspended by a crossbeam supported at both ends. The two positions of the weight were at roughly equal vertical distances above and below the center of the pendulum weights. For the gravitational gradient at the pendulum's location, switching the weight between these two positions has the effect of flipping the sign of the portion of the Γ_{yz} component due to the weights while leaving the Γ_{xy} component unchanged.

The actual change in the relevant gradient component, Γ_{yz}, at the pendulum's location due to moving the weight can be easily calculated. With weight of mass m and center of mass $\mathbf{x}_0 = (x_0, y_0, z_0)$, the gravitational

Parameter	Value	Description
m	25 kg	mass of test weight
ΔZ	18 cm	vertical displacement distance (from center)
R	42 cm	radial distance
ϕ	57°	angular position
$\Delta\Gamma_{yz}$	−32 Eo	calculated change in the gravitational gradient

Table 10.1: Test signal parameters.

acceleration at position $\mathbf{x}$ due to the weight is given by

$$\mathbf{g}(\mathbf{x}, \mathbf{x}_0) = -Gm\frac{\mathbf{x}_0 - \mathbf{x}}{|\mathbf{x}_0 - \mathbf{x}|^3} = -Gm\frac{(x - x_0)\hat{\mathbf{x}} + (y - y_0)\hat{\mathbf{y}} + (z - z_0)\hat{\mathbf{z}}}{\left[(x - x_0)^2 + (y - y_0)^2 + (z - z_0)^2\right]^{3/2}}. \tag{10.1}$$

To get the gradient component Γ_{yz}, we must take either the y-derivative of the z-component or *vice versa*, giving

$$\Gamma_{yz}(\mathbf{x}, \mathbf{x}_0) = \frac{\partial}{\partial y}g_z(\mathbf{x}, \mathbf{x}_0) = \frac{\partial}{\partial z}g_y(\mathbf{x}, \mathbf{x}_0) = -3Gm\frac{(y - y_0)(z - z_0)}{\left[(x - x_0)^2 + (y - y_0)^2 + (z - z_0)^2\right]^{5/2}}. \tag{10.2}$$

Taking the center of the pendulum to be at the origin and the upper and lower weight positions of $\mathbf{x}_\pm = (R\cos\phi, -R\sin\phi, \pm\Delta Z)$ in the Cartesian coordinates given in Figure 7.1, switching the weight between the two positions gives a change of the considered component of the gravitational gradient at the pendulum as

$$\Delta\Gamma_{yz}(\mathbf{x} = 0) \ = \Gamma_{yz}(\mathbf{0}, \mathbf{x}_+) - \Gamma_{yz}(\mathbf{0}, \mathbf{x}_-) = -6Gm\frac{\Delta Z R\sin\phi}{(R^2 + \Delta Z^2)^{5/2}}. \tag{10.3}$$

The measured parameters and resulting gradient change are given in Table 10.1.

10.2.2 Measurement results

After iterative improvements we reached a point where the torque noise when switching was comparable to the noise without switching, as shown in Figure 10.2. We then conducted a series of science runs, with the test weight position alternating for each run. This campaign resulted in thirteen runs with usable data before the actuator failed. (Some runs were unusable due to the pendulum being knocked out of lock, presumably due to activity in and around the room housing the setup.)

The resulting measurement for each usable run are shown in Figure 10.3. We can perform a fit against the weight position, a constant offset, and a linear drift in time, which results in the values shown in Table 10.2 and

