

FUNDAMENTAL AND APPLIED PHYSICS
WITH OPTICALLY LEVITATED MICROSPHERES

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

Charles Blakemore

August 2021

© 2021 by Charles Preston Blakemore. All Rights Reserved.
Re-distributed by Stanford University under license with the author.



This work is licensed under a Creative Commons Attribution-3.0 United States License.

<http://creativecommons.org/licenses/by/3.0/us/>

This dissertation is online at: <https://purl.stanford.edu/wk886kt8795>

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

Giorgio Gratta, Primary Adviser

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

John Fox

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

Leo Hollberg

Approved for the Stanford University Committee on Graduate Studies.

Stacey F. Bent, Vice Provost for Graduate Education

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

Preface

In this collection of work, I will first guide the reader toward an understanding of optical tweezers, a technique used to trap small refractive particles, such as glass microspheres or living cells. In 2019, Arthur Ashkin was awarded part of the Nobel Prize in Physics for his development of this technique, and he made many of the initial strides in the field.

With a general understanding of optical tweezers, I will discuss the specific setup and implementation constructed in our group at Stanford, especially those aspects that are particularly unique or novel to the apparatus. All of the work presented has been performed with optically trapped silica microspheres in vacuum, using a single-beam trapping technique with interferometric readout.

I will then move to a discussion of a number of technical developments with metrological applications including: force field microscopy, direct measurements of a levitated microsphere's mass and radius, control of the microsphere's rotational degrees of freedom, and residual gas sensing.

To conclude, I will explore some of the fundamental science goals of the project, including results from searches for fifth forces related to both dark energy and more general modifications to Newtonian gravity that might arise from new physics that manifest in the submillimeter regime.

Acknowledgments

Properly acknowledging every individual who has contributed to the completion of this PhD is likely impossible. It takes a village, as the saying goes. However, there are a number of people who have made incredible and noteworthy contributions, and it would be remiss not to mention them.

Not unlike most undergraduates, I didn't have a clear vision of what graduate school was really like. Coming to Stanford, two individuals really took it upon themselves to help my transition and prepare me for the endeavor that would dominate the next 6 years of my life. As an already established graduate student, Alexander D. Rider was not only a superb resource for information and advice, but also my primary teacher, research partner, and eventually, friend. I think it would be hard to separate our contributions to some of the work detailed here, as we were so critically involved in each other's work.

In a more senior role as a post-doctoral scholar, David C. Moore was like a combined encyclopedia and technical manual. Any question (and there were many) I had from fundamental physics, to statistics, to coding, Dave was willing and able to answer. I credit my approach to research and problem-solving in large part to Dave's instruction and the example he set for me.

Continuing up the food chain, we naturally arrive at Giorgio Gratta, my PhD advisor. Many things can be said about Giorgio, both as a scientist and as a person, but they all lead back in some way to an unceasing intellectualism and curiosity. It's hard for me to imagine a moment in time where Giorgio wasn't thinking about either physics, or a clearly first-principled viewpoint of some non-physics subject matter. As a well-established and successful faculty member, he was a superb role-model and resource during my time in graduate school, and I feel privileged to have had the opportunity to work with him.

Although Alex, Dave, and Giorgio were the people I interacted with most extensively,

my days at Stanford were filled with an uncountable number of conversations with colleagues. Particularly fruitful were my regular conversations with fellow group members on all subjects of physics, from design and construction of actual experiments, to fundamental theories, to statistical analysis, to scientific programming and data presentation, and much much more. The list is likely incomplete, but proceeding alphabetically: Evan Angelico, Jacopo Dalmasson, Alexander Fieguth, Dan Fudenberg, Emmett Hough, Mike Jewell, Akio Kawasaki, Scott Kravitz, George Laskaris, Brian Lenardo, Gaosong Li, Denzal Martin, Igor Ostrovskiy, Nadav Priel, Sandip Roy, Qidong Wang, Manuel Weber, and Shuoxing Wu.

Outside of our group, many aspects of my research tenure were made possible by amazing people all over Stanford. John Fox, with whom I served as a teaching assistant on multiple occasions, taught me a massive amount of material related to electronic circuit design, both analog and digital. In the Physics Machine Shop, Karlheinz Merkle and Mehmet Solyali (Mehmet in particular), helped teach me best practices for working in the shop, in addition to being available for regular discussions about machining topics in general. In administrative roles, Maria Frank, Marcia Keating, and Sha Zhang facilitated passage through the various bureaucratic hoops that Stanford maintains.

In the Stanford Nanofabrication Facility, Ludwig Galambos, Phil Himmer, Usha Raghuram, Maurice Stevens, and Uli Thumser were invaluable resources for process development and tool operation, while much of the work would have been impossible without the tireless efforts of the maintenance staff consisting of Elmer Enriquez, Jim Haydon, Gary Sosa, and Mario Vilanova. In the Stanford Nano Shared Facilities, Tom Carver and Richard Chin were very useful in relaying their extensive knowledge of evaporation and scanning electron microscopy, respectively.

Outside of the academic sphere, friends and family were perpetually available as emotional support and sanity maintenance, without which I would have been lost long ago. This is especially true of the unwavering support I received from my superb parents Lisa and Page Blakemore, and I credit much of my own success to my upbringing in their household, and the examples they set for me.

Contents

Preface	v
Acknowledgments	vii
1 Optical Tweezers	1
1.1 Ashkin and the First Optical Tweezer	1
1.2 Scattering of Light	3
1.3 Mie Scattering - The T -Matrix Method	10
1.4 Optical Tweezers as Force Sensors	15
1.5 Vertically Oriented Optical Tweezers	16
1.6 Imaging of Microsphere Motion: A Heterodyne Approach	23
1.7 Active Feedback Primer	29
2 The Optical Levitation Apparatus at Stanford	31
2.1 The Experimental Chamber	32
2.2 Source Light and Input Optics - Single-beam Trapping	36
2.3 The Optical Trap	41
2.4 Output Optics and Imaging	49
2.5 Auxiliary Optics	57
2.6 Microsphere Preparation	63
2.7 Feedback and Calibration	70
3 Three-dimensional Force-field Microscopy	87
3.1 Introduction	87
3.2 Experimental Setup	88
3.3 Sensor linearity and noise	90

3.4	Electric field from an overall bias voltage	94
3.5	Patch potential measurements	96
3.6	Conclusion	99
4	Direct Mass and Radius Measurements of Individual Microspheres	101
4.1	Introduction	101
4.2	Experimental Setup	102
4.3	Electrostatic Co-Levitation	103
4.4	Radius Measurements from SEM	109
4.5	Discussion	112
4.6	Conclusion	115
5	Electrically Driven Rotation and Residual Gas Sensing	117
5.1	Introduction	117
5.2	Experimental Setup	119
5.3	Libration	122
5.4	Gyroscopic precession	123
5.5	Torque Noise and Torsional Drag	125
5.6	Release from a Spinning Field	127
5.7	Equilibrium Phase Lag	133
5.8	Gas Composition and Changes in Dipole Moment	137
5.9	Capacitance Manometer Cross-Check	138
5.10	Non-ionizing Gas Analyzer	141
5.11	Conclusion	142
6	Beyond the Standard Model	145
6.1	The Problem With Gravity	145
6.2	Parameterizing New Forces	147
6.3	A Few Examples	149
6.4	Previous Searches for New Interactions	157
6.5	Extending the Experimental Reach	164
7	Search for Screened Interactions: The Chameleon Field	167
7.1	Introduction	167
7.2	Experimental Setup	168

7.3	Dipole Response Measurement	170
7.4	Probing for Chameleons	172
7.5	Conclusion	177
8	Search for non-Newtonian Interactions that Couple to Mass	179
8.1	The Experiment	180
8.2	Analysis	184
8.3	Results	190
8.4	Conclusion	193
	Conclusions	195
	Summary for Non-Scientists	203
A	The T-Matrix Method	209
B	Harmonic Oscillators	213
B.1	The Equation of Motion	213
C	Phasors and Heterodyne Interference	215
C.1	A Gaussian Beam	215
C.2	Two Gaussian Beams	217
D	System Control and Data Acquisition	219
D.1	Controllers and Their Assignments	219
D.2	LabVIEW Primer	225
D.3	Analog Signal Production	225
D.4	FPGA-based Digital Demodulation and Feedback	229
D.5	Digitization	239
D.6	Comment on Details Not Included	241
E	Custom Analog Circuits	243
E.1	Direct QPD Readout	244
E.2	Retroreflected Interferometry Readout	248

E.3	Feedback Dynamic Range Control	250
E.4	Thermistor Readout	252
F	The Cosmological Constant	255
F.1	Historical Context	255
F.2	A Friedman Universe	257
F.3	Observations in Our Friedman Universe	259
F.4	Relevant Length Scales	262
G	Attractor and Shield Fabrication	265
G.1	Basic Processes	266
G.2	Density-Staggered Attractor Recipe	282
G.3	Free-standing Electrostatic Shield Recipe	287
	Bibliography	313

List of Tables

2.1	Properties of the center of mass degrees of freedom, derived from fitting the amplitude and phase of a particular trapped microsphere's response to external forces, seen in Figures 2.20 and 2.21. The coefficients in the equation of motion (Equation (B.2)) are in units of rad/s, whereas the coefficients are reported in units of Hz: $f_0 = \omega_0/2\pi$ and $\Gamma = \gamma/2\pi$. The damping coefficients for x and y are determined almost entirely by the derivative gain terms in the feedback, whereas both the resonant frequency and damping for z are set by the feedback. Integral gain terms primarily affect stability below the frequencies used in the measurement in Figures 2.20 and 2.21.	84
4.1	Microsphere masses, m , averaged over all experimental conditions; radii, r , averaged from two distinct high magnification images of the individual spheres (for the $r = 2.35 \mu\text{m}$) or the population average (for the $r = 3.76 \mu\text{m}$); and the derived density, ρ , for the three microspheres caught on the silicon beam, as well the two microspheres of the larger variety. All measurements include statistical and systematic uncertainties, and the relative contributions are shown explicitly for the measured masses. Numbers 4 and 5 were not individually imaged with SEM, and have significantly different statistical uncertainties due to an improvement in the electrostatic drive system.	107

4.2	Systematic effects on the mass measurement. The amplifier discussed produces the voltage driving the electrodes, and thus the electric field. Geometric misalignment, including optical tolerances, can change the value of the electric field at the location of the trap. The application and subsequent measurement of the electric field is also subject to systematic effects, and each measurement channel can experience electrical pickup. Effects are determined empirically where possible, or are obtained from instrument specifications. . . .	113
5.1	Results of spindown pressure measurements, taken with microsphere No. 1. Statistical and systematic uncertainties are propagated independently from uncertainties on r , I , and τ	132
5.2	Calculated and measured calibration factors κ for the three microspheres used here, together with their measured masses, from which the values of r and then κ are determined. The microsphere density was assumed to be $\rho_{\text{ms}} = 1.55 \pm 0.05$ (stat.) ± 0.08 (sys.) from Reference [19]. The relatively large uncertainty in ρ_{ms} limits the precision with which κ can be calculated, whereas the precision of the measurement of κ is limited by uncertainties in the measured dipole moment and P_{max}	141
6.1	Example strengths and length scales for a set of moduli (i.e. scalar fields) from a specific supersymmetric string theory, and with natural assumptions made about symmetry-breaking energy scales. Reproduced from Reference [179]. . . .	153
7.1	Dipole moments and polarizabilities measured for each microsphere.	171
8.1	List of systematic uncertainties.	192
D.1	Components housed within the central PXI-1042 chassis and their functions.	220
D.2	PCIe components housed within the host PC and their functions.	221
D.3	Components housed within the auxiliary cFP-1808 chassis and their functions.	222
D.4	Instruments controlled with a single GPIB Hub, the NI-GPIB-HS, connected to the host PC via USB.	223
D.5	Instruments connected to the host PC via USB.	224
G.1	Cleaning processes.	267

G.2	Processes related to photolithography. All such processes should be undertaken in areas dedicated to photolithography, in order to avoid accidental exposure of the photoresist.	269
G.3	Processes related to deposition.	273
G.4	Processes related to etching. Many of the tools used have automated wafer handling and pre-programmed recipes. <i>Always</i> start from a pre-programmed recipe known to work for other users, and then tweak as necessary based on dummy wafer performance.	277
G.5	Miscellaneous processes.	281

List of Figures

1.1	A schematic depiction of one of the first experiments performed by A. Ashkin. (a) A single laser is brought to a focus within a glass cell containing latex microspheres suspended in water. The forward-scattered light is then viewed downstream in order to monitor microsphere motion. (b) A dual-beam setup where counter-propagating beams result in an equilibrium position in which the microsphere can be fully confined in three dimensions. High-angle scatter is used to monitor the motion of a trapped from a position orthogonal to the beam axes. In both iterations, it was clear to Ashkin that particles with an index of refraction higher than the surrounding medium were drawn to the region of highest intensity along the beam axis (<i>Fair-use educational reproduction from Reference [1]: freely available, Nobel Award winning research provided by the American Physical Society</i>)	2
1.2	A dielectric microsphere in the vicinity of the optical axis of a paraxial laser beam, where the reflection and refraction of two rays has been indicated. Reflection induces radiation pressure forces, while refraction induces deflection forces, both due to the change in photon momentum. The beam has an intensity distribution that is maximal along the optical axis, and falls off radially, as indicated by the inset axes on the left side, as well as the relative thicknesses of the two rays. The radial intensity gradient results in an overall force that tends to push the microsphere toward the optical axis and along the direction of propagation. (<i>Adapted from Reference [1]</i>)	5
1.3	A schematic depiction of the simplified momentum conservation argument for radial confinement. Deflection of the trapping beam results in a change of optical momentum, and thus requires a force since $F_{\text{ms}} = dp/dt$. An equal and opposite force is applied to the dielectric microsphere by the photons. . .	7

- 1.4 A dielectric microsphere in the vicinity of the focus of a highly convergent laser beam, where the reflection and refraction of two rays symmetric about the optical axis has been indicated. As before, reflection induces radiation pressure forces while refraction induces deflection forces due to the change in photon momentum. The sum of all optically induced forces shown will push the microsphere toward the focus, while radial confinement will keep the microsphere along the optical axis. The result is fully three-dimensional, optically induced confinement in a single-beam gradient force trap. (*Adapted from Reference [6]*) 9
- 1.5 (upper) Axial trapping efficiency for a microsphere (radius $r_{\text{ms}} \approx 0.8 \mu\text{m}$, index of refraction $n_{\text{ms}} = 1.33$) within an optical trap generated by a $\lambda = 1064 \text{ nm}$ beam in vacuum, focused with a numerical aperture $\text{NA} = 0.9$ and propagating in the $+z$ direction. The resultant optical force, shown by the secondary axis, corresponds to a total beam power of $\mathcal{P} = 1 \text{ mW}$. (lower) Associated potential energy, assuming the optical force is conservative, plotted with a harmonic approximation generated from the slope of the trapping efficiency curve at the axial equilibrium position. A secondary axis shows the scale of the force in units of thermal energy at temperature $T = 300 \text{ K}$. Two dashed lines help guide the eye to see the effective trap depth in units of $k_B T$. Data in both panels is expected to scale linearly with the optical power \mathcal{P} , so that the trap depth and spring constants can be tuned. 13
- 1.6 (upper) Radial trapping efficiency for a microsphere (radius $r_{\text{ms}} \approx 0.8 \mu\text{m}$, index of refraction $n_{\text{ms}} = 1.33$) within an optical trap generated by a $\lambda = 1064 \text{ nm}$ beam in vacuum, focused with a numerical aperture $\text{NA} = 0.9$ and propagating in the $+z$ direction. The resultant optical force, shown by the secondary axis, corresponds to a total beam power of $\mathcal{P} = 1 \text{ mW}$. (lower) Associated potential energy, assuming the optical force is conservative, plotted with a harmonic approximation generated from the slope of the trapping efficiency curve at the radial equilibrium position, $r = 0$. A secondary axis shows the scale of the force in units of thermal energy at temperature $T = 300 \text{ K}$. Data in both panels is expected to scale linearly with the optical power \mathcal{P} , so that the trap depth and spring constants can be tuned. 14

1.7	A vertically oriented optical trap, with a force applied to a microsphere confined within the trap. The resulting deflection of the trapping beam, $\Delta\theta$ is recollimated by a lens of focal length d to a displacement Δx , which can be measured to infer the applied force. Here it is assumed the optical power has been tuned to levitate the microsphere near to focus.	17
1.8	(upper) Axial trapping efficiency for a microsphere (radius $r_{\text{ms}} = 2.35 \mu\text{m}$, index of refraction $n_{\text{ms}} = 1.33$) within an optical trap generated by a $\lambda = 1064 \text{ nm}$ beam in vacuum, focused with a numerical aperture $\text{NA} = 0.12$ and propagating in the $+z$ direction. (lower) Required optical power to levitate the same microsphere, assuming a mass $m_{\text{ms}} = 84 \text{ pg}$, and with $g = 9.81 \text{ kg m/s}^2$ (See Chapter 4).	20
1.9	(upper) Radial trapping efficiency for a microsphere (radius $r_{\text{ms}} = 2.35 \mu\text{m}$, index of refraction $n_{\text{ms}} = 1.33$) within an optical trap generated by a $\lambda = 1064 \text{ nm}$ beam in vacuum, focused with a numerical aperture $\text{NA} = 0.12$ and propagating in the $+z$ direction. A secondary y -axis has been scaled by $\mathcal{P}_{\text{opt}}/c$, assuming an optical power of $\mathcal{P}_{\text{opt}} = 1.1 \text{ mW}$ calculated at $z = 0 \mu\text{m}$ (i.e. the focus, and based on the result shown in Figure 1.8) to demonstrate the expected optical confinement force vs radial position relative to the focus. (lower) Associated potential energy, assuming the optical force is conservative, plotted with the harmonic approximation generated from the slope of the trapping efficiency curve at $z = 0$	22
2.1	A photograph of the experimental chamber with various ports labeled. Components of the vacuum system will be discussed in Section 2.1.1, while the optical systems will be described in Sections 2.2 and 2.4. The gold-coated mechanical structure visible through the viewport is the centerpiece of the apparatus, supporting both the trapping lenses and shielding electrodes to be discussed in Section 2.3.	33

2.2	A block diagram of the vacuum system. The vacuum gauges directly on the main chamber are rated in Torr, although the measurements are usually converted to mbar, or Pa. The scroll pump for the main chamber is located outside of the lab, in a hallway closet, and connected to the vacuum system with a ~ 5 cm inner-diameter steel piping system embedded in the walls and ceiling. Similarly, the gas bottles and manifold are located outside of the lab, in the gowning area, connected with a ~ 1 cm inner-diameter steel piping system.	35
2.3	A schematic diagram of the fiber-optic system, depicting the source laser and all components up to the point of projection into free-space. The trap acousto-optic modulator (AOM) is usually operated at 150 MHz, while the reference beam AOM is operated at 149.375 MHz. All fiber components are single-mode and polarization-maintaining, and custom spliced for high power operation.	37
2.4	A schematic diagram of the input optics, from the point of projection into free-space, to the beam entering the experimental chamber.	39
2.5	(left) A schematic cross-section of the trapping lenses, shielding electrodes, and the housing structure for both. (upper right) The cubical cavity which constitutes the trapping region. (lower right) A photograph through the primary viewing window, where the microsphere is visible due to high-angle scattering of the trapping laser. The image was captured with a cell-phone camera, which retains some sensitivity at 1064 nm. In the schematic views, the foreground electrode has been suppressed, while in the photograph, it was intentionally removed in order to see the trapped microsphere.	43
2.6	A schematic cross-section of the stack of stages, with a typical attractor device mounted (see Chapters 7 and 8, and Appendix G). The electrode cube, including the bore through which the device is inserted into the trap, is visible on the right.	45

2.7	A schematic depiction of the knife-edge technique used to characterize the spot size and divergence of the optical trap. A gold-coated silicon cantilever is scanned across the trapping beam at various axial positions above and below the focus, while the output power from the trap is monitored. As the knife-edge covers and uncovers the beam, the photodiode measures the integral given by Equation (2.1), from which the beam profile can be recovered. . . .	47
2.8	(left) Typical results from a beam-profiling measurement for scanning the knife-edge in both x (top) and y (bottom), with the bottom surface of the knife-edge positioned nearest to the focus, at $h = 40 \mu\text{m}$. The data suggest the presence of a halo and/or diffraction rings with $\sim\text{few}\%$ of the maximum beam intensity. (right) Second moment of the marginalized intensity distribution from beam-profiling as a function of axial position over the full range of the knife-edge's motion, fit to Equation (2.3). The plot includes a dashed line demonstrating some interpretation of the ideal design performance, i.e. the second moment expected if the paraxial approximation was completely valid, the optical system was perfectly aligned, and the beam was actually an ideal Gaussian beam.	48
2.9	(left) A schematic depiction of the optical system used to separate light retroreflected by a trapped microsphere, combine it with a reference beam, and subsequently convert the optical interference to an electronic signal. (right) A photo of the same system, with components annotated. . . .	51
2.10	(left) A schematic depiction of the output optical system which handles the light transmitted through a trapped microsphere, de-magnifying the beam to reduce spot size, superposing with a reference beam, and finally, converting the optical interference to an electronic signal. (right) A photo of the same system, with components annotated.	54
2.11	Interference contrast vs radial position in the focal plane of the trap, for three distinct sets of measurements, taken for consistency, and shown with differing marker shapes. The solid curve represents the prediction of Equation. (2.5) calculated with measured beam waists and known focal lengths. Data is normalized to a maximum of one and centered. The width of the predicted profile is not fit to data. The non-Gaussian tails in the data are likely the result of the halo in the trapping beam, as seen in Figure 2.8.	56

2.12	A schematic depiction of the output optical system to separate and characterize the cross-polarized light generated by rotations of a trapped microsphere. Linearly polarized light enters the chamber and interacts with a trapped microsphere. Rotations of the microsphere in the horizontal plane couple some of the linearly polarized light into the orthogonal polarization, which can then be separated by a polarizing beamsplitter (PBS)	58
2.13	(left) A schematic depiction of the side-view microscope dedicated to light scattered at high-angle from a trapped microsphere. A single lens collimates the scattered light which is collected in a CMOS camera. (right) An example image captured with this microscope while a microsphere is within the trap. The circular feature visible is the opening of one of the truncated pyramidal electrodes.	60
2.14	A schematic depiction of the top-view microscope dedicated to light transmitted through a trapped microsphere, as well as the imaging of devices within the trapping region. (inset) An example image captured with this microscope while a microsphere is within the trap and a mechanical device is in proximity. Due to imperfect polarization, some amount of the trapping beam is visible and indicative of the position of the microsphere. The image is significantly aberrated by the aspheric trapping lens through which the light must pass.	62
2.15	Photograph of the bead dropper, the piezoelectric actuator, and the supporting structure. The bead dropper has been prepared as described in the text, and a layer of microspheres is visible near the far end. An SEM image of the same lot of microspheres is included.	65
2.16	Schematic cross-section of the electrode cube and trapping lenses, highlighting the bead dropper, bead catcher, and the slots through which they enter the beam path above and below the optical trap. The components are shown near to their full extension, and can be retracted by 12 mm using the translation stage shown.	67
2.17	Amplitude spectral density (ASD) of microsphere motion along the radial axis x , for various different microsphere heights within the trap. The signal is uncalibrated. The resonant frequency from a fit of the ASD to Equation (B.4) is projected onto the horizontal plane below the ASDs. It is clear that as the height is lowered (warmer colors) the radial resonant frequency is increasing.	69