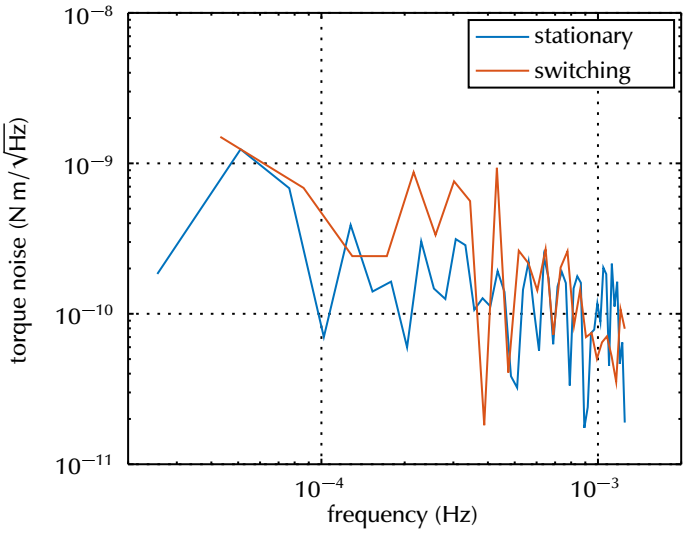

Figure 10.2: A comparison of the torque noise for the averaged cuts from two runs with the pendulum wheel stationary and switching. A signal at the full-rotation frequency (104 µHz) in the switching run was fit for and removed before taking the FFT.

Fit measurement	Value
$\Delta\Gamma_{yz}$	-55 ± 11 Eo
drift	2.1 ± 0.3 Eo/d
offset (weight centered)	-87 ± 6 Eo

Table 10.2: Measurement fit with drift and offset.

the lines shown in the plot. We resolved a signal with the expected sign and a magnitude slightly more than two standard deviations away from the expected value. This discrepancy could be due to measurement noise, miscalibration, or a drift in calibration.

The limiting noise in the apparatus may be due to floor tilt, which could be coupling to the pendulum's twist through the feedback system. Tilting the apparatus about the horizontal axis parallel to the electrode surface changes the pendulum-electrode spacing and should result in apparent torque noise. The magnitude of the tilt-to-torque coupling is set by the spacing and the pretorque. (The expected torque coupling for tilts about the axis perpendicular to the electrode surface, resulting in a change in the effective lever arm of the feedback electrode, is smaller by the ratio of lever arm to electrode separation—about an order of magnitude.) With a vertical distance L between the top fiber attachment point and the center of the electrode and floor tilt $\Delta\eta$, the

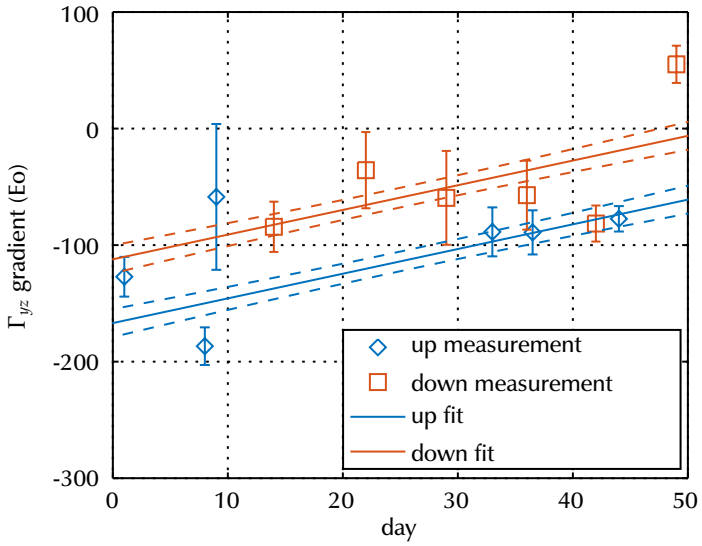

Figure 10.3: The measurements from the science campaign shown with 1-σ error bars. The position of the test weight was alternated for each measurement run, but some runs were excluded due to loss of feedback lock.

change in torque $\Delta\tau$ is given by

$$\tau_0 + \Delta\tau = -\frac{\varepsilon_0 A \ell K_S S_0}{2(d + L\Delta\eta)^2}, \tag{10.4}$$

with

$$\tau_0 = \kappa\phi_0 = -\frac{\varepsilon_0 A \ell K_S S_0}{2d^2}. \tag{10.5}$$