2.18	Block diagram of the feedback loop to stabilize a trapped microsphere’s center of mass motion. The x and y (radial) degrees of freedom only require damping and thus derivative feedback, while the z (axial) degree of freedom is stabilized with proportional, integral, and derivative gain. Both the measured power and retroreflected phase are used to construct a feedback signal for the axial degree of freedom, as there is only one actuator to control the trapping beam power. The H block assures that the two axial feedbacks don’t “see” each other, by removing the effect of the z -phase feedback signal from the power feedback error signal. The digital signal processing happens entirely within the FPGA.	74
2.19	Amplitude of microsphere response to an applied electric field as a function of time during discharge. The magnitude of the electric field remains constant throughout the measurement, while flashes of UV photons induce discrete changes in the response amplitude. The changes in amplitude are found to be integer multiples of a unit step, corresponding to the expected response of the microsphere with a single electron charge. The response amplitude here has been normalized into units of electrons, so that the discrete behavior is exceedingly obvious. Additionally, the response amplitude has been fit to a series of successive step functions, where the step sizes are constrained to be integer multiples of a single free parameter: the single-electron response. In this way, it is possible to measure the optical spring constant (along the direction of the applied field), and determine an empirical calibration of the microsphere response to external forces.	78
2.20	Amplitude a typical complex-valued, frequency-dependent transfer function, \mathbf{H}_{ij} , of an optically trapped microsphere. The amplitude scale is in units of [Arb/N], since the applied force is known in physical units, whereas the x , y , and z variables are in some units arbitrarily (but consistently) scaled by the FPGA. The rows don’t necessarily have the same vertical extent, as each response axis can be scaled independently. The on-diagonal components include a fit to the expected response of a damped harmonic oscillator.	82

2.21	Phase of a typical complex-valued, frequency-dependent transfer function, \mathbf{H}_{ij} , of an optically trapped microsphere. The relative phase of the off-diagonal response has not been ‘unwrapped’ for ease of representation, although an interpolation (discussed in the text) would require this. The on-diagonal components include a fit to the expected response of a damped harmonic oscillator.	83
3.1	Trap region: A drawing of the trap region is shown in the left panel, illustrating the 2-mm-diameter holes in the Au-coated pyramidal shielding electrodes through which the ACSB and the trapping beam are brought in and extracted. The trapping light is represented by the red conical feature. The panel to the right illustrates a detail of the end of the ACSB and trapped microsphere, along with the coordinate frame used in the data analysis. A separate arrow shows the direction of Earth’s gravity.	89
3.2	Linearity of the force sensor in the x direction, with residuals from perfectly linear behavior. Data are collected for a 41-Hz drive signal applied with a shielding electrode. The other two degrees of freedom have comparable linearity.	91
3.3	Force noise micrographs with a charged microsphere, represented with respect to a stationary ACSB at various relative positions. Practically, changes in the relative position are produced by displacing the ACSB, yet for ease of representation, the figure shows the ACSB as stationary. The force vectors represent the microsphere response at a frequency far from a single-frequency AC voltage drive applied to the ACSB during these measurements. For the data in the top (bottom) panel, the horizontal component of the force vector represents F_x , the force in the x direction (F_y , the force in the y direction), while the z component of the vector is always the force in the z direction. In both cases, force vectors measured at five y positions (with a different color indicating each y position shown in the legend) in a regularly spaced grid ranging from -40 to $40\ \mu\text{m}$ are plotted at each (x, z) location.	92
3.4	From top to bottom: histograms of the measured force noise in the x , y , and z directions at different grid locations. Also shown are the fits to normal distributions. The low noise and high linearity of the apparatus enable measurements with a dynamic range of over four orders of magnitude.	93

3.5	Vector field plots of (F_x, F_z) , top, and (F_y, F_z) , bottom, in an xz plane of relative positions. As before, changes in the relative position are produced by displacing the ACSB, yet for ease of representation, the figure shows the ACSB as stationary. The black arrows represent the measured force, while the red arrows represent the best fit from the FEA of the \vec{E} field produced by an ACSB with the same overall bias voltage as in the experiment. A few grid points are missing, due to data corruption introduced by the data acquisition instrumentation.	95
3.6	The RMS of the x and y components of the measured force on the microsphere in the plane of the ACSB, as a function of separation from the ACSB. The data are compared to a model described in the text. The bands represents the standard deviation of the model, using different random implementations of the patch potentials. The fit of $F_{x,\text{RMS}}$ ($F_{y,\text{RMS}}$) implies voltage patches of size $\sim 0.8 \mu\text{m}$ ($\sim 0.7 \mu\text{m}$), assuming $V_{\text{patch}} = 100 \text{ mV}$ [77, 78].	97
4.1	Illustration of the measurement technique. A charged microsphere is trapped by a Gaussian laser beam and held at fixed axial position with active feedback. A slowly varying electric field is applied, depicted with a black arrow. The active feedback reduces the optical power, indicated by the intensity of the trapping beam, such that the sum of the optical and electrostatic levitation forces opposing gravity is constant. The relation between optical power and applied field is then extrapolated to zero optical power, allowing a determination of mass from the implied electrostatic levitation field and the known charge.	104
4.2	(upper) Normalized optical power versus applied electric field for 100 50-s integrations with a single, $2.35 \mu\text{m}$ radius microsphere [76]. The extrapolation is performed separately for each integration. The mean of all extrapolations is shown with a dashed black line. The inset shows the distribution of the 100 extrapolated masses. (lower) The same style of measurement for 50 50-s integrations with a single, $3.76 \mu\text{m}$ radius microsphere [88], demonstrating broader applicability of the measurement technique.	106

4.3 Measured microsphere masses in chronological order are shown in the top panel. Unfilled markers indicate a low-vacuum environment, $P = 1.5$ mbar, while filled markers indicate high-vacuum environment, $P = 10^{-6} - 10^{-5}$ mbar. Black markers correspond to measurements with a negatively charged microsphere, while red markers correspond to measurements with a positively charged microsphere. Different microspheres are separated by vertical dashed lines. The mean mass for each microsphere weighted over all experimental conditions is shown in the bottom panel. The blue data points with cross-shaped markers indicate the three microspheres imaged by SEM following their mass measurement. 108

4.4 SEM images of the three microspheres collected on the silicon beam, at $\times 2500$ magnification. The left inset shows one microsphere at $\times 35000$ magnification, overlaid with the best-fit ellipse, and the right inset shows the $(1.000 \pm 0.005) \mu\text{m}$ diffraction grating [89], also seen at $\times 35000$ magnification. The diffraction grating serves as a calibration length scale for the high-magnification images of individual microspheres. 109

4.5 (upper) Distributions of microsphere radii measured by SEM for both conductively coated and uncoated microspheres of the smaller variety are shown in the top panel. Gaussian fits yield central values of $r_{\text{Au/Pd}} = (2.35 \pm 0.04) \mu\text{m}$ and $r_{\text{NC}} = (2.37 \pm 0.04) \mu\text{m}$ for the conductively coated and uncoated microspheres, respectively. Each distribution is generated from approximately 1000 distinct microspheres. Individual radius measurements from analysis of the microspheres seen in Figure 4.4 are shown in the panel immediately below. (lower) Distributions of microsphere radii uncoated microspheres of the larger variety, generated from only 127 distinct microspheres. A Gaussian fit yields a central value of $r_{\text{Au/Pd}} = (3.76 \pm 0.09) \mu\text{m}$. Individual microspheres of this size were not independently imaged, and thus only the population measurement is shown. 111

5.1	Top: typical amplitude spectrum of P_{\perp} for a microsphere prepared in a state of angular momentum pointing along the \hat{z} direction. The microsphere is spinning with $\omega_{\text{ms}} = 100\pi$ krad/s, driven by an electric field with $E = 27$ kV/m. The prominent line with sidebands are signatures of the microsphere rotation, with the position of the central line at twice the drive frequency. The sidebands are phase modulation of the rotation frequency, as shown in the bottom panel.	121
5.2	Harmonic oscillation frequency, ω_{ϕ} , versus driving electric field amplitude, E , for a microsphere spinning at $\omega_{\text{ms}} = 100\pi$ krad/s at a pressure of 2×10^{-6} mbar. The data is fit to Equation (5.6), obtaining $(d/I) = (108 \pm 2)$ A s/kg m) and $d = 127 \pm 14$ e μm	123
5.3	Top three panels: Spectra of the cross-polarized light intensity, \mathcal{P}_{\perp} , for a microsphere precessing about the \hat{x} axis while spinning at $\omega_{\text{ms}} = 100\pi$ krad/s). The modulation of the cross-polarized light occurs predominantly at twice the precession frequency, denoted by red diamonds. Bottom panel: Ω for different E . The slope of the fit (red line) provides $(d/I) = \omega_0\Omega/E = (106 \pm 2)$ A s/kg m, which is consistent with the measurement of d/I from the frequency of small oscillations.	124
5.4	Time evolution of ω_{ms} after the driving electric field is switched off. For $\omega_{\text{ms}} \gtrsim 150$ krad/s the angular velocity exhibits an exponential decay. For $\omega_{\text{ms}} \lesssim 150$ krad/s the dynamics are modified by torque that could be explained by a ~ 100 ppm degree of ellipticity in the 1 mW trapping beam and the $\eta \sim 10^{-2}$ phase retardation of the microsphere.	128
5.5	Damping time τ calculated from Equation (5.16) versus median integration time $\langle t_i \rangle$, for the four measurements with microsphere No. 1, indexed chronologically. The bands show the 1σ uncertainty propagated from the individual uncertainties of the values used to calculate τ	130

5.6	Exponential decay of microsphere No. 1 angular velocity due to torsional drag from residual gas after a 110 kHz driving field has been turned off. Data are shown for the same four measurements detailed in Figure 5.5. Dotted lines indicated realizations of Equation (5.15) with $\tau = \tau_j$ for measurements $j = 0, 1, 2, 3$, where τ_j is the average of the values plotted in Figure 5.5. The structure in the residuals, with amplitude $\sim 2\%$, is likely the result of a slowly fluctuating optical torque.	131
5.7	Top three panels: Equilibrium phase, ϕ_{eq} , versus chamber pressure for several magnitudes, E , of the driving field with $\omega_{ms} = 100\pi$ krad/s. For each value of E , ϕ_{eq} increases until the microsphere loses lock with the field, and the phase becomes random. For each E , a fit to Equation (5.5) (with the argument in Equation (5.17)) is shown in red. P_{max} is identified by a red diamond. Bottom panel: The linear relationship between P_{max} and E , with additional E included. The slope of the fit of E vs p_{max} is 639 ± 64 (kV/m)/mbar. Assuming the dipole moment measured from the frequency of small oscillations, this gives $k = \beta_{rot}/P = (4.1 \pm 0.6) \times 10^{-25}$ m ³ s, which is consistent with $k = \beta_{rot}/P = 3.4 \times 10^{-25}$ m ³ s predicted in References [123, 124].	135
5.8	Equilibrium phase lag between the orientation of the microsphere dipole moment and a rotating electric field, versus the residual gas pressure determined by two capacitance manometers. Three distinct measurements are shown with He, Ar, and SF ₆ residual gas. As the pressure is increased, the phase lags according to Equation (5.18), and P_{max} is measured from a fit to this expression. As ϕ_{eq} approaches $-\pi/2$, the microsphere rotation becomes unlocked from the driving field (ϕ becomes random) at a pressure slightly below P_{max} , due to librational motion of the microsphere and torque fluctuations from the increasing pressure of the surrounding gas.	136
5.9	The quantity P_{max}/d is plotted against the effective mass of the gas $m_{0,eff}$ for different species spanning the 4 to 150 amu range, for three different microspheres. With E and ω_0 known, the constant κ is extracted by fitting Equation (5.19) to the data, with a single free parameter, as shown for each microsphere. A χ^2 -minimization finds $\chi_{1,min}^2/N_{DOF} = 1.9/1$, $\chi_{2,min}^2/N_{DOF} = 3.5/1$, $\chi_{3,min}^2/N_{DOF} = 8.5/5$. Residuals shown below are plotted with the same units as the data.	139

6.1	Parameter space of Yukawa-type, $\alpha - \lambda$, modifications to gravity, including current experimental constraints. Gravitational strength interactions, $\alpha \approx 1$, are indicated with a dashed line. Adapted from References [141, 160], with significant additions in the submillimeter regime (the region of interest to this work) from more recent publications [193–199].	159
6.2	Schematic depiction of the torsion balance described by Reference [198]. The test and source masses are patterned platinum foils adhered to a glass substrate, with the epoxy serving as the low-density component of the modulation. The components are all gold-coated to reduce electrostatic interactions. A stationary electrostatic shield (not depicted) sits between the source and test masses to shield anomalous backgrounds. As the source mass rotates at angular frequency ω , an oscillating gravitational torque, plus any new interactions, causes the test mass assembly to torsionally oscillate at $n\omega$ and $m\omega$ where n and m are the number of full periods in the two distinct density modulations present. The density modulations depicted have 8-fold and 40-fold symmetry, while the actual device has 18-fold and 120-fold symmetry.	161
6.3	Schematic depiction of the microcantilever apparatus described by Reference [194]. The test mass consists of a solid gold block mounted to the end of a silicon cantilever, with some loaded resonant frequency ω_0 . A nanofabricated chip with a regular 25 μm density modulation serves as the source mass, and is driven to oscillate n periods of the modulation at a frequency ω such that $\omega_0 = n\omega$. A stationary electrostatic shield (not depicted) sits between the source and test masses to shield anomalous backgrounds. The high-density gold portion is also continuous, allow a current to be driven through it for electromagnetic calibrations.	163

7.1	(left) Schematic of the optical trap and shielding electrodes. The electrode in the foreground is removed to show the inside of the trap. (right) Zoom in on the region near the trap. A 4.7 μm diameter microsphere is suspended at the focus of an upward propagating laser beam. The 10 μm thick Au-coated Si cantilever is positioned at 20 μm to 200 μm separations from the microsphere and oscillated in the z direction using a nanopositioning stage. (This figure is nearly identical to Figure 3.1, aside from a redefinition of the coordinate system. Both figures are included in their respective chapters to maintain clarity and consistency.)	169
7.2	Measured response of microsphere #1 versus distance from the cantilever face as the cantilever is swept in z with a constant bias of 1,2,3,4, and 5 V. The data points are shown as dots and the best fit model as solid lines. (inset) Amplitude of the fit component $\propto \partial_z E_z$ (top) and the fit component $\propto E_z \partial_z E_z$ (bottom). Fits to the expected linear and quadratic dependence on the voltage are also shown (solid lines).	171
7.3	An example FEA calculation of the chameleon field. A cross-section in the yz -plane is shown, so that the cantilever appears edge-on and extends into and out-of the page. The residual vacuum pressure of $\sim 10^{-6}$ mbar does not affect the maximum value of the field here, which is instead limited by the finite size of the vacuum chamber, as the boundary condition enforces a low value at the vacuum-matter interface	173
7.4	Measured response for microspheres #1 (top), #2 (middle), and #3 (bottom) versus distance from the cantilever face as the cantilever is swept in z with a nominal bias of 0 V. The best fit electrostatic background-only model (dashed) and the amplitude of a chameleon force at the 95% CL upper limit for $\Lambda = 10$ meV (solid) are also shown.	175

7.5	Limits on Λ versus $1/\beta = M/M_{Pl}$ for the chameleon model discussed in the text. The 95% CL exclusion limits from this search are denoted by the dark (gray) region. Recent constraints from atom interferometry are shown by the light (blue) region [172, 210]. The horizontal line indicates $\Lambda = 2.4$ meV. Limits from neutron interferometry [211–213] and from the Eöt-Wash torsion balance experiment [175, 204, 217] are denoted by the hatched regions. These limits are shown only in the restricted regions of parameter space considered in Refs [211] and [175].	176
8.1	(a) Central portion of the experimental setup: a microsphere is trapped in an optical tweezer. A stationary shield centered about the trapped microsphere, with the closest surface within a few microns of the microsphere, and the attractor system is behind it. (b) Scanning electron microscope (SEM) image of the attractor system. The dark (bright) regions correspond to silicon (gold). (c) SEM image of the shield, viewed at a 40° angle to highlight the three-dimensional structure, with the vertical wall to the left.	181
8.2	Composite image of two nanofabricated devices in the vicinity of the trap, together with a trapped microsphere, captured by the side-view microscope. Device 1 is the attractor system, mounted on the main nanopositioning stage, and Device 2 is the electrostatic shield, mounted on the auxiliary nanopositioning stage. The L-shaped cross-section of the shield, which usually houses the attractor system, is not visible here. In this image, the attractor system is translated vertically from its nominal position and is out of focus, the latter reducing the apparent vertical extent. The main frame of the image shows the scattering of the trapping laser by the microsphere (A), which leads to saturation of several pixels of the camera. The inset is taken with a notch filter that has an optical depth of 6 for 1064 nm light (Thorlabs NF1064-44), in order to demonstrate the shadow of the microsphere (B) blocking the 870 nm illumination of the shield. The bright and blurry region (C) behind the device is caused by a reflection of the illumination from the handle of the attractor system.	183

- 8.3 Amplitude spectral density (attractor systemD) of the z component of the force on a $7.56\ \mu\text{m}$ diameter microsphere. The black (red) curve shows actual data with the attractor system stationary (scanning along y at 3 Hz with $202\ \mu\text{m}$ peak-to-peak amplitude). The blue bars show a comparison to an expected microsphere response produced by the potential described by Equation 8.1 with $\alpha = 10^{10}$ and $\lambda = 10\ \mu\text{m}$. The data displayed here is the average of 100 distinct 10-s integrations. 185
- 8.4 An example result of the mesh calculation for the Yukawa-modified gravitational force with $\alpha = 1$ and $\lambda = 10\ \mu\text{m}$, in all three orthogonal axes, on a $d = 7.5\ \mu\text{m}$ and sourced from the attractor system described. The expected force is shown for a few distinct horizontal separations, and a single fixed vertical separation, included to demonstrate exciting a response in the axial direction. 187
- 8.5 The single harmonic maximum likelihood estimator (MLE) $\hat{\alpha}_i$ for $\lambda = 10\ \mu\text{m}$ as a function of time for the six harmonics used in the analysis. Each harmonic f_i is evaluated separately taking into account its own phase response and noise level. Here, each estimation of $\hat{\alpha}_i$ comes from 5000 seconds of data. The error bars represent 95% confidence intervals about the MLEs. The panel to the right shows the MLE for each harmonic, integrating over the entire data set (note the expanded vertical scale). 188
- 8.6 Limit curve in the $\alpha - \lambda$ parameter space. The region above and to the right of the red and blue lines indicates the parameter space excluded by this experiment for positive and negative α , respectively, with a 95% confidence level. The gray region shows the parameter space covered by previous searches [194–196, 198]. The background-free sensitivity for this run, using the current attractor system-microsphere separation, noise conditions, and integration time is shown by dash-dotted orange line. In addition, the projected sensitivity for the next run, given the improvements outlined in the text, is shown by the dash-double-dotted purple line. This assumes a noise floor of $1 \times 10^{-18}\ \text{N}/\sqrt{\text{Hz}}$, attractor system-microsphere separation of $7.5\ (-5)\ \mu\text{m}$ in the $x\ (z)$ direction, and an integration time of 30 days. The reach could be extended further by using larger microspheres [34] or smaller separations [37]. 191

8.7	A sphere is shown confined by a optical tweezer. An applied force pushes the sphere to one side, resulting in a deflection of the laser beam used to from the optical tweezer. The deflection can then be used to monitor the motion of the sphere, and eventually measure the applied force.	205
D.1	A cartoon depiction of digital demodulation. The arbitrary phase ϕ from Equation. (D.1) was set to 0 for simplicity.	233
E.1	(upper) An LTSpice implementation of a single QPD quadrant and the first stage of amplification, including a realistic model of both the photodiode itself, as well as the inductor used to shunt the DC component of the photodiode signal. (lower) Bode plot of the amplitude and phase of the amplifier response as a function of frequency. A dashed vertical line indicates the chosen interference frequency of 625 kHz.	245
E.2	The second stage of amplification in the direct QPD readout board. An AC-coupled, non-inverting amplifier, with low-pass behavior for anti-aliasing. “No Connection” (NC) terminals of the op-amp have been terminated with a resistor and capacitor in parallel, following guidelines in the manufacturer datasheet.	246
E.3	A photograph of the QPD circuit, with the output beam reduction telescope visible on the left, as well as the reference beam collimator visible in the lower left. A superposition of the transmitted light and reference beam light is generated with the beamsplitter cube, and subsequently projected onto the QPD, which has a lens tube serving as shroud in this image.	247
E.4	(upper) An LTSpice implementation of the resonant non-inverting amplifier used as a second stage of the retroreflected interferometry readout (lower) Bode plot of the amplitude and phase of the amplifier response as a function of frequency. A dashed vertical line indicates the chosen interference frequency of 625 kHz.	249

E.5 A schematic of the level matching circuit implemented to take full advantage of the dynamic range of our analog drivers. The output voltage from a DAC within the FPGA is first buffered by an instrumentation amplifier, whose output drives a resistive divider in order to scale down the amplitude of the voltage. The output of the resistive divider is buffered by a unity-gain op-amp and sent to one of two inputs of a summing amplifier. A tunable DC bias is generated with a pair of precision 5 V voltage references, whose outputs are connected across a variable resistor such that the wiper terminal of the variable resistor can be anywhere in the range ± 5 V. 251

E.6 A schematic of the Wheatstone bridge readout for a single thermistor. The differential voltage is measured and amplified by an instrumentation amplifier, and is expected to be 0 V at room temperature and with the variable resistor set to 10 k Ω . The bridge can be biased to produce a positive-definite voltage by adjusting this variable resistor. 253

G.1 Selected steps from the attractor fabrication process, with numbers corresponding to the enumeration in the text. (1) trench lithography, (2) trench etch, (3) grow thermal oxide, (4) evaporate gold, (5) CMP topside, (6) electroform in trench with CMP after, (7) outline lithography, (8) outline etch, (9) backside PECVD oxide, (10) release lithography, (11) deep release etch, (12) HF vapor etch. Devices are then snapped free of the carrier wafer and sputter-coated. *Significantly* not-to-scale. 285

G.2 Annotated SEM micrographs of a completed attractor device, prior to deposition of the final metallic coating. Two details show the front face of the attractor with the characteristic scalloping from the deep silicon etching with the Bosch process [92], as well as one of the electroformed and polished gold portions of the density modulation. A few length scales are indicated. 286

- G.3 Selected steps from the free-standing electrostatic shield fabrication process, with numbers to match the text. (1) Grow thermal oxide, (2) “pocket” lithography, (3) etch pocket in oxide mask, (4) strip resist then outline lithography, (5) etch oxide hard mask and partial device layer, (6) strip resist, (7) etch pocket and outline simultaneously to buried oxide stop, (8) backside PECVD oxide, (9) release lithography, (10) deep release etch, (11) HF vapor etch. Devices are snapped free of the carrier wafer and sputter-coated. *Significantly not-to-scale.* 290
- G.4 (upper) A schematic depiction of the attractor and shield positioning near a trapped microsphere, emphasizing the L-shaped cross-section. (lower) Annotated SEM micrographs of a completed shield device, prior to deposition of the final metallic coating, with both images using a 40° inclination in viewing angle to make the three-dimensional structure clearly visible. A detail shows the end of the shield and the symmetric pockets for this particular device. A few length scales are indicated. 291

Chapter 1

Optical Tweezers

The interaction between photons and normal matter, composed of protons, neutrons, and electrons, has proven to be one of the most fruitful areas of science in the last half century, allowing the investigation of a wide variety of phenomena. Most often, these phenomena are made made accessible through the use of lasers, taking advantage of the inherent wavelength control, as well as the coherence of photons from a laser source.

In particular, lasers can be focused to extremely small spot sizes in order to generate optical fields with high intensity and sharp intensity gradients, as compared to sources of incoherent and/or broadband photons, amplifying any interactions between the laser photons and a target material. Of specific interest to this work is the interaction between a focused laser and particles with an index of refraction larger than that of the surrounding medium.

By virtue of this interaction, particles with high index of refraction are attracted to the region of space with the highest optical intensity, and can thus be localized and confined. This is generally referred to as an optical tweezer.

1.1 Ashkin and the First Optical Tweezer

Working at Bell Laboratories in the 1970s, Arthur Ashkin was developing methods with the ultimate aim to trap atoms using focused optical fields. Along the way, he had to build an understanding of the forces that a particle would experience in the vicinity of a tightly focused laser beam, as well as any effects from a surrounding medium. To that end, he was studying the behavior of micron-sized particles suspended in fluid, as well as aerosolized fluid droplets suspended in air.

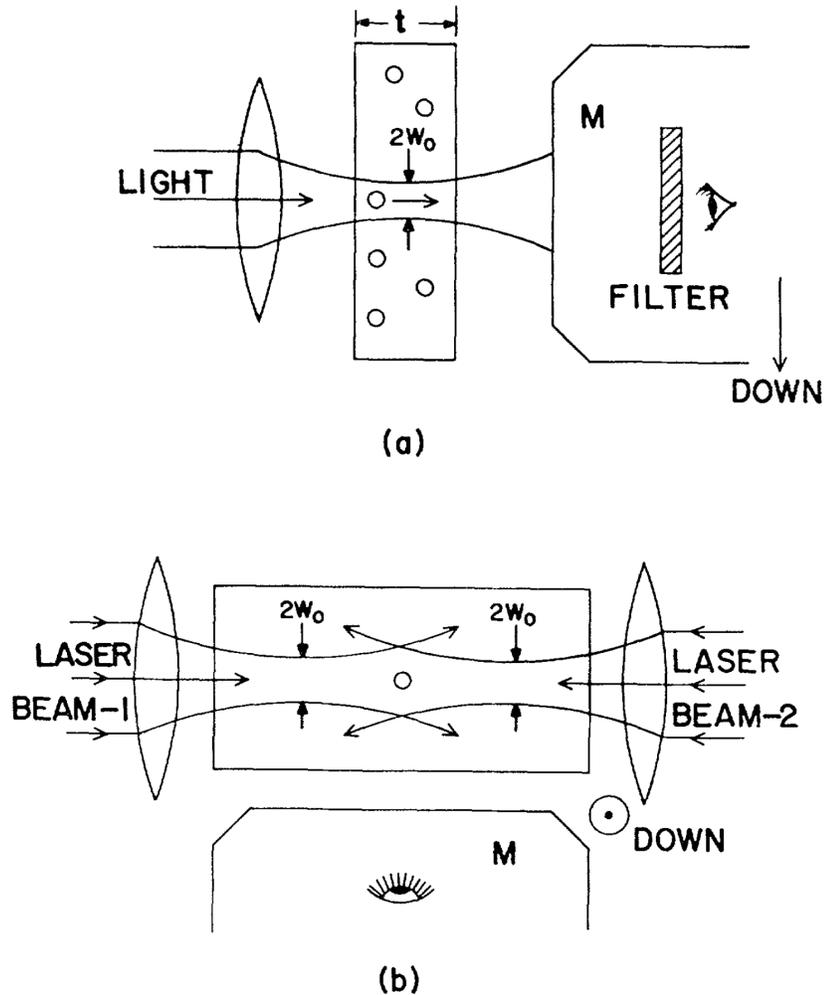


Figure 1.1: A schematic depiction of one of the first experiments performed by A. Ashkin. (a) A single laser is brought to a focus within a glass cell containing latex microspheres suspended in water. The forward-scattered light is then viewed downstream in order to monitor microsphere motion. (b) A dual-beam setup where counter-propagating beams result in an equilibrium position in which the microsphere can be fully confined in three dimensions. High-angle scatter is used to monitor the motion of a trapped microsphere from a position orthogonal to the beam axes. In both iterations, it was clear to Ashkin that particles with an index of refraction higher than the surrounding medium were drawn to the region of highest intensity along the beam axis (*Fair-use educational reproduction from Reference [1]: freely available, Nobel Award winning research provided by the American Physical Society*)

In the initial experiment [1], Ashkin had latex microspheres of various sizes ranging from $0.6\ \mu\text{m}$ to $2.7\ \mu\text{m}$ freely suspended in water within a glass cell. A $\lambda = 514\ \text{nm}$ laser was brought to a focus within this glass cell and positioned to impinge upon and interact with individual microspheres suspended within the cell. The laser light scattered by the microspheres was then observed with an objective lens. A depiction of the apparatus from Ashkin's manuscript is shown in Figure 1.1a.

With sufficient beam power, Ashkin immediately observed that microspheres near the beam were drawn radially inward toward the optical axis, where the intensity was highest, as well as being pushed along the direction of propagation by radiation pressure. When two identical beams were focused and co-aligned with one another, as in Figure 1.1b, he observed that the microspheres would eventually come to rest at an equilibrium position along the mutual optical axis, where the radiation pressure from both lasers was balanced.

And thus the first optical tweezer was born. To better understand the mechanism of this optically-induced confinement, we will appeal to two geometric optics arguments: one purely qualitative yet closer to reality, and another that allows for a simple quantitative parameterization.

1.2 Scattering of Light

Generally speaking, the scattering of photons off matter composed of protons, neutrons, and electrons is governed by Maxwell's equations [2, 3], which describe a classical field theory with gauge freedom. Limiting ourselves to the macroscopic formulation which averages over the behavior of individual atoms, the canonical four equations refined by Heaviside can be combined into two wave equations [4]. These are given by,

$$\left(v^2\nabla^2 - \frac{\partial^2}{\partial t^2}\right)\vec{E} = 0 \quad \left(v^2\nabla^2 - \frac{\partial^2}{\partial t^2}\right)\vec{B} = 0, \quad (1.1)$$

where $v = 1/\sqrt{\mu\epsilon} = c/n$ is the phase velocity, μ and ϵ are the permeability and permittivity of the medium, c is the speed of light, and n is the index of refraction. If the spectral theorem is assumed, the analysis can be restricted to monochromatic fields of the form $\vec{E}(\vec{x}, t) = \vec{E}(\vec{x})e^{-i\omega t}$ with ω the angular frequency of oscillation, which reduces the wave

equation to the Helmholtz equation for the electric (or magnetic) field,

$$(\nabla^2 + k^2) \vec{E} = 0 \quad \text{and} \quad \vec{B} = \frac{1}{ik} \nabla \times \vec{E}, \quad (1.2)$$

where $k = \omega/v = 2\pi n/\lambda$ is the wave vector, with λ the vacuum wavelength and v as before. Solutions to this equation can describe arbitrary electromagnetic fields in vacuum and homogeneous media.¹ With these equations in hand, the situation of greatest interest can be considered: the scattering of monochromatic laser light by spherical dielectrics. The nature of this scattering depends significantly on the ratio between the laser's vacuum wavelength λ and the particle diameter d , and can be understood in three distinct regimes.

For Rayleigh scattering, when $\lambda \gg d$, atomically or molecularly bound electrons are treated as harmonic oscillators subject to the electric field of incident photons, which cause the electrons to oscillate about the atomic nucleus. Their motion produces an oscillating dipole moment which then radiates the scattered electric field [4]. For geometric optics, when $\lambda \ll d$, solutions to Equation (1.2) are developed in the zero-wavelength limit, where material boundaries defined by discontinuities in $\epsilon = \epsilon(\vec{x})$ or $\mu = \mu(\vec{x})$ yield Fresnel reflection coefficients (for specular reflections) and the Snell-Descartes Law of Refraction [5]. These two relations can be used to understand a wide variety of optical systems with various reflecting mirrors and dielectric elements (lenses, prisms, etc.).

Before diving into the details of Mie Scattering, which describes the all-important regime where $\lambda \sim d$, we will make use of more approachable geometric optics arguments in order to understand optical trapping.

1.2.1 An Effective Geometric Optics Picture

Following Ashkin's original paper, consider a dielectric microsphere in the vicinity of a laser beam, but offset from the optical axis. Assume incident rays are parallel, but have an intensity distribution that is maximal along the optical axis, and falls off radially, such as in the case of a TEM₀₀ mode with a Gaussian intensity distribution. Reflection and refraction of the incident rays results in radiation pressure and deflection forces on the dielectric microsphere, since photons carry momentum and $\vec{F} = d\vec{p}/dt$. This scenario is shown schematically in Figure 1.2.

¹Although we will restrict ourselves to homogeneous media for simplicity, the above equations can be extended to inhomogeneous media simply by redefining the wave vector: $k \rightarrow k(\vec{x}) = \omega n(\vec{x})/c = 2\pi n(\vec{x})/\lambda$, where the inhomogeneity is parameterized by a spatially varying index of refraction.

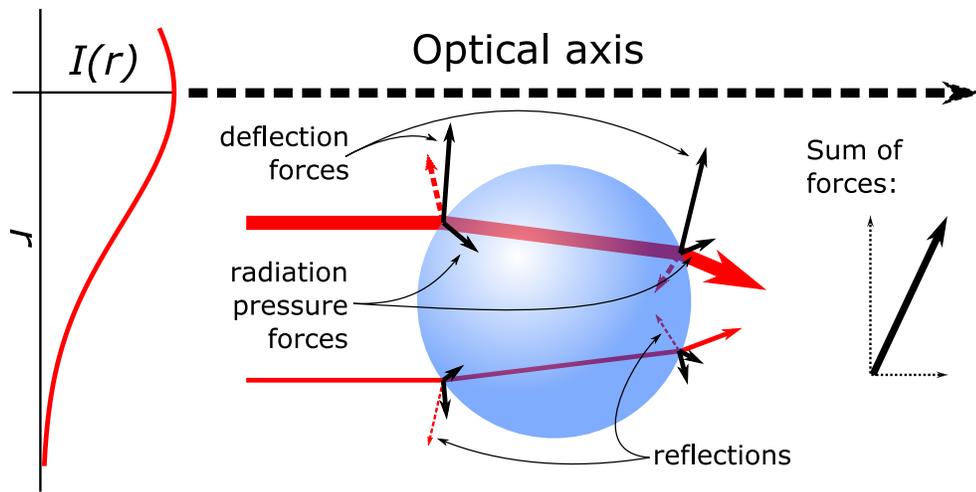


Figure 1.2: A dielectric microsphere in the vicinity of the optical axis of a paraxial laser beam, where the reflection and refraction of two rays has been indicated. Reflection induces radiation pressure forces, while refraction induces deflection forces, both due to the change in photon momentum. The beam has an intensity distribution that is maximal along the optical axis, and falls off radially, as indicated by the inset axes on the left side, as well as the relative thicknesses of the two rays. The radial intensity gradient results in an overall force that tends to push the microsphere toward the optical axis and along the direction of propagation. (Adapted from Reference [1])

Two rays are drawn, symmetrically placed above and below the center of the dielectric microsphere so that the angles of their reflection and refraction are mirrored. Because the upper ray has a higher intensity, indicated by its thickness, the deflection and radiation pressure forces induced by its refraction and reflection are larger than same forces produced by the less intense ray. The lengths of the resultant force vectors are indicative of their relative strengths, although the diagram is purely qualitative. The total force acting on the microsphere from these two rays simultaneously pushes the microsphere along the direction of propagation, and toward the region of highest intensity, as indicated, and can thus be thought of as a restoring force.

This argument can easily be extended to the entire beam and microsphere, noting that in the integral over the microsphere there will always be pairs of rays with mirrored deflection angles, but where one ray is more intense than the other, resulting in an overall force that draws a dielectric microsphere toward the optical axis of the trapping beam.

1.2.2 A Naïve (but Quantitative) Parameterization

With an even simpler geometric optics picture, it is possible to provide an approximate parameterization in order to quantify the scale of the optical confinement forces. Consider a collimated laser beam with total power \mathcal{P}_{opt} as a single ray (most of the power will be near the center for commonly used beams), and treat a dielectric microsphere as a basic thin lens, assuming a single point of deflection within the bulk of the microsphere. Now, offset the microsphere from the beam axis by some distance Δx . Under this simple construction, the beam will be deflected by an angle $\Delta\theta$, as shown schematically in Figure 1.3.

Because photons carry momentum, totaling $\mathcal{P}_{\text{opt}}/c$ for a beam with optical power \mathcal{P}_{opt} , a deflection implies a change in optical momentum flux and thus an associated force applied to the photons. By Newton’s third law, there is an equal and opposite force that acts on the microsphere, \vec{F}_{ms} , and provides a restoring force that confines the microsphere along the optical axis,

$$\vec{F}_{\text{ms}} = -\frac{d\vec{p}_{\text{photons}}}{dt} \approx \xi \frac{\mathcal{P}_{\text{opt}}}{c} \Delta\theta (-\hat{r}), \quad (1.3)$$

where ξ is a “trapping efficiency” factor, given that not all of the photons will interact with the microsphere since $d \not\gg \lambda$ (at least for the experimental parameters of interest), and \hat{r} is a unit vector pointing radially outward. Thus, the force in the $-\hat{r}$ direction is a

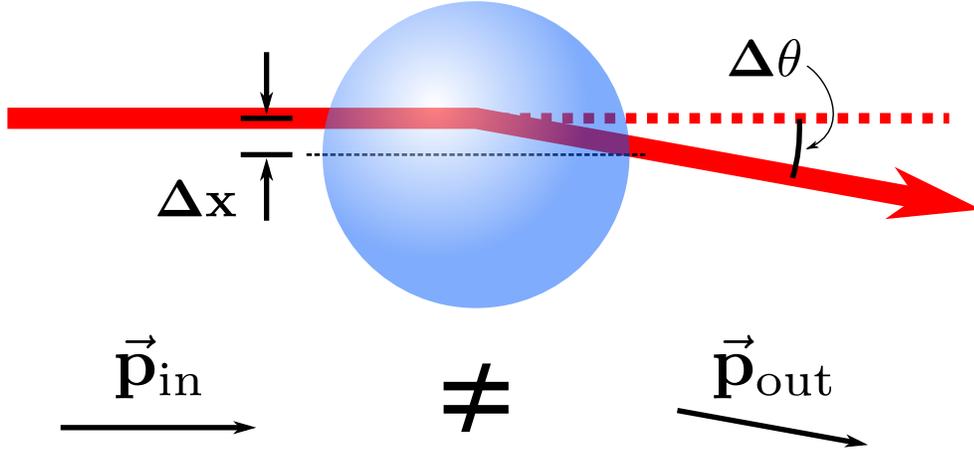


Figure 1.3: A schematic depiction of the simplified momentum conservation argument for radial confinement. Deflection of the trapping beam results in a change of optical momentum, and thus requires a force since $F_{\text{ms}} = dp/dt$. An equal and opposite force is applied to the dielectric microsphere by the photons.

restoring/confining force.

Assuming some representative numbers, $\mathcal{P}_{\text{opt}} \approx 1 \text{ mW}$, $\xi \approx 0.3$, and $\Delta\theta \approx 1 \text{ mrad}$, the optical force implied by this parameterization is,

$$\vec{F}_{\text{ms}} \approx (1 \times 10^{-15} \text{ N}) \left(\frac{\xi}{0.3} \right) \left(\frac{\mathcal{P}_{\text{opt}}}{1 \text{ mW}} \right) \left(\frac{\Delta\theta}{1 \text{ mrad}} \right) (-\hat{r}), \quad (1.4)$$

which gives a reasonably accurate indication of the scale of forces involved when optically trapping dielectric microspheres for easily measurable deflection angles and typical beam powers. It is expected that $\xi \sim d^\alpha$ with $\alpha > 0$ although the existence of resonances at various d complicate this picture slightly.

1.2.3 Comment on Axial Confinement

An important feature that has been glossed over significantly is the nature of the axial confinement (as opposed to radial confinement) induced by the interaction between a focused laser beam and a dielectric particle. Much like the radial confinement discussed previously, axial confinement can be understood in terms of a momentum conservation argument, wherein both reflection of refraction of incident photons induce forces as a result of the change in photon momentum. Let us consider a paraxial laser beam brought to a

highly convergent focus by a lens, and examine the interactions of the beam with a dielectric microsphere placed below the focus, making use of geometric optics to gain an effective understanding. This argument is presented schematically in Figure 1.4, which has been adapted from the first implementation of fully three-dimensional trapping with a single beam [6].

The reflection and refraction of two rays symmetrically placed about the optical axis, and thus assumed to have the same intensity (or nearly so for most laser beam profiles), produce forces that, overall, tend to push the microsphere toward the focus. The same argument holds for a microsphere placed above the focus, where the resulting force would be directed downward, and thus an equilibrium position at, or near to, the focus is expected. In their 1986 work [6], Ashkin *et al.* implemented an optical trap in which the main trapping beam was propagating downward (in the same direction as gravity), in order to unquestionably demonstrate that the axial force induced by a single, highly-focused trapping beam interacting with a dielectric microsphere was sufficient to generate a “full” optical tweezers exhibiting optically-induced, three-dimensional confinement, even with gravity and radiation pressure working in tandem against the confining gradient force, rather than in opposition to one another.

For single-beam optical tweezers with a low numerical aperture (NA), an axial equilibrium position does not exist in the absence of external forces. Essentially, with low NA, the Rayleigh range of the beam’s focus can be much larger than the trapped particle itself, and thus there is a vanishing axial intensity gradient in the vicinity of the focus. Within the Rayleigh range, rays can be thought of as propagating parallel to the optical axis, in which case the deflection and radiation pressure forces induced by the reflection and refraction of the parallel rays tends to push the microsphere along the beam axis, rather than confine it to some axial equilibrium position. This was the case for Ashkin’s original experimental apparatus [1], and is clearly visible in Figure 1.2, which was adapted from that work.

For optical systems where this is the case, other methods must be utilized for fully three-dimensional trapping. Two of the more common approaches were both developed originally by Ashkin, together with Dziedzic [1, 7]. As seen in Figure 1.1, it is possible to balance the on-axis radiation pressure force simply with a secondary, counter-propagating, trapping beam. Properly aligned, both beams contribute to radial confinement, while the combination of their opposing radiation pressure forces produces axial confinement. This configuration has the advantage that the effective optical spring-constant along each direction can be tuned

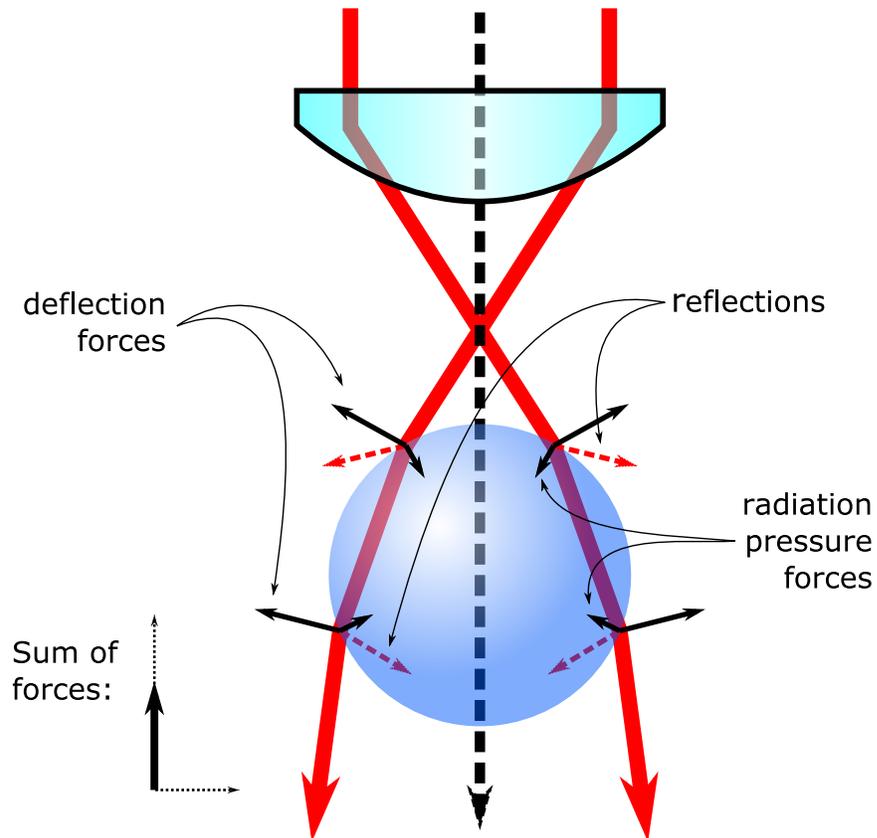


Figure 1.4: A dielectric microsphere in the vicinity of the focus of a highly convergent laser beam, where the reflection and refraction of two rays symmetric about the optical axis has been indicated. As before, reflection induces radiation pressure forces while refraction induces deflection forces due to the change in photon momentum. The sum of all optically induced forces shown will push the microsphere toward the focus, while radial confinement will keep the microsphere along the optical axis. The result is fully three-dimensional, optically induced confinement in a single-beam gradient force trap. (*Adapted from Reference [6]*)

by increasing the power of the trapping beams. Asymmetric adjustments to each beam's power, and thus radiation pressure, allows for tuning of the equilibrium position along the mutual optical axis of both beams. This solution also has clear downsides, as it not only requires the use of two focused laser beams, but also a precise co-alignment of their optical axes, which is by no means an easy task. It also complicates the required input and output optical systems that serve to shape and direct the incident and scattered beams.

The on-axis radiation pressure force of a low NA optical tweezer can also be countered with any other force, such as gravity due to interaction with the Earth. A vertically-oriented, upward-propagating optical trap does exactly this, using radiation pressure to levitate a dielectric microsphere against gravity, as Ashkin demonstrated numerous times [1, 8, 9]. In such a configuration, the axial equilibrium position is determined solely by the optical power. This can be understood by considering the axial position at which $|\vec{F}_{\text{rad}}| = |\vec{F}_{\text{grav}}|$, where F_{grav} is assumed constant and F_{rad} depends on the optical intensity at that axial position (and thus total power), which is maximized near the focus. This will be the approach taken in this work, and will be discussed again in more detail when appropriate.