Assuming $L\Delta\eta \ll d$ gives

$$\Delta\tau \approx -\frac{2L\Delta\eta}{d}\kappa\phi_0. \tag{10.6}$$

With $L = 56\,\text{cm}$ and the additional values given in Table 9.1 the tilt-to-torque coupling factor is expected to be $-39\,\mu\text{N m/rad}$. Applying this factor to tilt noise collected by a tiltmeter (AGI 755) clamped to the vacuum chamber's baseplate gives the the resulting expected torque noise shown in Figure 10.4, overlaid with the feedback torque spectrum taken at the same time, with the pendulum wheel stationary. The effective torque sensitivity limit of the autocollimator, given the fiber spring constant, and a thermal noise limit for $Q = 100$ [32] are shown for comparison. Suggestions for addressing the suspected tilt coupling are included in the next chapter.

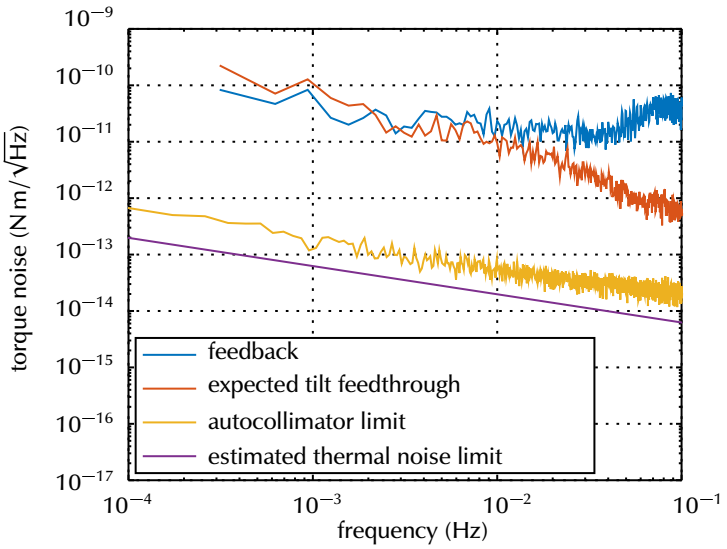

Figure 10.4: Comparison of the feedback torque, expected torque due to the tilt feedthrough, equivalent torque limit of the autocollimator, and an estimated thermal noise limit for $Q = 100$.

Chapter 11

Recommended next steps

The work to this point has served as a demonstration of the feasibility of implementing an active-mass torsion pendulum in a scientific application and also the applicability of such a pendulum to gravity gradient measurement. The observed noise in the device makes it unsuitable for the intended gradient monitoring, but the actuator design is immediately usable for occasional mass repositionings, and minor upgrades may produce a device that can be used for the motivating application. This chapter explains some of the changes and work needed to use such an apparatus as a gravity-gradient-monitoring system and to improve its sensitivity.

11.1 Measuring the full Γ_{21} gradient

As explained in section 1.1, the gravity gradient tensor in free space can be fully specified with five real scalar components. In Cartesian coordinates these correspond to the Γ_{xy}, Γ_{xz}, Γ_{yz}, and any two of Γ_{xx}, Γ_{yy}, and Γ_{zz}. Two of these, Γ_{xz} and Γ_{yz}, are related to the real and imaginary components of the Γ_{21} moment of the gravity gradient as expressed in spherical harmonics, which is of primary interest in terms of systematic errors in a measurement using a rotating torsion balance. The active-mass torsion balance prototype is only able to measure half of the Γ_{21}, since it measures only Γ_{xy} and Γ_{yz} (where $\hat{\mathbf{y}}$ is parallel to the weight rotation axis). One way to measure the full Γ_{21} would be to make a measurement with a second pendulum oriented perpendicularly to the first, but such an arrangement will not be measuring the two components at the same point. An alternative would be to place the apparatus on a simple turntable that cycles it between the two required orientations on a timescale longer than the wheel-rotation period.

In addition to providing a way to measure the full Γ_{21} moment, using this turntable would also allow for

measurements in the two opposite orientations, for a total of four different orientations separated by 90°. This would provide sufficient information to remove the gravitational effect of components of the torsion balance system on the measurement, giving a better measurement of the absolute gravitational gradient rather than just tracking changes in the relative gradient. This would only remove the effect of system components on the turntable, but since those would be the objects closest to the pendulum they would have a larger effect on the measurement than any stationary components.

Although the problem of measuring the full Γ_{21} moment at one location can be solved by rotating the apparatus, a similar problem lies in measuring the gravitational gradient at the actual point of interest, *e.g.*, the location of a science pendulum. This problem can be addressed to some extent by arranging a set of gradiometers around the point of interest and interpolating the value at the center. This would also solve the above problem of acquiring the full Γ_{21}, as its two components could simply be interpolated by separate pairs of gradiometers.

11.2 Recommended modifications and upgrades

A number of changes could be made to make deployment of an array of instruments more practical and effective, including reductions in size and cost, improved self-monitoring, and continuous calibration.

The system is currently housed in a large stainless steel bell jar. While this arrangement is convenient for prototyping, it is not practical either to purchase or to surround an experiment with a number of these. The bell jar could reasonably be replaced with an 8-in vacuum nipple, with a reducing flange and narrower nipple on the top to support the torsion fiber.

The coil, anticoil, and feedback electrode are currently driven using a precision function generator and three precision amplifiers. Only the feedback electrode needs a precision supply, and it does not need a high-current supply, though it does presently require a higher voltage than the DAQ card can provide. The coil drivers could be replaced with simple switching circuits on DC supplies, which should be properly protected against the inductive loads.

The reliability of the actuator could be improved either by lengthening the pendulum body to allow for a longer length of Nitinol or by reducing the wheel diameter. Reducing the wheel diameter will require reducing the width of the groove and the diameter of the pin, so the stiffness of the pin material sets a lower bound on the wheel size. Addition of a computer vision system to monitor the wheel angle and provide alerts when the wheel becomes stuck would be straightforward and would increase confidence in the system performance.

For precision monitoring an ongoing or periodic calibration would be warranted. Although the expected

calibration could be monitored indirectly, through the same calibration method already implemented here or by monitoring the pendulum-electrode separation, a true gravitational calibration would be more robust against changes. This could be implemented using a mass multipole rotated or repositioned at a frequency above the rotation rate of the wheel. Care would have to be taken to not decrease the sensitivity of the balance by introducing extra noise.

11.3 Improving the sensitivity

As discussed in section 10.2.2, the sensitivity of the prototype appears to be limited by the tilt-to-torque coupling due to the feedback electrode. As shown there, the coupling is proportional to the pretorque, and the pretorque was set to be much higher than in our other balances in order to give the single-electrode feedback enough dynamic range to keep control of the pendulum when the wheel was powered. The anticoil was added to assist in controlling the pendulum, but the pretorque was kept at the same level. It is possible that the system was capable of maintaining control of the pendulum at a lower pretorque, but since it was decommissioned prior identifying tilt coupling as the likely limiting factor, we cannot now test this.

If the system were to be recommissioned, the next step would be to investigate whether the noise level is a function of the pretorque. This could be done without the wheel turning. If this were found to indeed be the case, efforts should be made to both decrease the low-frequency tilt and to upgrade the anticoil control system to allow for proportional feedback, rather than just using a set frequency and amplitude for each switching routine. If this were still not sufficient, a second anticoil could be installed to provide an opposing torque, forming a bidirectional high-gain feedback system used only during switching. In tandem with these measures, a continued investigation of the root cause of the disturbance necessitating the anticoil may be fruitful.

The other known noise sources are the thermal noise inherent in the pendulum system and the autocollimator noise. Estimated torque spectrums for these are shown in Figure 10.4. Converting into gradient sensitivities using the proper scaling factor gives an autocollimator limit of 5 Eo/$\sqrt{\text{Hz}}$ with the same 400-s switching time. This could be improved by a factor of ~40 by using one of our group's multi-slit autocollimators [15]. The sensitivity can also be improved by reducing the mass of the pendulum, which would allow for a smaller diameter and weaker fiber, and by increasing the wheel's gravitational quadrupole by moving the masses farther from the center.