1.3 Mie Scattering - The T -Matrix Method

Although the arguments presented in the previous sections are useful to get an intuitive understanding, accurate quantitative predictions are difficult when $\lambda \sim d$, which is exactly the relevant regime with a 1064 nm trapping laser and micron-scale dielectric microspheres. Analytical solutions exist for simple cases, such as homogeneous spherical particles, and are referred to as generalized Lorenz-Mie scattering theory [10–12], often requiring infinite summations of slowly converging series. Numerical techniques reign supreme here, allowing for the development of quantitatively accurate approximate solutions.

Following the construction of the Helmholtz equation (Equation (1.2)), again appealing to the spectral theorem to write both incident and scattered optical fields as superpositions of a complete set of basis functions that are themselves solutions to the Helmholtz equation [13],

$$E_{\text{inc}} = \sum_n^{\infty} a_n \psi_n^{(\text{inc})} \quad \text{and} \quad E_{\text{scat}} = \sum_m^{\infty} p_m \psi_m^{(\text{scat})}, \quad (1.5)$$

where a_n and p_m are coefficients of the basis functions. The summations are not usually taken to infinity, rather, they are truncated at some value set by the convergence behavior

of the coefficients for a particular scenario. Assuming the dielectric microsphere has a linear response to the incident optical fields, the scattered field is related to the incident field by a simple matrix equation,

$$p_m = \sum_n^{\infty} T_{mn} a_n. \quad (1.6)$$

This formalism allows the calculation of the shape of the electromagnetic field of a focused laser beam that has been scattered by a dielectric microsphere confined near the focus. By analyzing the change in optical momentum between the incident and scattered beams, it is possible to infer the optical forces that would be applied to a trapped microsphere. Some details of this calculation are discussed in Appendix A.

1.3.1 Example Result from the Optical Tweezers Toolbox

Although a number of optical tweezer simulations exist, our focus is restricted to a particularly well-commented code library that has proven to be both useful and quantitatively accurate: the Optical Tweezers Computational Toolbox [13]. This code library is a robust implementation of the T -matrix method described in the previous section and Appendix A, and the developers have numerous publications detailing the methods, as well as the relevant advantages and drawbacks [13–15].² The library is made freely available as a collection of MATLAB functions in which the user can specify arbitrary particle shapes, particle positions, and input beam modes. By simulating the interaction between the defined beam and sphere for a range of relative displacements about the focal point, the scattering behavior of a trapped microsphere can be characterized. Furthermore, by examining the incoming and outgoing momentum flux of the incident and scattered optical fields, respectively, any change in optical momentum can be used to infer optical forces applied to the microsphere for a particular relative displacement.

Importantly, the results from a specific utility function are considered, wherein a user defines input beam and scatterer parameters as well as the desired positions over which to perform the computation, and then the toolbox calculates a “trapping efficiency parameter”, ξ , which can be both positive and negative. This quantity is a measure of the strength and

²These are the three major works directly related to the fundamentals and operation of the Optical Tweezers Computational Toolbox itself, but there is a plethora background and supplementary material available in their bibliographies.

direction of the optical interaction for a particular relative displacement, where the optical force is given by $F_{\text{opt}} = \xi \mathcal{P} / n_m c$ for some assumed beam power \mathcal{P} and index of the medium n_m (in vacuum $n_m = 1$). For example, consider a Cartesian coordinate system and let z be along the optical axis. If $\xi_z(z) > 0$, then the axial (z) optical force points in the $+z$ direction, and vice versa for $\xi_z(z) < 0$. Thus, if ξ_{x_i} crosses from positive to negative for an increase in the coordinate x_i , then the zero-crossing represents an equilibrium position since displacements away from this position result in restoring forces. The trapping efficiency is usually characterized along orthogonal axes passing through the focal point, which is often near to the equilibrium.

In the work detailed in later chapters, a laser with vacuum wavelength $\lambda = 1064 \text{ nm}$ is used in order to trap amorphous silica microspheres primarily with diameter $d \approx 4.7 \text{ }\mu\text{m}$, using an optical tweezer in vacuum with low numerical aperture. However, for reasons that will become apparent shortly, these parameters are not ideal for generating sufficiently “generic” optical potentials. Instead, for the following example calculations, a silica microsphere with diameter $d \approx 1.6 \text{ }\mu\text{m}$ and index of refraction $n \approx 1.33$ is used, which is less than that of fused silica,³ confined to an optical trap generated by a linearly polarized, paraxial TEM₀₀ mode focused with a numerical aperture of NA = 0.9.

With these parameters, the axial trapping efficiency is computed and shown in Figure 1.5, as well as the radial trapping efficiency shown in Figure 1.6 for the Cartesian axis aligned with the incident beam polarization (the orthogonal radial axis yields identical results). The trapping efficiency curves in both figures include a secondary axis in force units, assuming $\mathcal{P} = 1 \text{ mW}$ and $n_m = 1$. The axial trapping efficiency has a clear asymmetry associated to the direction of propagation, although there is a stable equilibrium position just past the focus, as well as an unstable equilibrium slightly farther along the beam, in the direction of propagation $+z$.

If it is assumed that this optical interaction is conservative, the force can be integrated as a function of position to compute an associated potential energy U , since $F_{x_i} = -dU/dx_i$. This is shown in the lower panels of both Figures 1.5 and 1.6, with a secondary axis in units of $k_B T$, where $k_B = 1.38 \times 10^{-23} \text{ kg m}^2/\text{s}^2 \text{ K}$ is the Boltzmann constant. This allows for a naïve comparison to the expected thermal energy of a trapped particle at temperature T .

³Silica microspheres grown with the Stöber process often have porous internal structures [16–19]. This results in a lower mass density and thus a lower number density of electrons. Assuming the classical electron oscillator model for the polarizability of the material: $(n^2 - 1) \sim \xi \propto n_e$, with ξ the dielectric susceptibility and n_e the electron number density.

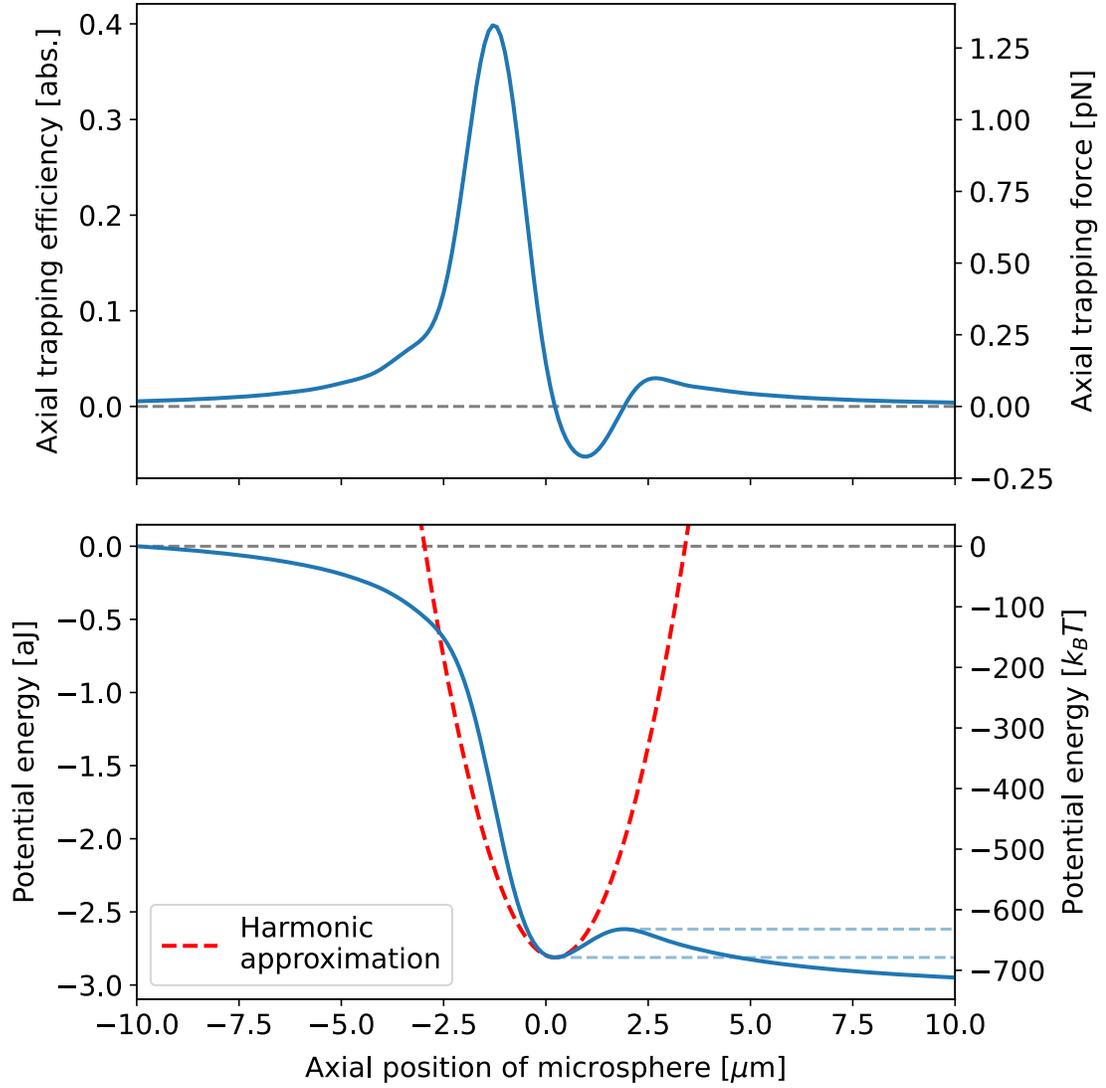


Figure 1.5: (upper) Axial trapping efficiency for a microsphere (radius $r_{\text{ms}} \approx 0.8 \mu\text{m}$, index of refraction $n_{\text{ms}} = 1.33$) within an optical trap generated by a $\lambda = 1064 \text{ nm}$ beam in vacuum, focused with a numerical aperture $\text{NA} = 0.9$ and propagating in the $+z$ direction. The resultant optical force, shown by the secondary axis, corresponds to a total beam power of $\mathcal{P} = 1 \text{ mW}$. (lower) Associated potential energy, assuming the optical force is conservative, plotted with a harmonic approximation generated from the slope of the trapping efficiency curve at the axial equilibrium position. A secondary axis shows the scale of the force in units of thermal energy at temperature $T = 300 \text{ K}$. Two dashed lines help guide the eye to see the effective trap depth in units of $k_B T$. Data in both panels is expected to scale linearly with the optical power \mathcal{P} , so that the trap depth and spring constants can be tuned.

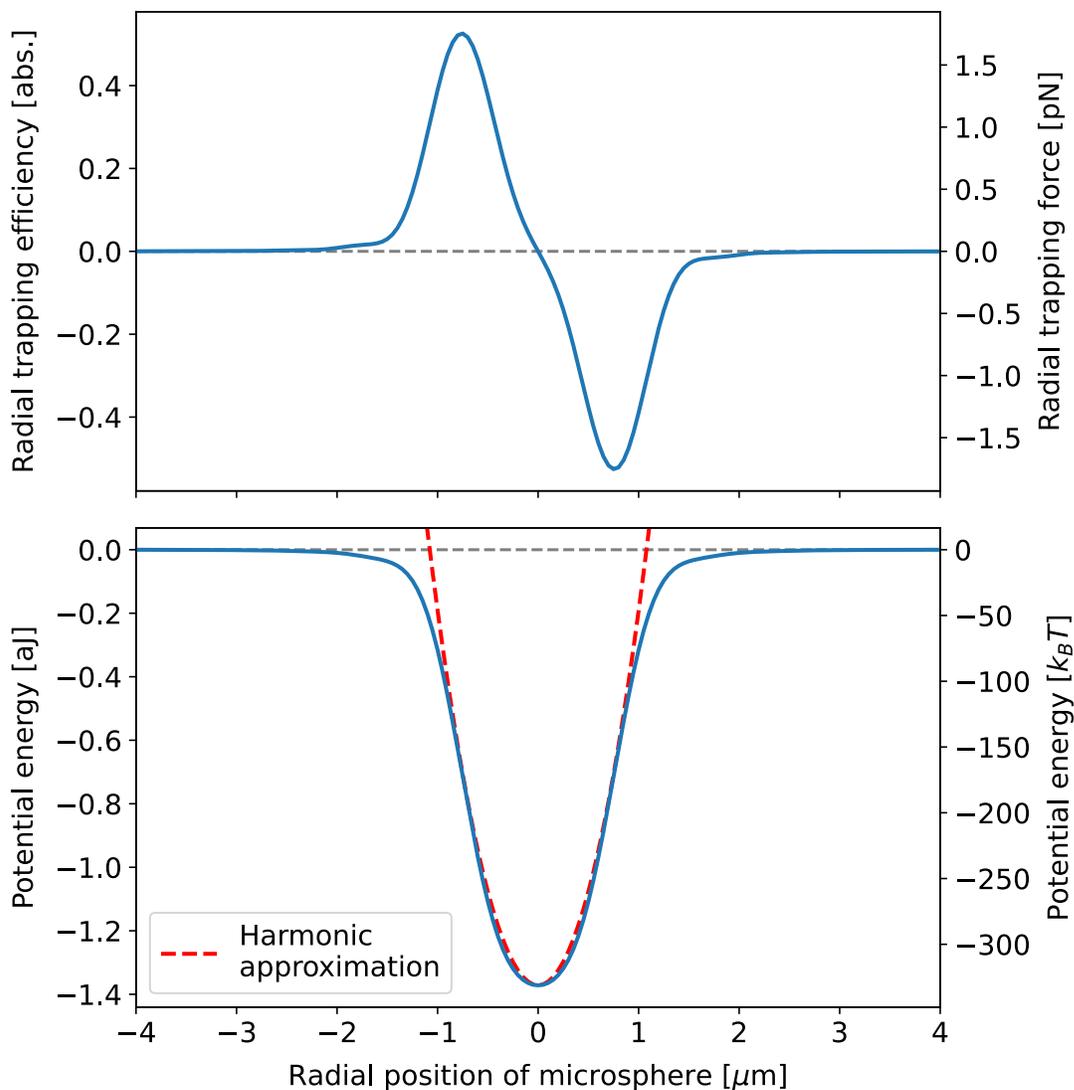


Figure 1.6: (upper) Radial trapping efficiency for a microsphere (radius $r_{\text{ms}} \approx 0.8 \mu\text{m}$, index of refraction $n_{\text{ms}} = 1.33$) within an optical trap generated by a $\lambda = 1064 \text{ nm}$ beam in vacuum, focused with a numerical aperture $\text{NA} = 0.9$ and propagating in the $+z$ direction. The resultant optical force, shown by the secondary axis, corresponds to a total beam power of $\mathcal{P} = 1 \text{ mW}$. (lower) Associated potential energy, assuming the optical force is conservative, plotted with a harmonic approximation generated from the slope of the trapping efficiency curve at the radial equilibrium position, $r = 0$. A secondary axis shows the scale of the force in units of thermal energy at temperature $T = 300 \text{ K}$. Data in both panels is expected to scale linearly with the optical power \mathcal{P} , so that the trap depth and spring constants can be tuned.

Given the asymmetry of the axial potential energy, two dashed horizontal lines indicate the local energy minimum associated with the equilibrium position, as well as the nearby local maximum, so that the trap depth can be inferred.

Further assuming motions of the microsphere will be “small” about the equilibrium, the force can be approximated as scaling linearly with displacement, suggesting a parameterization of this interaction with a Hooke’s Law relation:

$$F_{\text{opt},x_i} = -k_i x_i = -(m_{\text{ms}}\omega_i^2) x_i, \quad (1.7)$$

where k_i is an “optical spring constant” for motions along the x_i axis, and ω_i is the natural frequency of oscillation for a trapped microsphere along this axis, which is now being described as a simple harmonic oscillator. A harmonic approximation for both axial and radial curves is plotted together with the data, where the spring constants are calculated from the slope of the trapping efficiency and force curves at their equilibrium positions.

1.4 Optical Tweezers as Force Sensors

At this juncture, there is a clear application of optically trapped particles as force sensing elements. Consider a dielectric microsphere confined to the equilibrium position of an optical trap. Making the assumption of quasi-static equilibrium, an external force applied to the trapped microsphere would produce a displacement, such that the optical restoring force balances the applied force. Force sensing then becomes an exercise in displacement sensing. In this collection of work, we aim to use an optically-levitated microsphere as a force sensor, with an emphasis on short-distance force detection of interactions that couple to mass. Under this paradigm, mechanical access to the trapping region, without significantly perturbing the optical fields, is of paramount importance in order to excite possible interactions of interest without generating background forces or producing anomalous scattering of light. This places some constraints on possible imaging schemes.

Displacement of an optically trapped microsphere can be observed in a number of ways. A common practice is the use of auxiliary imaging beams, wherein motion of the microsphere scatters light from an imaging beam(s), which is then used to infer the microsphere’s motion. However, the use of auxiliary imaging beams has some fundamental drawbacks, most notably that the beams themselves perturb the optical trap and necessarily apply forces to the

trapped particle. Additionally, any imaging beam that is not coaxial with the trapping beam limits mechanical access, since any source mass introduced to excite an interaction must also not block/interfere with either the imaging or trapping beams.

Instead, a more direct approach will be used, where the displacement of the trapped microsphere, and thus the applied force, will be inferred by measuring by the scattering of the trapping beam itself. For example, as was argued from a geometric optics perspective before, a radial displacement (read: externally applied radial force) will deflect the light transmitted through the trapped microsphere. If the beam transmitted through the trapped microsphere is collimated, then this deflection will be converted to a displacement of the collimated beam, which can easily be measured with segmented photodetectors. This simple conceptualization is shown schematically in Figure 1.7, where a paraxial beam has been treated as a single ray, and a force applied to a trapped particle induces a deflection of the transmitted light and thus a displacement of the collimated light.

The situation is slightly more complicated for forces applied in the axial direction, although there are a few commonly used methods. As a dielectric microsphere moves axially, it changes the divergence of transmitted light. If an aperture is then positioned downstream, the amount of light transmitted through the aperture depends on the divergence of the outgoing beam, and thus depends on the axial position of the trapped microsphere. One can similarly consider the light retroreflected by the trapped microsphere. Not only does the divergence of the retroreflected light change depending on the axial position of the microsphere, but the *phase* of the retroreflected light also changes. It is this latter property that will be exploited to measure the axial position in this work.

1.5 Vertically Oriented Optical Tweezers

A clear requirement of any apparatus aiming to probe forces and interactions that couple to mass at short-distance, is that sources of the interaction to be probed need to be brought in proximity to a trapped microsphere. This must be done without significantly perturbing the trapping beam, which can simultaneously affect the behavior of a trapped microsphere, as well as produce anomalous scattering of light that could contaminate attempted measurements. The difficulty is compounded if the trapping beam is highly convergent, where any mechanical device intended to be in proximity to a microsphere (assumed to be trapped at or near to the focus) must be narrow along the dimension parallel to the optical axis, given

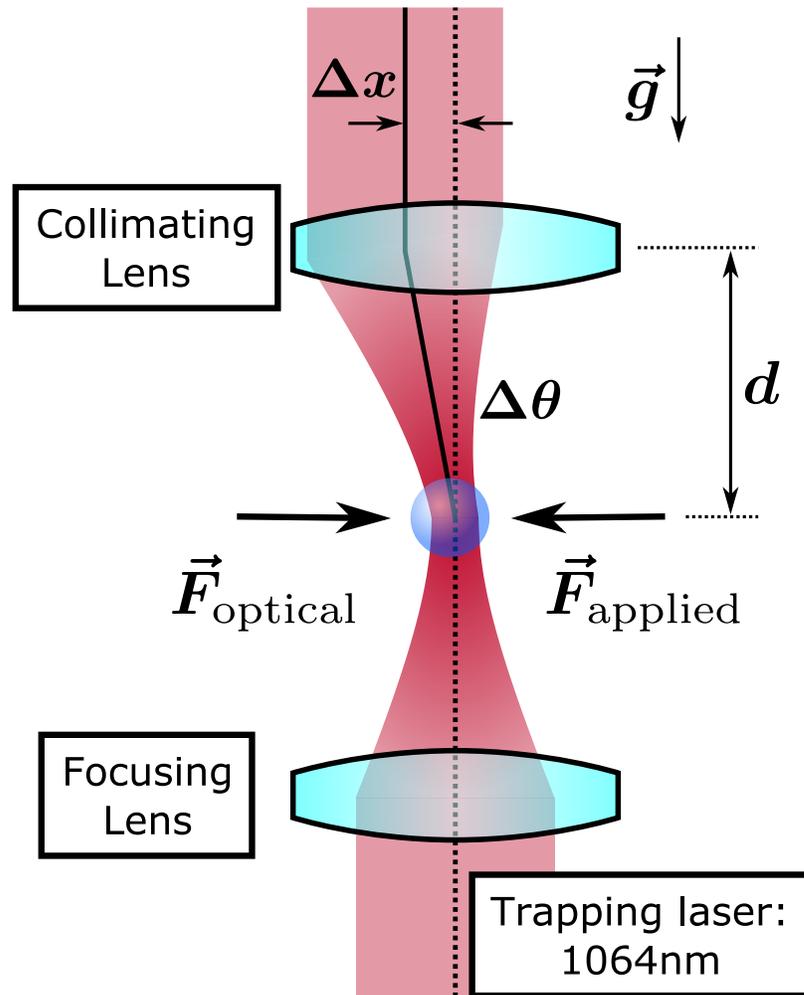


Figure 1.7: A vertically oriented optical trap, with a force applied to a microsphere confined within the trap. The resulting deflection of the trapping beam, $\Delta\theta$ is recollimated by a lens of focal length d to a displacement Δx , which can be measured to infer the applied force. Here it is assumed the optical power has been tuned to levitate the microsphere near to focus.

that the lateral extent of the trapping beam increases rapidly as you move away from the focus. Generally speaking, a highly convergent trap also limits the amount of space that can be dedicated to electrostatic shielding, since lenses with high NA usually have short working distances.

The most naïve solution to make this problem tractable is the use of a trapping beam with a low NA, so that there is an ample working distance between the physical surface of any lenses and the trapped particle, as well as a long Rayleigh range about the focus.⁴ The long working distance allows for effective electrostatic shielding of the trapping region, and more importantly, when combined with the long Rayleigh range, there is sufficient space in which to bring mechanical objects close to a trapped microsphere. One might immediately recall, however, that with a single beam that has low NA, there is no optically-induced equilibrium point in the axial direction for a dielectric microsphere. The interaction between such a beam and a microsphere tends to confine the microsphere radially to the optical axis, whilst exerting a radiation pressure force along the direction of propagation, as shown schematically in Figure 1.2.