References

[1] J. D. Jackson, *Classical Electrodynamics* (Wiley, 1998).

[2] *Basic measurement units*, European Space Agency, http://www.esa.int/Our_Activities/Observing_the_Earth/The_Living_Planet_Programme/Earth_Explorers/GOCE/Basic_measurement_units (visited on 02/21/2017).

[3] Analysis and decision: application number 3020291 (Seattle Department of Construction and Inspections, 2016), http://www.seattle.gov/dpd/LUIB/AttachmentProject3020291ID72123020291.pdf.

[4] 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**, 041101 (2008) 10.1103/PhysRevLett.100.041101.

[5] Y. Su, B. R. Heckel, E. G. Adelberger, J. H. Gundlach, M. Harris, G. L. Smith, and H. E. Swanson, "New tests of the universality of free fall", Phys. Rev. D **50**, 3614–3636 (1994) 10.1103/PhysRevD.50.3614.

[6] L. Eötvös, "A Föld alakjának kérdése [The problem of the shape of the earth]", Természettudományi Közlöny **33**, 321–28 (1901), http://mek.oszk.hu/03200/03286/html/eotvos1/foldalak.html, translated quote from https://en.wikipedia.org/wiki/Lorand_Eotvos.

[7] M. I. Evstifeev, "The state of the art in the development of onboard gravity gradiometers", Gyroscopy and Navigation **8**, 68–79 (2017) 10.1134/S2075108717010047.

[8] D. J. Kapner, T. S. Cook, E. G. Adelberger, J. H. Gundlach, B. R. Heckel, C. D. Hoyle, and H. E. Swanson, "Tests of the gravitational inverse-square law below the dark-energy length scale", Phys. Rev. Lett. **98**, 021101 (2007) 10.1103/PhysRevLett.98.021101.

[9] W. A. Terrano, E. G. Adelberger, J. G. Lee, and B. R. Heckel, "Short-range, spin-dependent interactions of electrons: a probe for exotic pseudo-goldstone bosons", Phys. Rev. Lett. **115**, 201801 (2015) 10.1103/PhysRevLett.115.201801.

[10] C. A. Hagedorn, "A sub-millimeter parallel-plate test of gravity", PhD thesis (University of Washington, 2015), http://hdl.handle.net/1773/34135.

[11] W. Kliemann, and S. Namachchivaya, *Nonlinear dynamics and stochastic mechanics* (Taylor & Francis, 1995).

[12] G. L. Smith, C. D. Hoyle, J. H. Gundlach, E. G. Adelberger, B. R. Heckel, and H. E. Swanson, "Short-range tests of the equivalence principle", Phys. Rev. D **61**, 022001 (1999) 10.1103/PhysRevD.61.022001.

[13] R. V. Jones, "Some developments and applications of the optical lever", Journal of Scientific Instruments **38**, 37–45 (1961) 10.1088/0950-7671/38/2/301.

[14] R. Cowsik, R. Srinivasan, S. Kasturirengan, A. Senthil Kumar, and K. Wagoner, "Design and performance of a sub-nanoradian resolution autocollimating optical lever", in APS meeting abstracts (Apr. 2006), http://meetings.aps.org/link/BAPS.2006.APR.L11.3.

[15] T. B. Arp, C. A. Hagedorn, S. Schlamminger, and J. H. Gundlach, "A reference-beam autocollimator with nanoradian sensitivity from mHz to kHz and dynamic range of 10^7", Review of Scientific Instruments **84**, 095007–095007 (2013) 10.1063/1.4821653.

[16] P. B. Dixon, D. J. Starling, A. N. Jordan, and J. C. Howell, "Ultrasensitive beam deflection measurement via interferometric weak value amplification", Phys. Rev. Lett. **102**, 173601 (2009) 10.1103/PhysRevLett.102.173601.