Appealing to one of Ashkin’s initial solutions, a fully three-dimensional trap can be constructed with a low NA trapping beam by orienting the beam such that it propagates vertically upward, so that the axial radiation pressure force opposes gravity, and an axial equilibrium position exists. As the trapping beam is focused to a spot, the intensity, and by extension the radiation pressure force, is maximized at the focus. Thus, the axial position at which gravity is equal in magnitude to the radiation pressure force can be controlled by tuning the optical power. One might then expect that the optical power required for levitation is minimized at or near to the focus, where the lateral extent of the beam is minimized and the closest mechanical approach is possible.

1.5.1 Levitation via Radiation Pressure

The power required to levitate a microsphere against the force of Earth’s gravity with a vertically-oriented, low NA trapping beam can be calculated with the Optical Tweezers Computational Toolbox, but it is pedagogical to first consider order of magnitude arguments. As before, optical forces associated to the interaction between a focused laser beam and a dielectric microsphere, such as those required to levitate a microsphere against gravity, are

⁴Beams with a low NA can also be more readily approximated as paraxial Gaussian beams, so that the concept of “Rayleigh range” is appropriate in this context.

induced by changes in photon momentum from their scattering. In a laser beam with total optical power \mathcal{P}_{opt} , there is a momentum flux of $\mathcal{P}_{\text{opt}}/c$ [4, 5, 20]. In the limit of plane wave illumination, perfect reflection, and a planar reflector, the change in momentum flux from a reflection of this beam induces an axial force equal to $2\mathcal{P}_{\text{opt}}/c$, which is a strict upper limit on the levitation force a beam with power \mathcal{P}_{opt} can exert.

In most realistic situations, the majority of the trapping beam is transmitted through a trapped microsphere, as the microsphere is neither planar, nor does a vacuum-to-silica material boundary reflect significantly at typical operating wavelengths in the visible and near-infrared. Thus, the levitation force is usually parameterized as $F_{\text{opt}} = \xi_z \mathcal{P}_{\text{opt}}/c$ where $\xi \sim 0.1$ is a reasonable order of magnitude for relevant microsphere and beam parameters. Additionally, the complexities of Mie Scattering contribute higher-order effects, such that this efficiency ξ_z is not monotonic as a function of microsphere radius, microsphere index of refraction, or beam NA, to name a few parameters. These effects can be understood by using the Optical Tweezers Computational Toolbox.

Consider a silica microsphere with radius $r_{\text{ms}} = 2.35 \mu\text{m}$, mass $m_{\text{ms}} = 84 \text{ pg}$, and index of refraction $n_{\text{ms}} = 1.33$ at $\lambda = 1064 \text{ nm}$, which are realistic values for a particular batch of microspheres that were used in nearly all of the results presented in later chapters.⁵ Using the Optical Tweezers Computational Toolbox and assuming an optical trap in vacuum, with wavelength $\lambda = 1064 \text{ nm}$ and focused with a numerical aperture $\text{NA} = 0.12$, the axial trapping efficiency ξ_z can be calculated as a function of microsphere position within the beam, allowing an inference of the required beam power to levitate the sphere against gravity by setting $F_{\text{grav}} = F_{\text{opt}}$. This yields,

$$\mathcal{P}_{\text{opt}} = \frac{m_{\text{ms}} g c}{\xi_z}. \quad (1.8)$$

The result of this calculation, for the given microsphere and beam parameters, is shown in Figure 1.8. There is a minimum in required levitation power of $\sim 1.0 \text{ mW}$ just below the focus, where the asymmetry is similar to the result in Figure 1.5. One can immediately see that there is no optically-induced equilibrium position: the axial trapping efficiency ξ_z is positive definite. This is synonymous with the statement that for a microsphere anywhere along the optical axis, the force always points along the direction of propagation, at least for these particular parameters. In this case, an external force must be applied to stabilize

⁵The empirical determination of these values is detailed in Chapter 4.

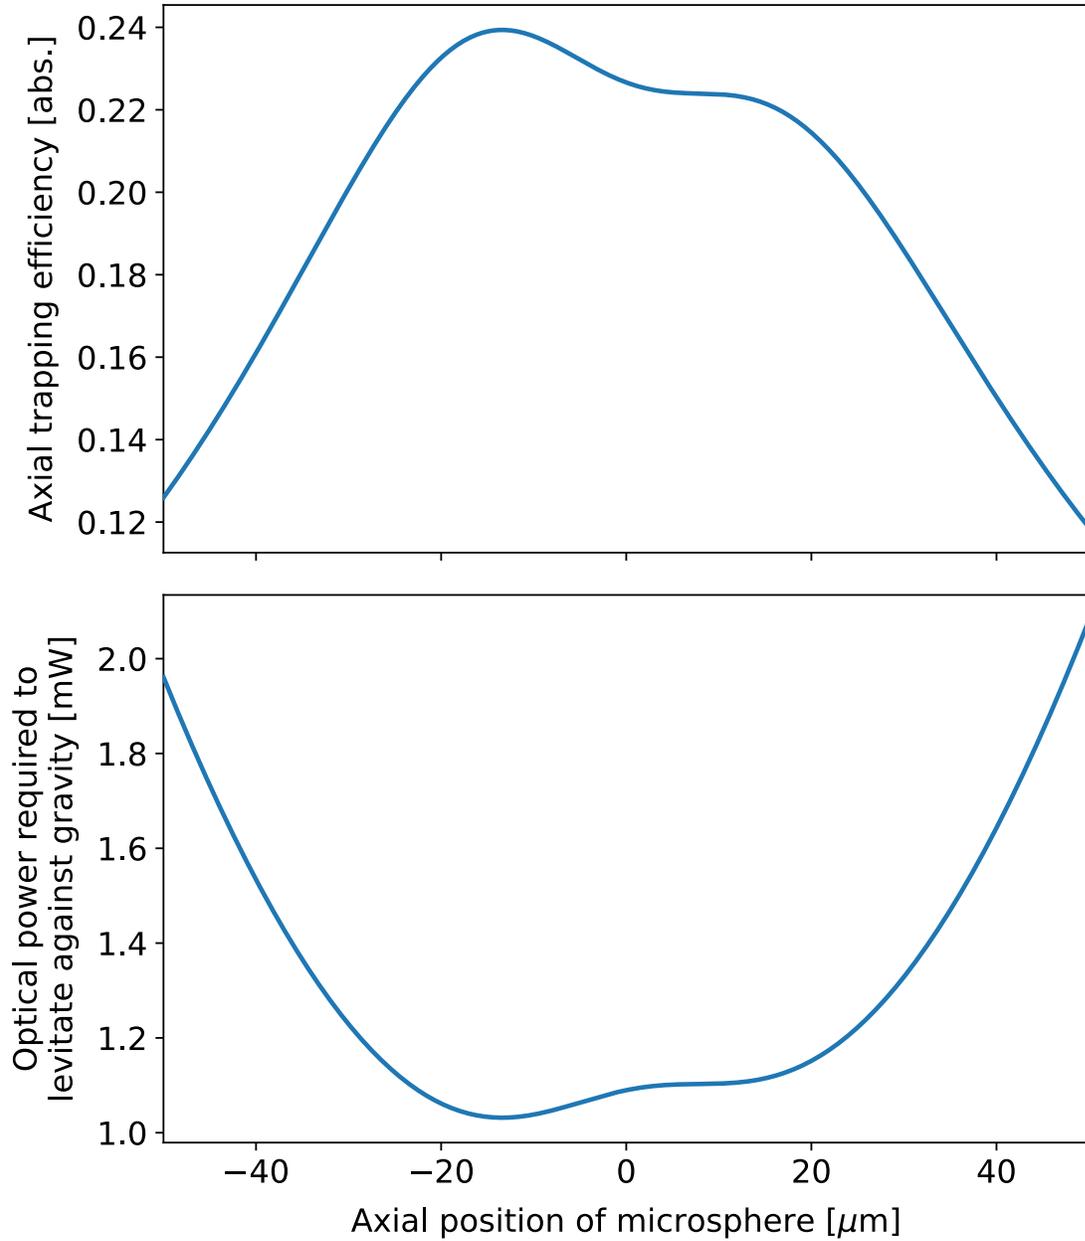


Figure 1.8: (upper) Axial trapping efficiency for a microsphere (radius $r_{\text{ms}} = 2.35 \mu\text{m}$, index of refraction $n_{\text{ms}} = 1.33$) within an optical trap generated by a $\lambda = 1064 \text{ nm}$ beam in vacuum, focused with a numerical aperture $\text{NA} = 0.12$ and propagating in the $+z$ direction. (lower) Required optical power to levitate the same microsphere, assuming a mass $m_{\text{ms}} = 84 \text{ pg}$, and with $g = 9.81 \text{ kg m/s}^2$ (See Chapter 4).

the trap axially, which can be done with Earth's gravity.

1.5.2 Radial Equilibrium

In the radial direction, the expected radial trapping efficiency can be computed just as before, but now with parameters relevant to the work at hand. Since the apparatus is intended to operate with a radially symmetric beam, it suffices to consider the trapping behavior along a single Cartesian axis, although that need not be the case in general. With the microsphere parameters ($r_{\text{ms}} = 2.35 \mu\text{m}$, $m_{\text{ms}} = 84 \text{ pg}$, and $n_{\text{ms}} = 1.33$), and input beam configuration ($\lambda = 1064 \text{ nm}$, and $\text{NA} = 0.12$), the radial trapping efficiency is shown in the upper panel of Figure 1.9, with a second y -axis scaled by $\mathcal{P}_{\text{opt}}(z = 0)/c$ to demonstrate the expected radial confinement force with the microsphere positioned at the focus of the trapping beam.

As before, the associated potential energy from integrating the force as a function of position is plotted in the lower panel of Figure 1.9, with a secondary axis in units of $k_B T$ with $T = 300 \text{ K}$. A trap depth of $\mathcal{O}(100k_B T)$ suggests that, at room temperature residual gas collisions are unlikely to eject a microsphere of the given dimensions from an optical trap with the given configuration.

The harmonic approximation clearly breaks down for large microsphere displacements of order $\Delta r \sim 0.5r_{\text{ms}}$, but is otherwise extremely useful in understanding the dynamics of an optically trapped microsphere. Certainly, in searches for new interactions involving the detection of extremely small forces, only the linear response need be considered. Further details and derivations of the harmonic oscillator response to applied forces, including damping and thermal forces from collisions with residual gas, can be found in Appendix B.

Making use of the results in that appendix, the power spectral density of the displacement of a thermally driven, optically trapped microsphere can be calculated. It is usually assumed (and occasionally demonstrated [21, 22]) that the force applied to a trapped microsphere by the momentum transfer of colliding gas particles exhibits a frequency-independent power spectral density [23]. If this thermal drive is applied to the equations of motion for a damped harmonic oscillator, the resulting motion along a single axis (x , for example) has a power spectral density, in units of m^2/Hz , of the form,

$$S_{xx}(\omega) = \frac{2k_B T}{m_{\text{ms}}} \frac{\gamma_x}{[(\omega_x^2 - \omega^2)^2 + \omega^2 \gamma_x^2]}, \quad (1.9)$$

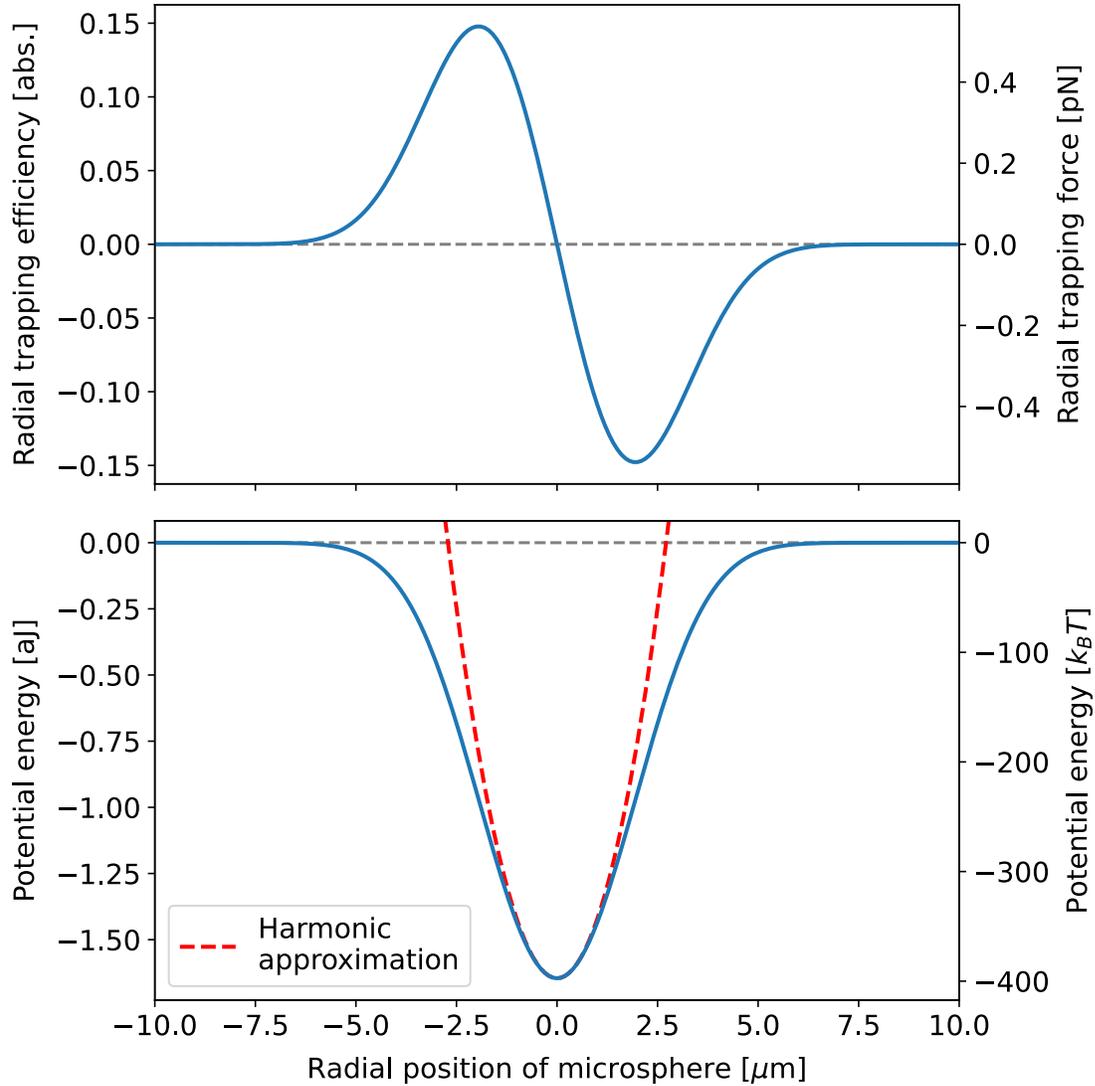


Figure 1.9: (upper) Radial trapping efficiency for a microsphere (radius $r_{\text{ms}} = 2.35 \mu\text{m}$, index of refraction $n_{\text{ms}} = 1.33$) within an optical trap generated by a $\lambda = 1064 \text{ nm}$ beam in vacuum, focused with a numerical aperture $\text{NA} = 0.12$ and propagating in the $+z$ direction. A secondary y -axis has been scaled by $\mathcal{P}_{\text{opt}}/c$, assuming an optical power of $\mathcal{P}_{\text{opt}} = 1.1 \text{ mW}$ calculated at $z = 0 \mu\text{m}$ (i.e. the focus, and based on the result shown in Figure 1.8) to demonstrate the expected optical confinement force vs radial position relative to the focus. (lower) Associated potential energy, assuming the optical force is conservative, plotted with the harmonic approximation generated from the slope of the trapping efficiency curve at $z = 0$.

written as a function of angular frequency ω , where T is the temperature of the residual gas (assumed to be in equilibrium with the trapped particle), γ_x is the damping coefficient in the x direction (see Appendix B for conventions, such as the definition/normalization of γ , given that a number of different conventions exist), and k_B , m_{ms} and ω_x are as before.

By surrounding a trapped microsphere with a non-negligible and known pressure of residual gas, it is possible to measure just such a response spectrum, allowing an empirical determination of the resonant frequency, the damping, as well as the scale of the fluctuations, the latter of which can be used to infer the ratio of the microsphere mass to the temperature. This will be discussed again in Chapters 2, 4, and 5.

1.6 Imaging of Microsphere Motion: A Heterodyne Approach

Since the displacement can be used to infer the applied force, given the analysis in the previous sections, we now consider how one might detect and measure these displacements. As was mentioned, the design of the experimental apparatus is meant to minimize the interaction between optical fields and any mechanical devices brought in proximity to a trapped microsphere. One way in which this is accomplished is by avoiding the use of auxiliary imaging beams. Then, in order to measure the displacement of a trapped microsphere, the scattering of the trapping beam itself must be considered.

This task can be partitioned into two sub-tasks: measurement of the transmitted/forward-scattered light, and measurement of the retroreflected light. It is expected that, given the complexities of Mie scattering, motion along any direction will induce a change in both the transmitted light and the reflected light, necessitating development of imaging techniques that have some degree of orthogonality *a priori*.

Orthogonality can be achieved to a certain degree by employing interferometric imaging techniques, where the light transmitted through and retroreflected by a trapped microsphere is superposed with stationary reference beams. In the axial direction, the phase of retroreflected light is modulated by axial motions of the microsphere, which immediately suggests a phase-sensitive imaging technique may be appropriate. In the radial directions, the transmitted trapping beam is deflected by radial motions (collimated to displacements), which in turn changes the spatial distribution of the transmitted light. By interfering the transmitted light with a stationary reference beam, an amplitude-sensitive interference technique is appropriate. For small displacements about an equilibrium position, axial motions

should not deflect the transmitted light in either radial direction, and radial motions should not alter the phase of the retroreflected light.

1.6.1 Reflected Light Interferometry

First, let us consider motion along the axial direction. Some amount of the trapping light is retroreflected by the microsphere, primarily by the bottom surface, in addition to other high-angle scatter. As a trapped microsphere moves up and down within the optical trap (i.e. along the optical axis), the path length of the retroreflected light changes, which is necessarily encoded in the phase of the retroreflected light.

Consider a planar reflector being translated axially by some distance δz . One might expect a phase change $\delta\phi$ of the light retroreflected by this planar device to be given by,

$$\delta\phi = 2\pi \left(\frac{2\delta z}{\lambda} \right), \quad (1.10)$$

with λ being the wavelength of the trapping light. The factor of 2 in the numerator is due to the reflection, as any motion δz induces twice that value in the change of the optical path-length. In practice, this relation does not appear to hold exactly for retroreflected light from microspheres. Instead, an empirical calibration of every degree of freedom is performed, as discussed in Chapter 2, although it is pedagogical to consider an ideal case.

Let $I(t)$ be the intensity of the superposition projected onto a photodiode. The intensity is necessarily a function of the radial coordinates x and y , but these arguments are suppressed given that the photodiode performs a spatial integral over the intensity impinging on its active surface. This results in a geometric factor that sets the overall amplitude of the interference, but has little effect on the phase. Proceeding heuristically, let's consider each beam (i.e. the retroreflected light and a reference beam) in the superposition as being perfectly collimated and aligned, such that they have flat wavefronts at the surface of the photodiode. Then,

$$\begin{aligned} I(t) &\propto |E_{\text{tot}}(t)|^2 = |E_{\text{trap,r}}(t) + E_{\text{ref}}(t)|^2 \\ &= |E_{\text{trap,r}}e^{i(\omega_t t + \phi_t + \delta\phi)} + E_{\text{ref}}e^{i(\omega_r t + \phi_r)}|^2, \end{aligned} \quad (1.11)$$

where $E_{\text{trap,r}}$ and E_{ref} (both functions of the beam radial coordinates) are the electric fields of the reflected trapping beam and the reference beam at the surface of the photodiode,

respectively, $\omega_t = \omega_{\text{opt}} + \omega_1$ and $\omega_r = \omega_{\text{opt}} + \omega_2$ are the frequencies of the trapping and reference beams including frequency shifts $\omega_1 = 2\pi(150 \text{ MHz})$ and $\omega_2 = 2\pi(149.375 \text{ MHz})$ introduced by acousto-optic modulators, ϕ_t and ϕ_r are fixed, but arbitrary, phases of the trapping and reference beams (the “arms” of the interferometer do not have their lengths controlled), and $\delta\phi$ is the phase shift induced by motion of the microsphere. Ignoring the constant terms, the intensity of the interference is

$$\begin{aligned} I_{\text{IF}}(t) &\propto E_{\text{trap,r}} E_{\text{ref}} \operatorname{Re} \left[e^{i(\Delta\omega t + \Delta\phi + \delta\phi)} + e^{-i(\Delta\omega t + \Delta\phi + \delta\phi)} \right] \\ &\propto 2E_{\text{trap,r}} E_{\text{ref}} \cos(\Delta\omega t + \Delta\phi + \delta\phi), \end{aligned} \quad (1.12)$$

where $\Delta\omega = 625 \text{ kHz}$ is the heterodyne interference frequency, and $\Delta\phi$ is a fixed phase difference between the interfering optical fields. The photodiode performs a spatial integral over the $E_{\text{trap,r}} E_{\text{ref}}$ term, the resulting magnitude of which depends on how well the spatial modes of the two beams coincide, and converts the oscillating optical power, $\int dA I_{\text{IF}}$, to an oscillating photocurrent, $i_{\text{photo}}(t) \sim i_{\text{photo},0} \cos[\Delta\omega t + \Delta\phi + \delta\phi(z)]$. If the phase $\delta\phi$ of this oscillating photocurrent can be measured, then the microsphere axial position can be inferred. The DC terms dropped when writing Equation 1.12 are also present on the photodiode, so that overall, there is never a negative intensity on the photodiode. Indeed, with perfect mode matching of two ideal and symmetric Gaussian beams, the interfering intensity modulates between zero and four times the individual beam intensity (see Appendix C).

As will be discussed in Chapter 2 and Appendix D, the oscillating photocurrent generated by interference of the retroreflected light is first amplified to a voltage, and then digitally demodulated with a phase-locked ADC in order to recover the phase of the oscillating photocurrent ($\Delta\phi + \delta\phi(z)$), where $\Delta\phi$ is assumed to be relatively constant, but can fluctuate if the path length of either the reference beam or the retroreflected light fluctuates. An empirical calibration then allows $\delta\phi$ to be related directly to an applied axial force on a trapped microsphere. For this phase measurement of the retroreflected light, the amplitude of the interference photocurrent is not a critical quantity: it must be large enough that the phase can be determined with some degree of robustness, which may depend on broadband electrical noise, as well as potential amplitude fluctuations of the signal oscillating at $\Delta\omega$.

1.6.2 Transmitted Light Interferometry

Much of the trapping beam light is transmitted through a trapped microsphere, rather than reflected. As the microsphere moves radially within the optical trap, it deflects some of the transmitted light, which, as the reader may recall, is the source of the restoring force which keeps the microsphere confined. A converging lens above the trap converts these beam deflections into displacements, and the displaced beam is made to interfere with a stationary reference beam. If the transmitted trapping beam and the reference beam are perfectly aligned to the center of a position-sensitive photodetector (for example, a quadrant photodiode), then one naïvely expects the interference photocurrent to be centered and equally distributed. As the transmitted beam is displaced relative to the reference beam by radial motions of the microsphere, the distribution of interfering optical fields, and thus interference photocurrent, is also displaced, which can then be measured.

Naively, a highly-pixelated detector such as a camera would be ideal for this purpose, as the user would then be able to observe any changes in the interference pattern, which one might expect contains more information, in addition to the expected deflection of the centroid of the interference. However, nearly all commercially available cameras are not well-suited to operate continuously at ~ 500 kHz near the interference frequency. Continuous operation is required, as microspheres have been found to require active feedback under high vacuum conditions in order to remain in the trap for long periods of time [19, 24–29]. Thus, we must appeal to higher-bandwidth solutions such as quadrant photodiodes.

To simplify the problem, consider an infinite photodetector, segmented into two halves. Let $\vec{x} = \{x, y\}$ be a Cartesian coordinate system for positions on the photodetector, with the two segments corresponding to $x < 0$ and $x > 0$. Assume two superposed beams, the transmitted light that has been perturbed by the microsphere and a second reference beam, with ideal Gaussian intensity distributions and perfect alignment such that any impinging optical fields have a flat wavefront at the surface of the photodetector. At its focus, a TEM₀₀ Gaussian beam has the form,

$$\begin{aligned} \vec{E}(\vec{x}, t) &= E(\vec{x}) \cos[\Phi(t)] = \text{Re} \left(E(\vec{x}) e^{i\Phi(t)} \hat{p} \right) \\ \text{with } E(\vec{x}) &= e^{-\frac{|\vec{x}|^2}{w^2}} \quad \text{and} \quad \Phi(t) = \omega t + \phi, \end{aligned} \tag{1.13}$$

where $E(\vec{x})$ is the spatial mode of electric field with peak value E and beam waist w (the radius at which the intensity has dropped by $1/e^2$); $\Phi(t)$ is the temporal phase, with

$\omega = 2\pi c/n\lambda$ the optical frequency, c the speed of light, n the index of refraction, λ the vacuum wavelength, and ϕ an arbitrary constant phase (included for generality); and \hat{p} is the polarization of the beam, assumed orthogonal to the direction of propagation.

For brevity, define numeric subscripts 1 = trap and 2 = ref for the transmitted trapping beam and reference beam, respectively. As before, the intensity of the superposition can be computed as the square of the total electric field,

$$\begin{aligned} I(\vec{x}, t) &= \frac{1}{2\eta} \left| \vec{E}_1(\vec{x}, t) + \vec{E}_2(\vec{x}, t) \right|^2 \\ &= \frac{1}{2\eta} \left\{ E_1^2 \exp\left(\frac{-2|\vec{x}|^2}{w_1^2}\right) + E_2^2 \exp\left(\frac{-2|\vec{x}|^2}{w_2^2}\right) \right. \\ &\quad \left. + 2E_1E_2 \exp\left(\frac{-|\vec{x}|^2}{w_1^2} - \frac{|\vec{x}|^2}{w_2^2}\right) \cos[(\omega_1 - \omega_2)t + (\phi_1 - \phi_2)] \right\}, \end{aligned} \quad (1.14)$$

where $\eta = 1/c\epsilon$ is the wave impedance, c is the speed of light, and ϵ is the dielectric constant. The factor of (1/2) in the first line comes from an implicit time average over the terms oscillating at $\sim 2\omega_{\text{opt}}$ (see Appendix C). Now, let's apply an offset δx in the x direction for the trapping beam, which is assumed to be the result of microsphere motion, and compute the spatial integral over one of the two segments of the photodetector, $x < 0$. Ignoring terms at DC, the total power of the interference on one half of our detector is

$$\begin{aligned} \mathcal{P}_{\text{IF}}(t | x < 0) &= \int_{-\infty}^0 dx \int_{-\infty}^{\infty} dy I_{\text{IF}}(\vec{x}, t) \\ &= \frac{E_1E_2}{\eta} \cos(\Delta\omega t + \Delta\phi) \int_{-\infty}^0 dx \int_{-\infty}^{\infty} dy \exp\left(\frac{-|\vec{x} + \delta x \hat{x}|^2}{w_1^2} - \frac{|\vec{x}|^2}{w_2^2}\right) \\ &= \frac{E_1E_2}{\eta} \cos(\Delta\omega t + \Delta\phi) \left\{ \frac{\pi w_1^2 w_2^2}{2(w_1^2 + w_2^2)} e^{-\frac{\delta x^2}{(w_1^2 + w_2^2)}} \left[1 + \text{Erf}\left(\frac{\delta x}{w_1^2} \sqrt{\frac{w_1^2 w_2^2}{w_1^2 + w_2^2}}\right) \right] \right\} \\ &\approx \frac{\pi w_1^2 w_2^2 E_1 E_2}{2\eta(w_1^2 + w_2^2)} \cos(\Delta\omega t + \Delta\phi) \left(1 + \frac{2\delta x}{\sqrt{\pi} w_1} \sqrt{\frac{w_2^2}{w_1^2 + w_2^2}} \right), \end{aligned} \quad (1.15)$$

where we've assumed δx is small compared to the beam diameters, and thus expanded the exponential and the error functions, retaining only terms linear in δx . Since the segmented photodetector consists of two half-infinite planes, an increase in \mathcal{P}_{IF} on one segment is mirrored by a decrease in the other, naively assuming that the total \mathcal{P}_{IF} is approximately constant for small δx . This implies that $\delta\mathcal{P}_{\text{IF}}(t, \delta x | x < 0) \approx -\delta\mathcal{P}_{\text{IF}}(t, \delta x | x > 0)$. Then,

the difference in interfering optical power between the two halves can be calculated as,

$$\frac{\delta\mathcal{P}_{\text{IF}}(t)}{\delta x} = \frac{2\sqrt{\pi}w_1w_2^3E_1E_2}{\eta(w_1^2 + w_2^2)^{3/2}} \cos(\Delta\omega t + \Delta\phi). \quad (1.16)$$

It's not particularly convenient to express this in terms of the peak electric fields E_1 and E_2 . The time-averaged power of an ideal Gaussian beam can be related to the electric field as $\mathcal{P} = \pi w^2 E^2 / 4\eta$ (see Appendix C). Substituting this expression into Equation (1.16) for each beam,

$$\frac{\delta\mathcal{P}_{\text{IF}}(t)}{\delta x} = \frac{8w_2^2}{(w_1^2 + w_2^2)^{3/2}} \sqrt{\frac{\mathcal{P}_1\mathcal{P}_2}{\pi}} \cos(\Delta\omega t + \Delta\phi). \quad (1.17)$$

And thus, we've derived a relation between the difference in interfering optical power $\delta\mathcal{P}_{\text{IF}}$ between two halves of a segmented photodetector as a function of the incident beam's lateral displacement and some assumed beam shape and beam power parameters. The beam displacement δx is related to the beam deflection $\delta\theta$, which is in turn related to the microsphere displacement δx_{ms} and thus the optical force holding the microsphere in quasi-equilibrium against an externally applied force. Assuming the trapped microsphere behaves as a harmonic oscillator, this can be summed up as,

$$\begin{aligned} \frac{\delta\mathcal{P}_{\text{IF}}(t)}{\delta x_{\text{ms}}} &= \frac{\delta F_{\text{opt}}}{\delta x_{\text{ms}}} \cdot \frac{\delta\theta}{\delta F_{\text{opt}}} \cdot \frac{\delta x}{\delta\theta} \cdot \frac{\delta\mathcal{P}_{\text{IF}}(t)}{\delta x} \\ &= (-k_{\text{opt},x}) \cdot \left(\frac{\xi\mathcal{P}_1}{c}\right) \cdot \left(\frac{1}{d}\right) \cdot \frac{\delta\mathcal{P}_{\text{IF}}(t)}{\delta x} \\ &= \frac{-8\xi\mathcal{P}_1k_{\text{opt},x}w_2^2}{cd(w_1^2 + w_2^2)^{3/2}} \sqrt{\frac{\mathcal{P}_1\mathcal{P}_2}{\pi}} \cos(\Delta\omega t + \Delta\phi), \end{aligned} \quad (1.18)$$

where d is the focal length of the recollimating lens, and $k_{\text{opt},x}$ is the optical spring constant in the x direction. This relation implies that if the interfering optical power is measured, then the amplitude of the difference in interfering power between two segments is directly related to the microsphere displacement.

As will be discussed in Chapter 2 and Appendix D, the segmented photodetector used is a quadrant photodiode (QPD), where the incident interfering optical power is converted to a photocurrent and amplified to a voltage *independently* for each quadrant. The oscillating voltages from the quadrants are each digitally demodulated with a phase-locked ADC in order to recover their amplitudes, and lateral displacements of the microsphere can then be

inferred by differences in the amplitude of the demodulated interference across the QPD. By construction, the amplitude and phase are relatively independent, and thus phase fluctuations from changes in optical path length should not alter the demodulated amplitude, assuming the fluctuations are slower than the demodulation sampling frequency.

Although the math detailed here is useful to understand the fundamentals of the interference, the spatial modes of the transmitted trapping beam and reference beams are not perfectly Gaussian (especially the transmitted trapping beam which has been perturbed by the microsphere), the curvature of the two beams is not perfectly matched, and our actual detector is not infinite in extent. Thus, as before, an empirical calibration will be used to relate δx and δy , the horizontal displacements of the microsphere, to any applied forces.

1.7 Active Feedback Primer

The dynamics of particles confined to optical traps are necessarily dominated by the optical interactions themselves, and throughout this chapter we've made the implicit assumption that any laser beams involved are stationary apart from scattering induced by the trapped particles. Any fluctuations of the beam pointing, beam power, or even lasing frequency, some which are fundamental to the laser system itself, can apply non-negligible forces to a trapped microsphere. There are also pedestrian sources of fluctuations such as air currents around the free-space optics, impacts (both of tool and body) from users working on the optical trap, as well seismic vibrations from general human activity in and around the building, with an earthquake sprinkled in here and there.

Generally these anomalous fluctuations (i.e. not thermal in nature) are impulse-like, and would immediately eject a microsphere from the trap if sufficiently strong. However, it is also possible that both thermal and anomalous fluctuations drive a microsphere in non-conservative orbits [30, 31], accumulating energy in the center of mass degrees of freedom and eventually allowing the microsphere to escape the optical potential well. Due to the destabilizing nature of these fluctuations, optical traps often make use of active feedback to control the position of the microsphere. Often, the active feedback will only be viscous in nature, to dampen both thermal and anomalous fluctuations. By comparing real-time estimates of the microsphere position and velocity to some desired setpoints, where the real-time signals are derived from imaging techniques such as those described above, forces can be applied to stabilize the microsphere motion about the setpoints. Given that linear control

theory is well-established, the only difficulty that arises here is in the implementation of the feedback actuator itself.

A variety of feedback architectures have been implemented by others in the field, but generally involve one of two mechanisms: either an optical interaction between a trapped particle and a modulated/deflected laser beam, or an electrostatic interaction with a non-zero charge on the trapped particle. For the former, one can manipulate the optical power and beam pointing of the trapping beam itself [26, 32–34] in order to modulate the optical spring constant(s) or change the equilibrium position(s). It is also possible to use auxiliary cooling beams where the power of the cooling beams is modulated to apply optical forces [21]. In the case of electrostatic feedback, a charged and trapped particle is subjected to electric fields from nearby electrodes, which exert forces on the particle. The latter technique has been used in recent times to cool the center-of-mass motion of levitated nanoparticles [35], in some cases reaching the motional quantum ground state [36].

In the apparatus described by Chapter 2, we will take the approach of manipulating the trapping beam in order to produce feedback forces. Angular deflections of the trapping beam prior to being injected into the chamber generate radial displacements in the focal plane of the trap, so the that microsphere is pushed to the displaced equilibrium position. For the axial degree of freedom, power modulations of the beam adjust the equilibrium position through the levitation constraint $F_{\text{opt}} = F_{\text{grav}}$, operating under the assumption that $F_{\text{opt}} \propto \mathcal{P}_{\text{opt}}$. In this way, measurements of the microsphere displacements described in the previous sections can be used to inform an linear control algorithm, which in turn generates the required feedback forces. This will be discussed in more detail in Section 2.7.

Chapter 2

The Optical Levitation Apparatus at Stanford

Since approximately 2013 through 2021, there have been four major iterations of the experimental apparatus at Stanford. The first was constructed and dismantled prior to the beginning of this PhD, while the second was relatively short-lived and produced only a single result [25]. Both will thus be mentioned primarily in relation to improvements made in the third iteration. Skipping ahead, the fourth and most recent iteration (at the time of writing) has been described in great detail in Reference [37], and has produced the first measurement of non-Newtonian gravity with optically levitated microspheres.

Almost the entirety of the work presented throughout this thesis was performed on the third iteration of the apparatus, which has been partially described in Refs. [19, 26–28, 38]. For the sake of completeness, this chapter will provide all of the pertinent details, and will be referred to heavily throughout the remainder of this work. It should be noted that unlike many of the other chapters, this one will contain a significant amount of heuristic logic and empirical evidence for the “proper” operation of the apparatus. Much of the material is original and has been written solely for this thesis, although some specific sections may be adapted (or simply copied) from Reference [26], for which the author was a primary contributor.

The optical trap itself is housed within a vacuum chamber that constitutes the centerpiece of the experiment. The light that creates the trap originates in a fiber-optic laser system, and is subsequently directed and controlled with a collection of free-space optics, deemed the input optics, before being projected into the vacuum chamber. Photons that are transmitted

through and reflected by the trapped microsphere are collected both by a separate set of free-space optics, deemed the output optics, and the input optics, respectively, in order to monitor the position of a trapped microsphere.

2.1 The Experimental Chamber

The main experimental chamber is a “Spherical Cube” vacuum vessel (Kimball Physics, MCF600-SphCube-F6C8 [39]), with six 6.00” ConFlat (CF) ports and eight 2.75” CF ports allowing for a variety of optical and mechanical access points, vacuum pump ports, and electrical feedthroughs. The bottom 6” CF port has a blank flange that serves as the baseplate for the optical trap. Two opposing 6” CF ports in the horizontal plane are used for the input and output of the trapping laser, which passes through AR coated windows (Edmund Optics, 11-904), mounted at a 5° angle to emulate the affect of a wedge window and reduce cavity affects between the window surfaces. The orthogonal pair of opposing 6” CF are used for 1) an electrical feedthrough for the primary nanopositioning stage and 2) a large quartz viewing window. The window port is also the primary access point when the chamber needs to be opened, such as during the initial preparation of a trapped microsphere, described in Section 2.6. The final top 6” CF port is connected to two capacitance manometers for pressure monitoring, discussed in detail Section 2.1.1.

Four of the smaller 2.75” CF ports are used as electrical feedthroughs for the shielding electrodes surrounding the trap, the motorized stages used to control mechanical devices in close proximity to the optical trap, as well as the piezoelectric actuator used to propel microspheres into the trap initially. The remaining four 2.75” CF ports are used for (1) the primary pumping port (discussed in Section 2.1.1), (2) a Pirani thermal conductivity vacuum gauge (Kurt J. Lesker Co., KJLC 275i), (3) a Residual Gas Analyzer (MKS, EV2-220-000FT [40]), and (4) a Teflon-based fiber-optic feedthrough [41] used to flash the optical trap with UV photons. A schematic view of the chamber, including the optical trap within, is shown in Figure 2.1, together with a photograph.

In general, CF ports that are opened frequently make use of Viton rubber gaskets, while CF ports that are rarely opened, or even closed indefinitely, make use of the usual copper gaskets. In both cases, manufacturer specifications suggest base pressures of 1×10^{-8} mbar are achievable, and in practice the base pressure of the experimental chamber is limited at a much higher pressure of $\sim 1 \times 10^{-6}$ mbar due the presence of various plastics in some of

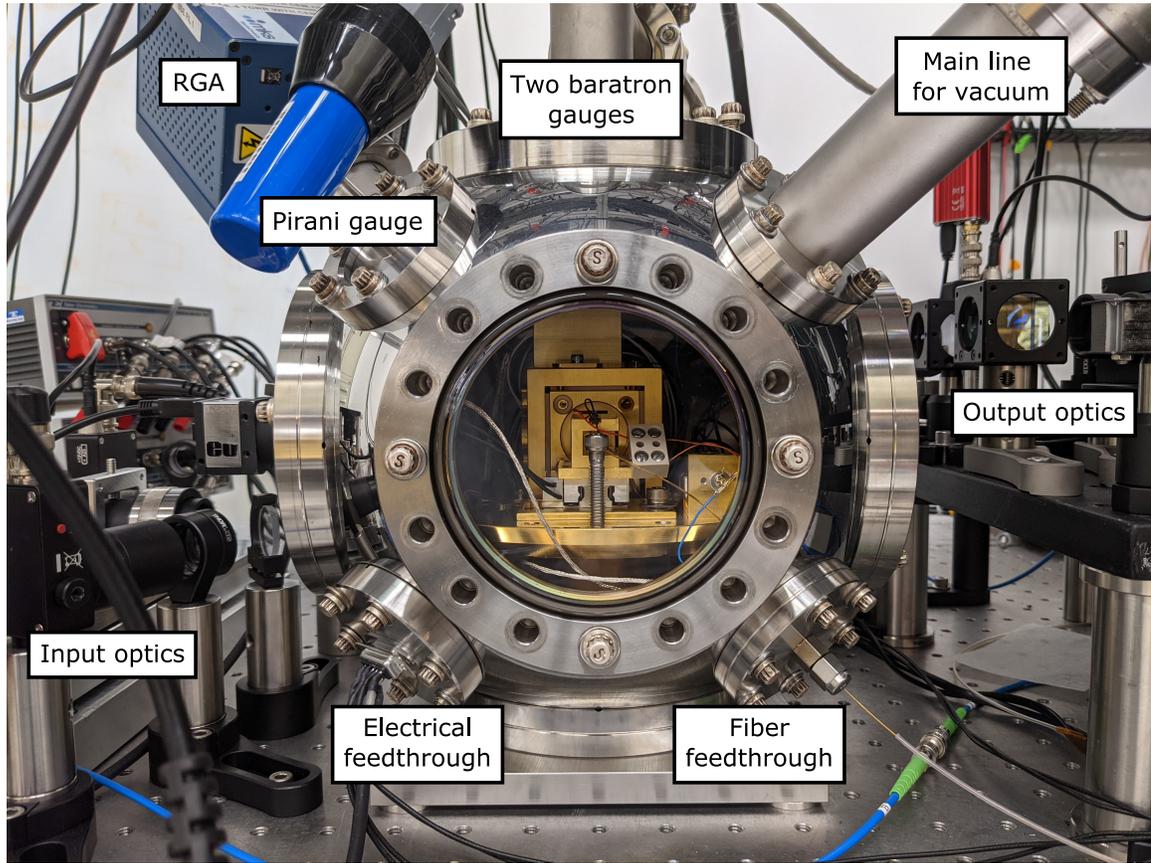


Figure 2.1: A photograph of the experimental chamber with various ports labeled. Components of the vacuum system will be discussed in Section 2.1.1, while the optical systems will be described in Sections 2.2 and 2.4. The gold-coated mechanical structure visible through the viewport is the centerpiece of the apparatus, supporting both the trapping lenses and shielding electrodes to be discussed in Section 2.3.

the cabling used.

2.1.1 Vacuum System

In order to perform sensitive measurements with optically levitated microspheres, thermal noise due to collisions with residual gas particles must be removed. It can also be advantageous to introduce specific gas species into the experimental chamber. A vacuum and gas handling system accomplishes both tasks.

High vacuum within the experimental chamber is achieved with a turbomolecular pump (Pfeiffer Vacuum, HiPace 80 [42]) whose exhaust is roughed by a dry scroll pump (Edwards

Vacuum, XDS5-C). The turbomolecular pump is first connected to a 4.5" CF gate valve, and subsequently adapted to a 2.75" CF cross. One port of the CF cross is connected to the chamber through a flexible 1 m long, ~ 38 mm inner diameter vacuum hose with 2.75" CF terminations. A 10 cm rigid pipe and a ceramic break, both with 2.75" CF terminations, adapt the hose to the experimental chamber, the latter electrically isolating it from the turbo pump.

The final two ports of the CF cross are used for (1) a bypass pumping path so that the chamber can be roughed to moderate vacuum without exposing the turbo pump to atmosphere, and (2) a computer controlled leak valve (Pfeiffer Vacuum, EVR 116) connected to a gas manifold in order to introduce specific gases into the chamber. The bypass path includes a manually actuated leak valve in series with a pneumatically actuated angle valve, so that the pumping rate can be controlled and the bypass path can be closed completely. A schematic overview of the vacuum system is shown in Figure 2.2.

The residual gas pressure is controlled and measured between $\sim 2 \times 10^{-6}$ mbar and 1 mbar. The vacuum pressure is tuned by introducing or removing N_2 gas and is measured by a cold-cathode gauge for pressures below 10^{-4} mbar, a capacitance manometer for pressures between 10^{-4} and 10^{-2} mbar, and a Pirani gauge for pressures between 10^{-3} and 1 mbar. The cold cathode gauge is found to affect the charge the MS, so it is only used to measure the $\sim 2 \times 10^{-6}$ mbar base pressure of the vacuum system after an experiment. The capacitance manometer does not cover the full range of vacuum pressures, so the Pirani gauge is calibrated against the more accurate capacitance manometer, where there is overlap. This system is capable of measuring the pressure to an accuracy of 10% for N_2 over the range of interest [40, 43].