[17] D. J. Starling, P. B. Dixon, A. N. Jordan, and J. C. Howell, "Optimizing the signal-to-noise ratio of a beam-deflection measurement with interferometric weak values", Phys. Rev. A **80**, 041803 (2009) 10.1103/PhysRevA.80.041803.

[18] J. C. Howell, D. J. Starling, P. B. Dixon, P. K. Vudyasetu, and A. N. Jordan, "Interferometric weak value deflections: quantum and classical treatments", Phys. Rev. A **81**, 033813 (2010) 10.1103/PhysRevA.81.033813.

[19] Y. Aharonov, D. Z. Albert, and L. Vaidman, "How the result of a measurement of a component of the spin of a spin-1/2 particle can turn out to be 100", Physical Review Letters **60**, 1351–1354 (1988) 10.1103/PhysRevLett.60.1351.

[20] Y. Aharonov, S. Popescu, and J. Tollaksen, "A time-symmetric formulation of quantum mechanics", Physics Today **63**, 27–32 (2010) 10.1063/1.3518209.

[21] N. W. M. Ritchie, J. G. Story, and R. G. Hulet, "Realization of a measurement of a 'weak value'", Phys. Rev. Lett. **66**, 1107–1110 (1991) 10.1103/PhysRevLett.66.1107.

[22] A. M. Steinberg, "Quantum measurement: a light touch", Nature **463**, 890–891 (2010) 10.1038/463890a.

[23] O. Hosten, and P. Kwiat, "Observation of the spin Hall effect of light via weak measurements", Science **319**, 787– (2008) 10.1126/science.1152697.

[24] C. A. J. Putman, B. G. de Grooth, N. F. van Hulst, and J. Greve, "A detailed analysis of the optical beam deflection technique for use in atomic force microscopy", Journal of Applied Physics **72**, 6–12 (1992) 10.1063/1.352149.

[25] F. Higgs, *A quick test of the printable stepper thrust plate*, (2009) https://vimeo.com/3169018.

[26] G. B. Kauffman, and I. Mayo, "The story of Nitinol: The serendipitous discovery of the memory metal and its applications", The Chemical Educator **2**, 1–21 (1997) 10.1007/s00897970111a.

[27] E. G. Adelberger, C. W. Stubbs, B. R. Heckel, Y. Su, H. E. Swanson, G. Smith, J. H. Gundlach, and W. F. Rogers, "Testing the equivalence principle in the field of the Earth: Particle physics at masses below 1 μeV?", Phys. Rev. D **42**, 3267–3292 (1990) 10.1103/PhysRevD.42.3267.

[28] S. Schlamminger, C. Hagedorn, M. Famulare, S. Pollack, J. Gundlach, S. M. Merkowitz, and J. C. Livas, "High sensitivity torsion balance tests for LISA proof mass modeling", in AIP conference proceedings, Vol. 873, 1 (2006), pp. 151–157.

[29] C. Hagedorn, *The PlateWash analysis repo*. https://github.com/4kbt/PlateWash/.

[30] D. Ramadanoff, and S. W. Glass, "High-altitude brush problem", Transactions of the American Institute of Electrical Engineers **63**, 825–830 (1944) 10.1109/T-AIEE.1944.5058806.

[31] *Acceptable methods, techniques, and practices: Aircraft inspection and repair*, Advisory circular 43.13-1B (United States Federal Aviation Administration Regulatory Support Division, 1998), https://www.faa.gov/regulations_policies/advisory_circulars/index.cfm/go/document.information/documentID/99861.

[32] P. R. Saulson, "Thermal noise in mechanical experiments", Phys. Rev. D **42**, 2437–2445 (1990) 10.1103/PhysRevD.42.2437.

Typeset using pdfLaTeX with the following key packages:

- Font: `epigrafica`
- Automatic handling of SI units: `siunitx`
- Optical schematics: `pst-optexp`
- References: `biblatex` with options `backend=biber, style=phys, articletitle=true, biblabel=brackets, doi=true, url=true`.

LaTeX-based plots output from Octave using `print -dpdflatex`.