The gas manifold connects the experimental chamber to six pressurized gas cylinders containing He, N_2 , Ar, Kr, Xe, and SF_6 . Downstream of the requisite pressure regulators, the gas cylinders are each connected by pneumatically actuated bellows valves to a common system of stainless steel tubing, deemed the gas manifold. The manifold itself is connected the back of the computer controlled leak valve so that, with the proper valves open, the contents of each gas cylinder can be individually leaked into the experimental chamber. The manifold is also connected to a dedicated dry scroll pump allowing the entire system to be evacuated prior to filling with a specific gas, in order to ensure purity. Connections between components are made with a variety of VCR, NPT, and SwageLok gaugeable tube fittings. A capacitance manometer also connected to the manifold suggests base pressures

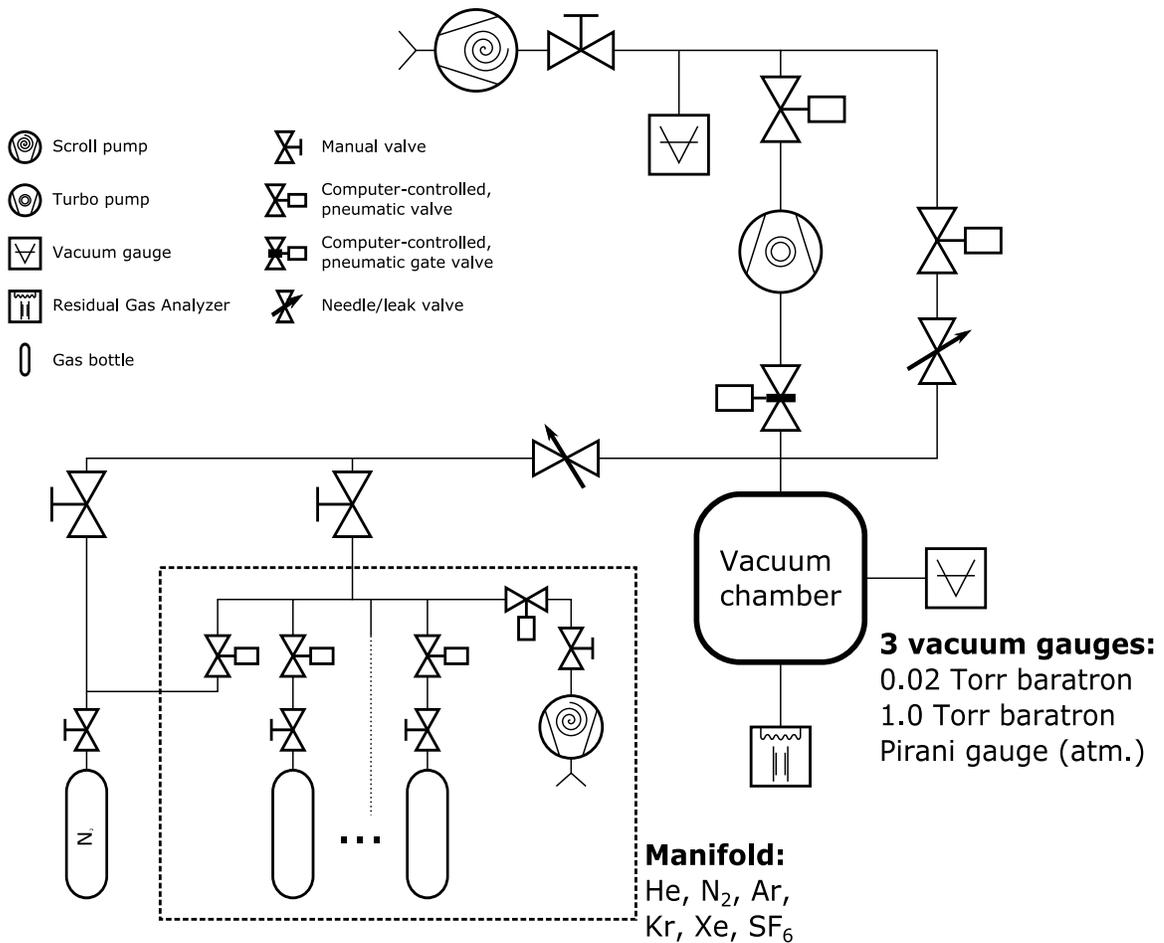


Figure 2.2: A block diagram of the vacuum system. The vacuum gauges directly on the main chamber are rated in Torr, although the measurements are usually converted to mbar, or Pa. The scroll pump for the main chamber is located outside of the lab, in a hallway closet, and connected to the vacuum system with a ~ 5 cm inner-diameter steel piping system embedded in the walls and ceiling. Similarly, the gas bottles and manifold are located outside of the lab, in the gowning area, connected with a ~ 1 cm inner-diameter steel piping system.

of $\sim 1 \times 10^{-2}$ mbar are readily achievable within the manifold, which is usually pressurized to ~ 1 bar when in use.

2.2 Source Light and Input Optics - Single-beam Trapping

With the knowledge and insight from two previous iterations of optical traps, the apparatus being described was designed with a number of specific considerations. Most significant was the advent of a single-beam trap, wherein the trapping, imaging, and feedback control are all accomplished with the same beam. This has immediate advantages as extra lasers are not required to operate the system, allowing for the mechanical access mentioned previously, as well as vast reduction in the number of optical components required.

To best make use of this configuration, an interferometric readout was proposed. Incident trapping light that is transmitted through a trapped microsphere and refracted by its motion is collected and made to interfere with a reference wavefront, while incident light that is retroreflected by the microsphere is also collected and made to interfere with a second reference wavefront. The implementation of interferometry is described in detail in Section 2.4, although it should be noted here, as the choices of laser system and steering optics were influenced by this design goal.

2.2.1 The Fiber Optic Laser System

The source light for the optical trap needs a long temporal coherence, in order to implement an interferometric readout scheme. For the same interferometry, the source light needs to be split into three different paths: one to generate the trap itself and two to use as reference wavefronts for the interferometry. Finally, the amount of light directed to the vacuum chamber to create the trap should have a maximal power of approximately 1 W (an empirically determined value), so that it is possible to catch and then levitate silica microspheres with a radius of up to $\sim 10 \mu\text{m}$ in a vertically-oriented optical trap stabilized by gravity.

Initially, the light is generated in a ytterbium-doped fiber laser (Orbits Lightwave, Eternal SlowLight) that lases at 1063 nm with an output power of ~ 70 mW. To optimize linewidth, and thus temporal coherence, the laser is operated in a constant temperature mode. Although this has not been measured, the manufacturer reports that this system should have a linewidth of ~ 1 kHz, corresponding to a temporal coherence preserved over

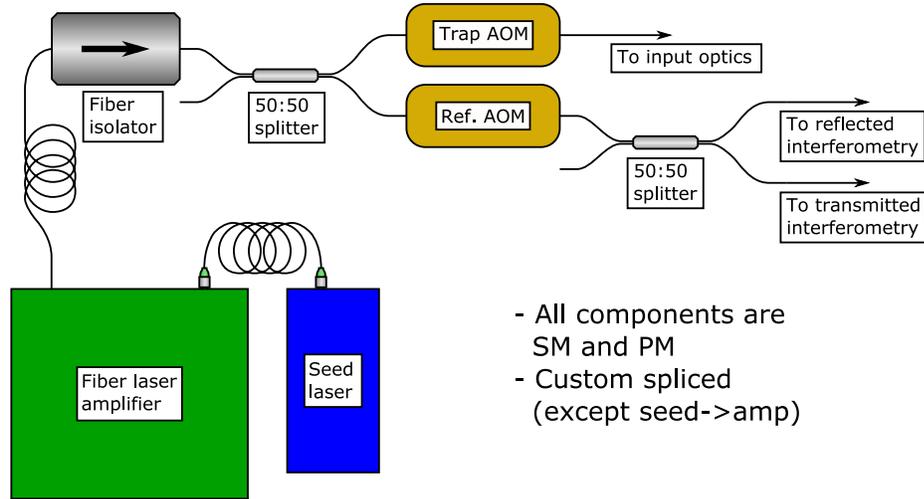


Figure 2.3: A schematic diagram of the fiber-optic system, depicting the source laser and all components up to the point of projection into free-space. The trap acousto-optic modulator (AOM) is usually operated at 150 MHz, while the reference beam AOM is operated at 149.375 MHz. All fiber components are single-mode and polarization-maintaining, and custom spliced for high power operation.

~300 km of propagation in vacuum. The light exits the laser module via a polarization-maintaining (PM), single-mode fiber (PM980-XP) terminating in a ferrule connector/angled physical contact (FC/APC) connector. The same fiber type, PM980-XP, is used throughout the remainder of the fiber optic system.

The output of the fiber laser seeds a ytterbium-doped fiber amplifier (Nufern NuAMP), with a maximal output power of 10 W, saturating the input stage to ensure operational stability. The high output power usually requires free-space operation, as 10 W of optical power within a single-mode fiber is beyond the damage threshold of typical fiber connectors (such as FC/APC). Another way to overcome this limitation is by directly splicing fibers together, avoiding fiber connectors entirely. This can be accomplished with a number of tools, but is especially difficult when the fibers are both single-mode and PM, as the cores are only a few microns in diameter, and have a specific rotational alignment along the axis of the fiber. Despite this, single-mode PM fiber splicing can be accomplished with commercially available instruments, in particular a Glass Fiber Processor (Vytran Corporation, GPX-3000), which has components for both cleaving and splicing fibers.

The output from the fiber amplifier is directly spliced to a high-power fiber-optic isolator (Thorlabs, IO-L-1064), providing an extra layer of protection at the cost of some optical

loss. The output fiber from the isolator is spliced to a 50:50 fiber-optic beamsplitter (BS) (Gooch and Housego, FFP-8MC264G10 []), with one output for the trapping beam, and the other for both reference beams. The trapping beam output is spliced to a fiber-coupled acousto-optic modulator (AOM, Gooch and Housego, T-M150-0.4C2G-3-F2P) operated at 150 MHz, whose output fiber pigtail is used to project the trapping light to free-space with a fiber mount and an objective lens, discussed in the next section. The RF port of the AOM is driven by a signal generator (Stanford Research Systems, SG382) which has been amplified (Mini-Circuits, ZHL-03-5WF+) in order to take full advantage of the dynamic range of the AOM. This will be discussed in more detail in Appendix D.

The reference beam output from the first 50:50 fiber-optic BS is spliced to an identical fiber-coupled AOM which is operated at 149.375 MHz in order to generate a differential frequency shift between the trapping and reference beams, which will be critical for the heterodyne readout discussed in Section 2.4. This AOM is driven by a separate signal generator (Hewlett-Packard, 8643A) that has also been amplified (EIN, 400AP) for the same reason. The output of this AOM is spliced to a second 50:50 fiber-optic beamsplitter to source both reference beams. The reference beam for the retroreflected light interferometry is created by coupling one of the beamsplitter outputs, a fiber pigtail, to an aspheric lens, generating a beam with a Gaussian waist of 3 mm. Both the lens and fiber are mounted in an SM1 threaded lens tube, the latter on an externally threaded fiber-optic bulkhead (Thorlabs, SM1F1-250). The other BS output is first spliced to a single-mode PM fiber patch cable (Thorlabs, P3-1064PM-FC), terminating in an FC/APC connector. The reference wavefront for the transmitted light is then created by coupling the patch cable to free space with a zoom lens collimator (Thorlabs, ZC618APC-C), which has adjustable output beam diameter and divergence.

The entirety of the fiber optic system, excepting the free-space projection, is depicted schematically in Figure 2.3.

2.2.2 Free-space Optics

In some of the initial science results with previous iterations of the apparatus [24, 25], it was suspected that extensive halos and stray light were contributing to background forces. Halos are generated from diffusive scattering, primarily in the various refractive optics used for beam expansion, as well as from repeated reflections from beam-steering and alignment mirrors. Stray light is generated from specular reflections of the trapping beam, which can

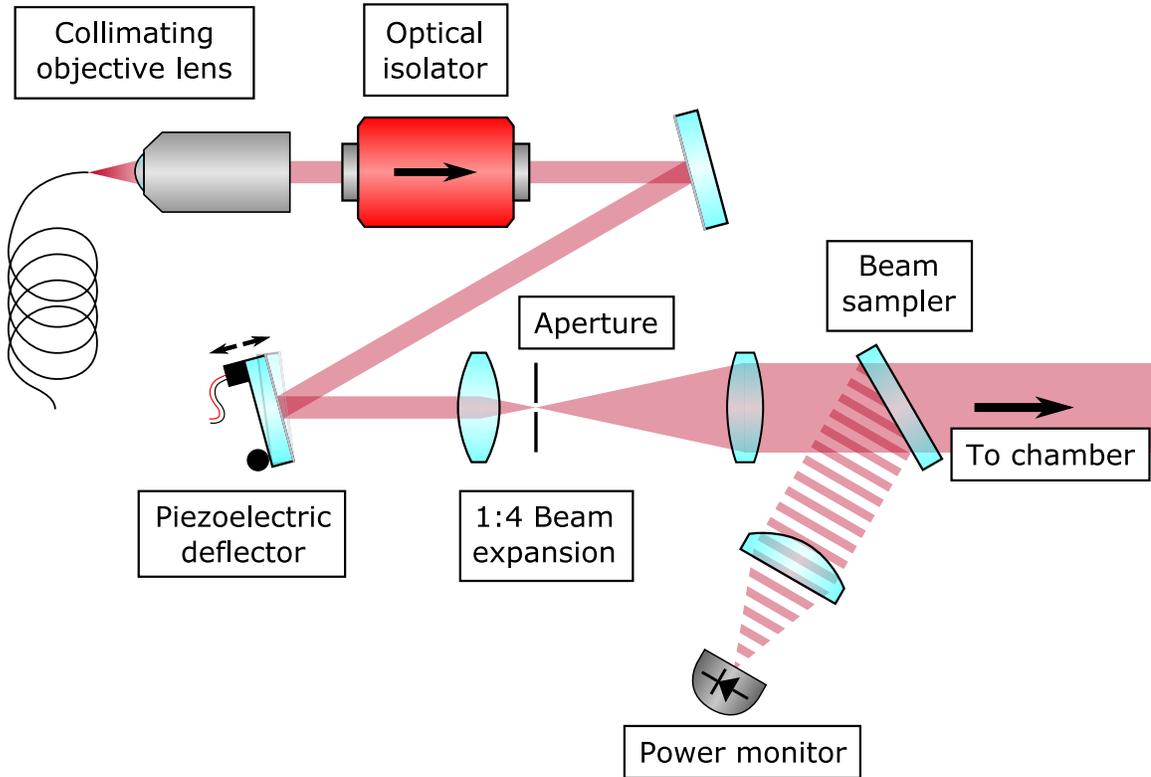


Figure 2.4: A schematic diagram of the input optics, from the point of projection into free-space, to the beam entering the experimental chamber.

trace their origin to non-ideal reflections from internal surfaces of mirrors as well as refractive interfaces, often referred to as ghost beams. With this in mind, the free-space optics were designed with a minimal number of components. A schematic depiction of the system is shown in Figure 2.4

The fiber pigtail output from the trapping beam AOM is projected to free-space and collimated with an infinity-corrected microscope objective (Thorlabs, LMH-20X-1064), which generates a beam with a Gaussian waist of $\sim 800\ \mu\text{m}$. The objective lens is threaded into an aluminum housing, and the relative alignment with the fiber pigtail is achieved with a 3-axis translational flexure stage (Newport, 466A), together with an optical fiber rotator (Newport, 466A-718) to align the output polarization. The collimated beam is sent through a free-space Faraday isolator (Electro-Optics Technology Inc, customized PAVOS Series), which includes polarizing optics to separate the forward and backward propagating photons. This is crucial for the retroreflected interferometry, which uses light reflected from

the lower surface of the microsphere in order to measure the microsphere's axial (vertical) position within the trap.

The output beam from the free-space isolator is directed toward a pair of protected silver mirrors, one 25 mm diameter for static alignment, and a second 7 mm diameter which is mounted on a piezoelectric deflector system with two degrees of freedom for active beam deflection (Thorlabs, ASM003). The first mirror is used to assure the beam is centered on the second mirror, which is then used to provide stabilizing feedback to a trapped microsphere by displacing the position of the trap to counteract forces from collisions with residual gas, pointing fluctuations of the trapping beam, and acoustic vibrations from the environment, among others. The deflection mirror is placed in a Fourier plane of the trap, so that deflections of the collimated beam produce horizontal displacements of the optical trap's focal point. In practice, deflections required are of the order of $\sim 10 \mu\text{rad}$ at a moderate vacuum of 1.5 mbar, and $\sim 100 \text{ nrad}$ at the base pressure of the experimental chamber.

Both the deflection system and the Faraday isolator have relatively small clear apertures, 5 mm for both, requiring a beam with narrow lateral extent. However, this is suboptimal for generating an optical tweezer, as it significantly limits the achievable numerical aperture with reasonable focal lengths. Thus, the beam is expanded by a factor of four with a pair of 25 mm diameter, bi-convex lenses in a refracting telescope configuration: the primary lens having a focal length of $f = 50 \text{ mm}$ and the secondary with $f = 200 \text{ mm}$ to achieve the desired magnification. The telescope lenses are co-aligned with a lens-tube and cage system, while their relative displacement along the optical axis is fine-tuned with a pair of High-Precision Zoom Housings (Thorlabs, SM1ZM).

The cage system, used as a break between two lens-tubes where each tube terminates in a zoom housing with one of the telescope lenses, allows for the placement of a pinhole at the telescope focus. The pinhole acts as a spatial filter, reducing both halo and any stray light. Although not directly measured, the spot size within the telescope should be of the order of a few microns, so that a pinhole with a $50 \mu\text{m}$ diameter does not significantly perturb the output mode, even when the beam is deflected to provide stabilizing feedback. The pinhole is aligned to the focal point of the telescope with axial adjustment via translation along the cage system, and radial adjustment facilitated by a 2-axis flexure mount (Thorlabs, CP1XY).

All of the free-space optics discussed require careful co-alignment in order to best preserve the output spatial mode from the optical fiber, which is nearly Gaussian. Once aligned to one

another, the beam entering the experimental chamber must also be carefully aligned to the aspheric trapping lenses, as these are exceptionally sensitive to misalignments, which cause a number of aberrations. The optical components themselves are all mounted on a custom 5-axis stage, so that following their co-alignment, the entire configuration can be adjusted relative to the trapping lenses. Final alignment is determined iteratively by measuring and then optimizing both the spot size and divergence angle of the optical trap, via adjustments of the 5-axis stage. The measurement technique and a typical beam profile and divergence are shown in Section 2.3.2.

Following magnification, the trapping beam passes through a $\sim 99:1$ beam-sampler (Thorlabs, BSF10-C), with the majority of the light transmitted through to the chamber input window, discussed in Section 2.1. The $\sim 1\%$ reflected from the beam sampler is directed to a power monitoring photodiode (Thorlabs, DET10A) with a 1064 nm line filter (Thorlabs, FL051064-10) and a 25 mm focal length plano-convex lens to ensure all of the incident light is collected onto the diode. The photocurrent is converted to a voltage and amplified by a transimpedance amplifier (Thorlabs, AMP120), before being sent to an ADC. By comparing this measured optical power to an arbitrarily defined reference value, the RF power driving the trapping beam AOM can be modulated to feedback on any measured amplitude fluctuations (see Section 2.7).

2.3 The Optical Trap

The optical trap is a vertically oriented, single-beam trap, which has been generally described in Chapter 1, and depicted in Figure 1.7. The input optics discussed in Section 2.2 generate a 1064 nm, horizontally polarized, horizontally propagating, and collimated beam, with a Gaussian waist of $w_0 = (3.1 \pm 0.1)$ mm. This trapping beam enters the chamber and directly impinges on a right-angle prism mirror (Thorlabs, MRA25-P01), reflecting the beam vertically toward the first trap lens. The lens that focuses the laser light to form the trap is a 25 mm diameter, 25 mm focal-length Best Form Aspheric Lens (Edmund Optics, 89-439), which results in a trap with a $w_0/f = 0.12$ numerical aperture given the measured beam waist and known focal length.¹ The diverging output beam is then recollimated by

¹The usual formula for numerical aperture of a single lens is $NA = n \sin \theta \approx nD/2f$ where D is the lens pupil diameter, f is the focal length, and n is the index of refraction of the medium (in this case vacuum and thus $n = 1$) where the approximation $\sin \theta \approx \theta$ has been made. Here, the lens pupil is not filled, and thus we approximate $D \approx 2w_0$.

an identical lens, and reflected horizontally by an identical right-angle prism mirror to exit the chamber toward the output optics.

The lenses are epoxied within retaining rings that are, in turn, supported by a pair of interlocking u-shaped housings, which fix the concentricity and distance between the lenses. Their relative displacement along the optical access can be fine tuned with shim washers. The lens housings also support six pyramidal electrodes to shield the trapping region, with bore holes for optical and mechanical access from all six directions, where the top of the pyramids are truncated and form a nominally cubical cavity around the focal point of the trap lenses. Each electrode is isolated from the housing with a MACOR spacer, allowing each to be independently biased. The housings, lens retaining rings, and electrodes are all machined in-house from aircraft grade aluminum (alloy 7075), and commercially electroplated with gold. A cross-sectional view of the lenses, electrodes, and the housing for both is shown in Figure 2.5, together with a zoomed view of the trapping region, and a photograph of a trapped microsphere, captured with a cellphone camera.

The converging input and diverging output beams propagate through the bottom and top electrodes, respectively, whose bore holes have been coated with colloidal graphite to reduce reflections from stray light. The bore holes in two of the four electrodes in the horizontal plane are dedicated for mechanical access of various fabricated devices, which are used to source electromagnetic interactions, as well as possible new interactions from theories extending the current Standard Model of particle physics (see Chapters 7 and 8). The bore hole in a third electrode is used for a rudimentary microscope with a view in the xz -plane (see Section 2.5), which allows for basic monitoring of the trapping region, such as whether a microsphere is trapped or not. The final electrode has a much narrower bore, specifically designed to house a multi-mode fiber and a ball lens, in order to deliver UV photons into the trapping region.

The upper and lower electrodes, as well as the fiber delivery electrode, are all modified with horizontal slots that allow two quartz panes to be inserted above and below the trap simultaneously with a linear translation stage (Newport, AG-LS25V6), with minimal perturbation of the optical mode. The panes, their custom mounting structures, and the corresponding slots are partially visible in the photograph in Figure 2.5. The upper quartz pane is long, narrow, and cantilevered from the surface of a piezoelectric actuator. The lower surface of this pane is covered with microspheres that adhere via van der Waals forces, and the piezoelectric actuator is used to excite mechanical modes within the quartz pane

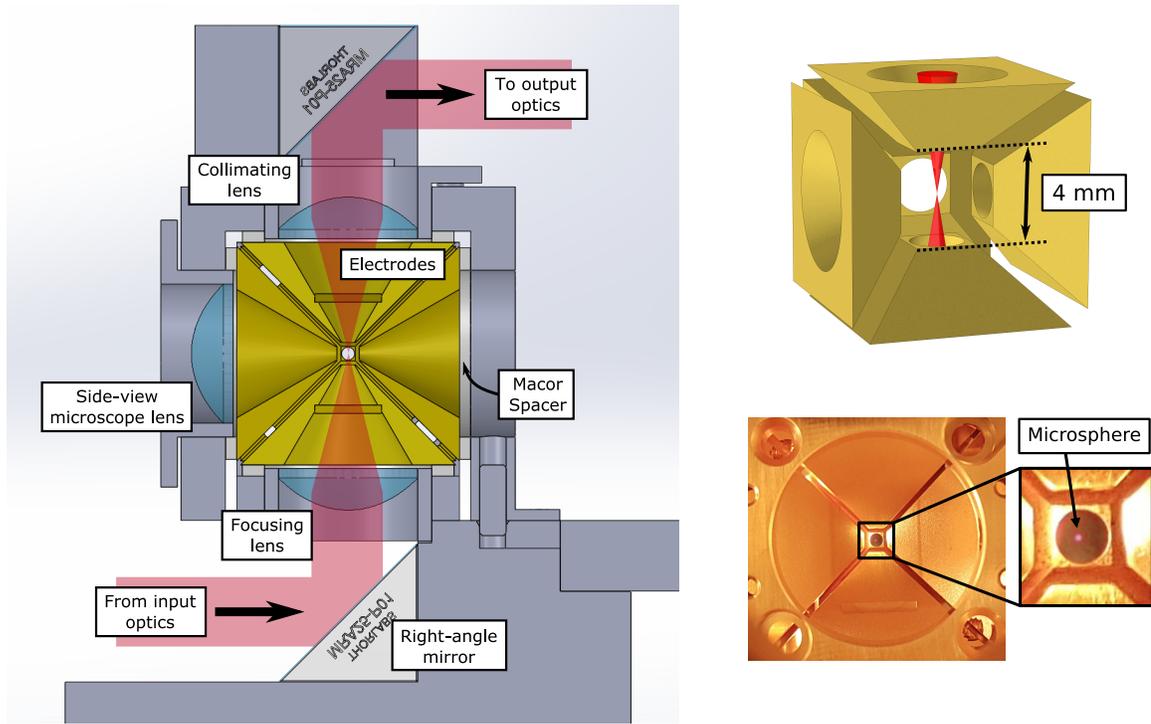


Figure 2.5: (left) A schematic cross-section of the trapping lenses, shielding electrodes, and the housing structure for both. (upper right) The cubical cavity which constitutes the trapping region. (lower right) A photograph through the primary viewing window, where the microsphere is visible due to high-angle scattering of the trapping laser. The image was captured with a cell-phone camera, which retains some sensitivity at 1064 nm. In the schematic views, the foreground electrode has been suppressed, while in the photograph, it was intentionally removed in order to see the trapped microsphere.

in order to eject microspheres toward the optical trap. The quartz pane below the trap is much wider, and serves to protect the lower lens from accumulation of microspheres that fail to be trapped following ejection from the upper pane. Details of the preparation of a trapped microsphere are included in Section 2.6.

2.3.1 Nano-positioning Stage and Mechanical Devices

As stated previously, a significant advantage of the optical levitation apparatus at Stanford is the use of a single-beam for trapping, feedback, and imaging, together with trapping lenses that have a relatively long focal length, 25 mm, compared to typical optical tweezers. These two properties together allow mechanical devices, to be introduced into the immediate vicinity of the focus of the optical trap with minimal disturbance of the optical mode, as there are no additional beams and the diffraction-limited divergence of the trapping beam is only $\sim 6.5^\circ$. These mechanical devices have a number of potential uses: knife edges for beam-profiling to characterize the optical mode of the trap (see Section 2.3.2); near-field electrodes for system calibration and technology demonstrations (see Chapter 3); as well as source masses to excite possible new interactions for beyond the Standard Model physics searches (see Chapters 7 and 8).

In each application, precise positioning of the mechanical devices is critical. This is accomplished by having the devices mounted on a series of stages, two for coarse manual alignment and device insertion, and one for computer-controlled and piezoelectrically-driven fine alignment. An individual device is first mounted to a custom-machined conical snout (gold-plated 7075 aluminum) using a small mechanical spring together with vacuum-compatible colloidal silver paint to ensure a conductive connection, as nearly all of the devices used are metal-coated. The snout itself is mounted on the actuated face of a three-axis, piezoelectrically-driven flexure stage (Newport, NPXYZ100SGV6) which has three orthogonal axes of translation with $80\ \mu\text{m}$ of travel in each direction, and strain gauges for closed-loop control providing 4 nm resolution and 30 nm repeatability. A custom MACOR spacer between the stage and snout ensures electrical isolation, where screws hold the spacer to the stage, while the same colloidal silver paint attaches the snout to the MACOR spacer, with significant clearance between the silver paint and the heads of the screws.

The piezoelectrically-driven flexure stage is mounted to a linear translation stage (Newport, AG-LS25V6) with a custom-machined L-bracket that ensures clearance for the motion of the flexure stage, as well as registered kinematic mounting for consistent alignment. The

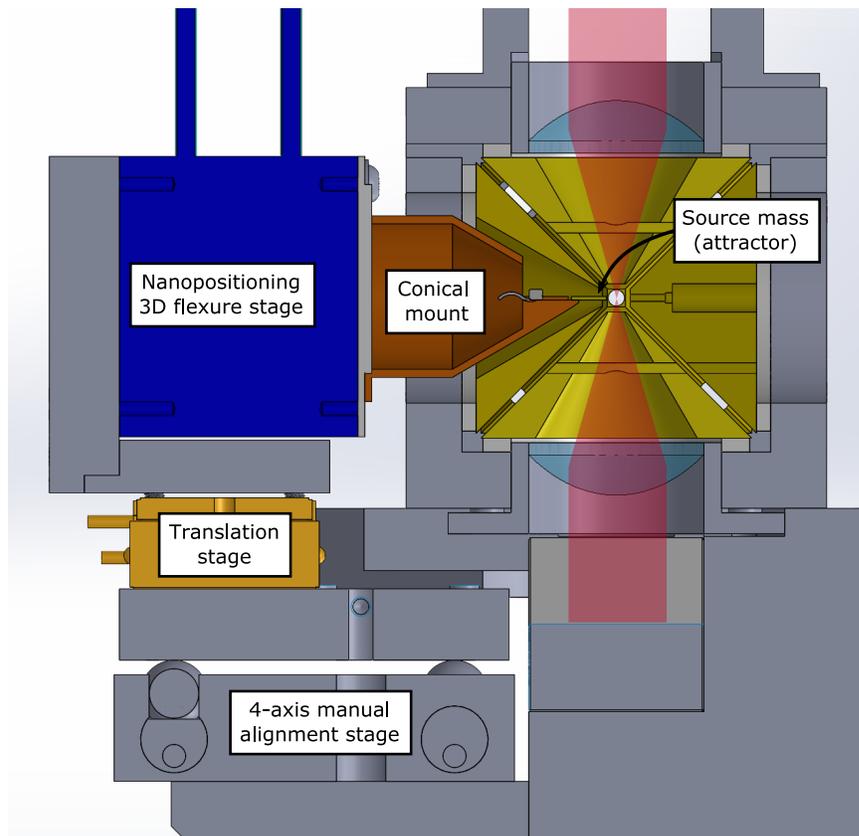


Figure 2.6: A schematic cross-section of the stack of stages, with a typical attractor device mounted (see Chapters 7 and 8, and Appendix G). The electrode cube, including the bore through which the device is inserted into the trap, is visible on the right.

linear translation stage has a full travel of 12 mm so that the mechanical devices can be removed fully from the trapping region when they are not in use. This entire stack sits atop a four-axis tip-tilt stage (Newport, 9071-V) which is used to coarsely align the position of the mechanical device relative to the optical focus, i.e. where a trapped microsphere usually resides. A cross-section of the entire assembly of stages and mounting components is detailed in Figure 2.6.

Generally speaking, the mechanical devices used in this work resemble cantilevers, such as those used for atomic force microscopy. They are fabricated in silicon and have a handle element shaped as a rectangular prism with characteristic dimensions $7\text{ mm} \times 2\text{ mm} \times 0.5\text{ mm}$, off of which is cantilevered a narrow beam-like structure with characteristic dimensions $1000\text{ }\mu\text{m} \times 500\text{ }\mu\text{m} \times 10\text{ }\mu\text{m}$. Where appropriate, the specific geometry of a device will be

indicated, as well as the material composition and any underlying structure or unique coatings. The fabrication of these devices is performed in-house at the Stanford Nanofabrication Facility by the author, together with colleagues. Procedures and recipes for the fabrication process are detailed in Appendix G.

2.3.2 Characterizing the Trap

The spatial mode and divergence of the optical trap are measured with a knife-edge technique, where a silicon cantilever measuring $475\ \mu\text{m} \times 500\ \mu\text{m} \times 10\ \mu\text{m}$ in the x , y , and z directions is scanned across the laser beam at various axial positions above and below the focus. At the same time, the light exiting the chamber is focused down to a single photodiode to monitor the power. The concept of this measurement is illustrated in Figure 2.7. Let $I(x, y, z)$ be the intensity of light in the vicinity of the optical trap, in units of Power/Area² so that integrating over x and y at a particular z yields the total power of the beam \mathcal{P}_0 . By positioning the cantilever appropriately and scanning it across the trap in a single direction, x for example, the photodiode measures the cumulative intensity

$$\mathcal{P}(x, z) = \int_{-\infty}^x dx' \int_{-\infty}^{\infty} dy I(x', y, z), \quad (2.1)$$

which can then be differentiated to recover a projection of the intensity onto the y -axis. By scanning the attractor in the y direction, the same can be done to measure the projection of the intensity function onto the x -axis. The result from the differentiation is sometimes referred to as a beam profile, which can be used to estimate spot size and beam quality. As is discussed in great detail in Reference [44], measuring laser beam size and quality is not only difficult, but the choice of methodology can be contentious. Two distinct characterizations are used, and the results are compared.

The most straightforward technique is to fit the measured beam profile to a Gaussian function. A naïve, unweighted fit will prioritize minimizing residuals near the beam center where the intensity is most significant, yielding one measure of the beam's extent. The two left panels of Figure 2.8 demonstrate a typical beam profile near the focal point for both x and y , each with a fit to a Gaussian intensity of the form $I(x) = I_0 \exp\left(\frac{-2x^2}{w^2}\right)$. The shape of the optical mode near the trap center necessarily has the most impact on the behavior of a trapped microsphere, as the microsphere remains confined to this region. It is clear that while this method accurately describes the observed intensity in the vicinity of the

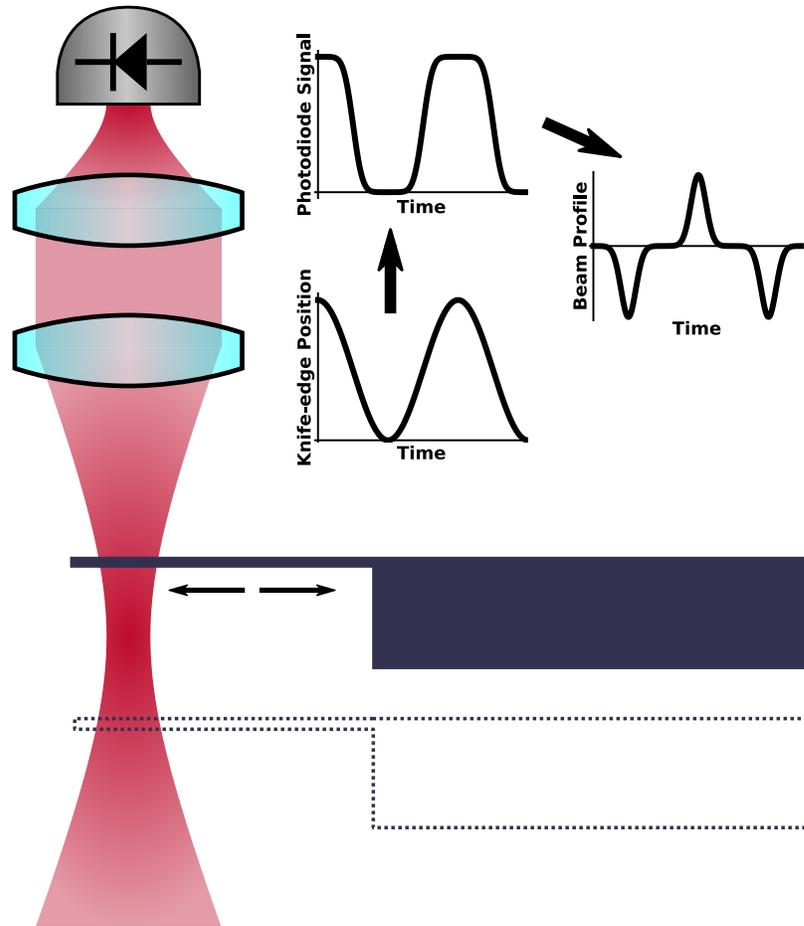


Figure 2.7: A schematic depiction of the knife-edge technique used to characterize the spot size and divergence of the optical trap. A gold-coated silicon cantilever is scanned across the trapping beam at various axial positions above and below the focus, while the output power from the trap is monitored. As the knife-edge covers and uncovers the beam, the photodiode measures the integral given by Equation (2.1), from which the beam profile can be recovered.

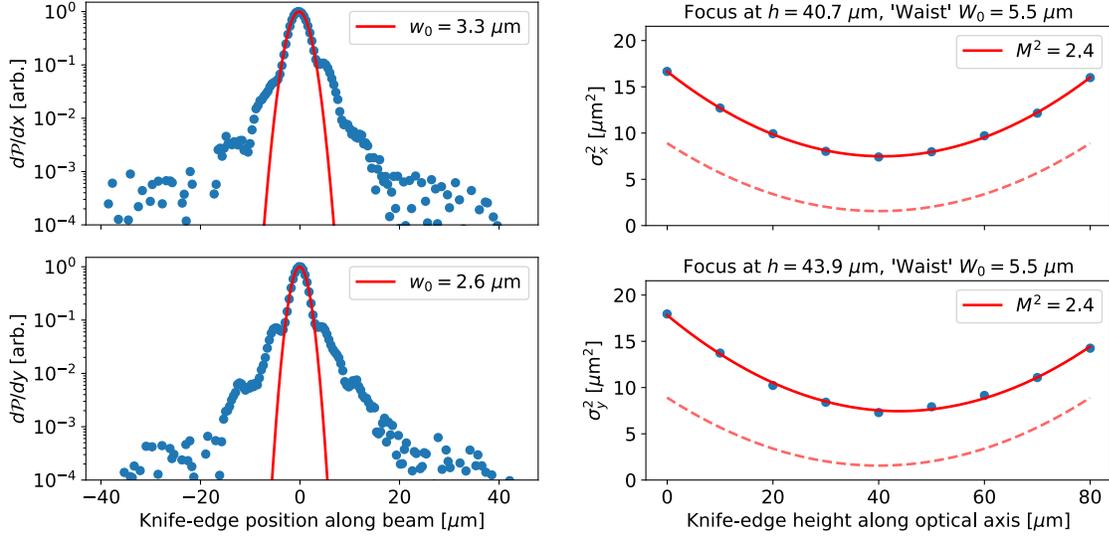


Figure 2.8: (left) Typical results from a beam-profiling measurement for scanning the knife-edge in both x (top) and y (bottom), with the bottom surface of the knife-edge positioned nearest to the focus, at $h = 40 \mu\text{m}$. The data suggest the presence of a halo and/or diffraction rings with \sim few% of the maximum beam intensity. (right) Second moment of the marginalized intensity distribution from beam-profiling as a function of axial position over the full range of the knife-edge's motion, fit to Equation (2.3). The plot includes a dashed line demonstrating some interpretation of the ideal design performance, i.e. the second moment expected if the paraxial approximation was completely valid, the optical system was perfectly aligned, and the beam was actually an ideal Gaussian beam.

beam center, there is significant structure in the tails of the beam, requiring an alternative characterization.

First, let us limit ourselves to beams that satisfy the paraxial approximation, and define the second moment of the beam intensity distribution, $I(x, y, z)$, as a function of axial position,

$$\sigma_x^2(z) = \frac{\int_{-\infty}^{\infty} \int_{-\infty}^{\infty} I(x, y, z) (x - \bar{x})^2 dx dy}{\int_{-\infty}^{\infty} \int_{-\infty}^{\infty} I(x, y, z) dx dy}, \quad (2.2)$$

where, as written, this second moment is associated to the beam width along the x axis, with a similar definition for the beam width along the y axis. For a symmetric Gaussian beam with the conventional definition of the intensity distribution $I(r) = I_0 \exp\left(\frac{-2r^2}{w^2}\right)$, the beam waist w is equivalent to the two times the square-root of the second moment, i.e. $w = 2\sigma$ as defined. Thus, for non-Gaussian beams, two times the second-moment is

often referred to as the spot size, written as W to differentiate it from the usual Gaussian waist. With these definitions, the spot size of any laser beam, W , satisfying the paraxial approximation can be described by the expression,

$$W^2(z) = W_0^2 + M^4 \left(\frac{\lambda}{\pi w_0} \right)^2 (z - z_0)^2, \quad (2.3)$$

where z is the axial coordinate, z_0 is the position of the focus along the optical axis, $W_0 = 2\sigma_0$ is the spot size at the beam's focal point, λ is the laser's wavelength, and M is a unitless parameter related to the beam propagation and divergence. Appealing to Siegman's explanation,

The second-moment-based beam widths W_x and W_y defined above will then propagate with distance in free space exactly like the Gaussian spot size $w(z)$ of an ideal Gaussian beam, except for the insertion of an M^2 multiplication factor in the far-field spreading of the beam [44].

With these tools in hand, the calculation of the second moment can be done for each axial position where a beam-profiling measurement is performed, allowing for an estimation of M^2 , as well as some measure of the beam extent at the focus W_0 , by fitting the observed spot size to Equation (2.3). The results of this procedure for beam-profiling in both x and y are shown in the two right panels of Figure 2.8, together with the expected second-moment for a perfectly aligned, ideal Gaussian beam, i.e. the design performance.

This characterization technique is crucial for optimizing the alignment of the stage on which the input optics reside, discussed in Section 2.2.2. Minute adjustments are made, and then a series of beam profiles at different heights are measured. This procedure is iterated until the optimal beam profile and divergence are achieved. As the input optics stage has 5-axes of control, adjustments can be made to change the angle of the trap without translating the focal point in the horizontal plane, and vice versa, for both horizontal directions. The data shown in Figure 2.8 were collected after just such an alignment procedure.

2.4 Output Optics and Imaging

For the same reasons that the number of free-space optical elements on the input side was minimized, the output free-space optics are similarly frugal. The beam exits the chamber through an AR coated window, mounted at a slight angle (see Section 2.1), immediately

passing through a polarizing beamsplitter cube (Thorlabs, CCM1-PBS25-1064) separating any of the incident trapping light scattered into the polarization orthogonal to the trapping light’s polarization. Imaging the cross-polarized light is critical to understanding the rotational dynamics of an optically trapped microsphere, as any birefringence in the microsphere will modulate the power in the orthogonal polarization at twice the rate of rotation, as demonstrated in Refs. [28, 38] and discussed in Chapter 5.

The output trapping beam is then de-magnified by a factor of three, again with a pair of 25 mm diameter, bi-convex lenses in a refracting telescope configuration: the primary lens having a focal length of $f = 150$ mm and the secondary with $f = 50$ mm. The output from the telescope passes through a 50:50 non-polarizing beamsplitter cube (Thorlabs, CCM1-BS014), where it is superposed with a reference wavefront, generated with a fiber zoom-lens collimator as discussed in Section 2.2.1. The collimated reference beam can be translated and deflected with a 6-axis mount (Thorlabs, K6XS), in order to optimize mode-matching with the trapping beam and thus interference contrast. The resulting superposition is projected onto a quadrant photodiode (QPD, Hamamatsu, S5980), with a custom biasing and readout circuit.

On the input side of the chamber, there is another 50:50 non-polarizing beamsplitter cube which combines the light retroreflected by the microsphere and separated by the Faraday isolator with a second reference wavefront, generated by coupling a fiber optic pigtail to an aspheric lens as discussed in Section 2.2.1. Two mirrors on adjustable mounts allow for a similar degree of mode-matching between the collimated reference beam and the retroreflected light. Their superposition is projected onto a single photodiode, where the phase of the interfering optical fields encodes the axial position of the microsphere, as discussed below.

2.4.1 Reflected Light Interferometry - Implementation

Some amount of the trapping light is retroreflected by the bottom surface of the microsphere, in addition to other high-angle scatter. Assuming the retroreflected light comes from a point source, the first trapping lens (below the sphere) will collimate the reflection, which then propagates backward along the beam line. The Faraday isolator in the input optics separates this back-propagating component, which is then superposed with a reference beam, the “local oscillator” in the usual language of interferometry, and projected onto a single photodiode (Thorlabs, DET100A). As a trapped microsphere moves up and down within the optical

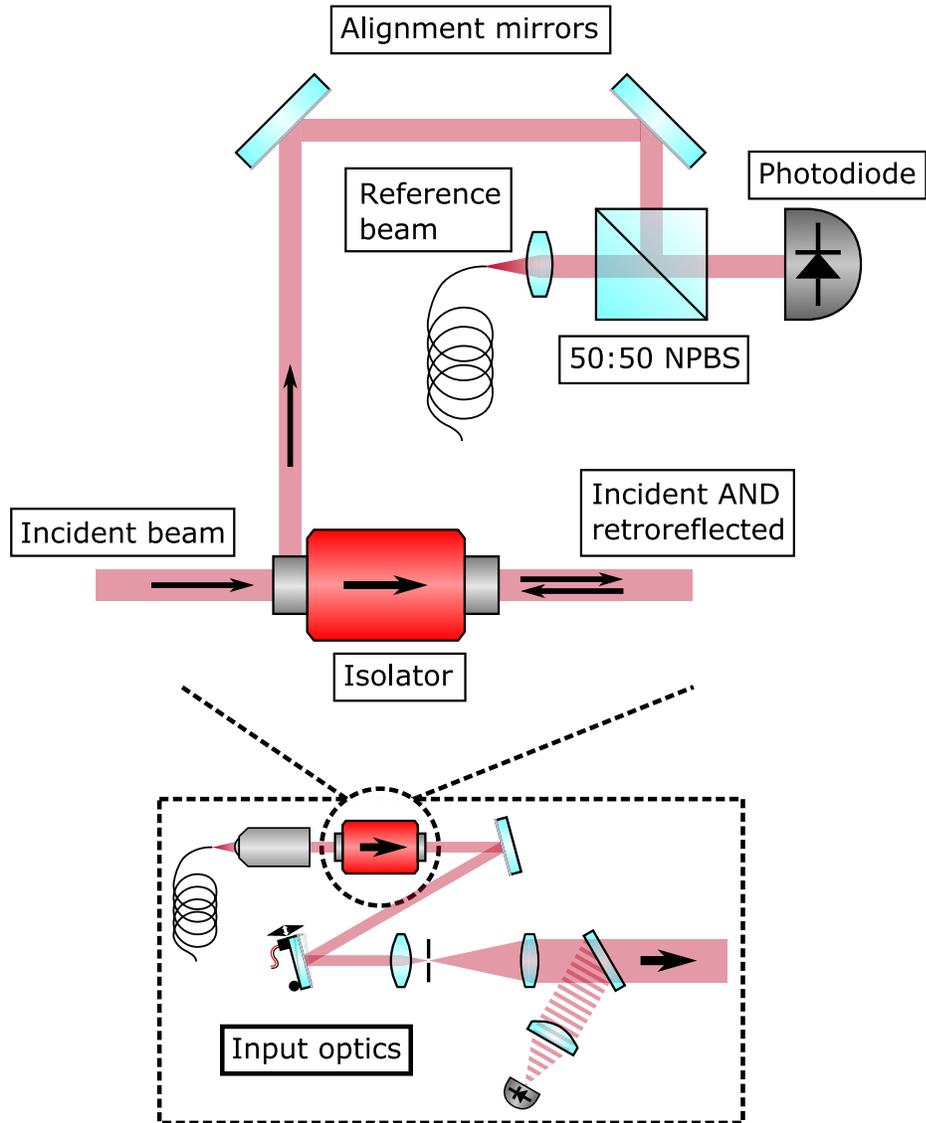


Figure 2.9: (left) A schematic depiction of the optical system used to separate light retroreflected by a trapped microsphere, combine it with a reference beam, and subsequently convert the optical interference to an electronic signal. (right) A photo of the same system, with components annotated.

trap (i.e. along the optical axis), the path length of the retroreflected light changes, which is necessarily encoded in the phase of the retroreflected light, and thus in the phase of the interference signal, as discussed in Chapter 1. A schematic depiction of this optical system, together with an annotated photograph of the same components, is shown in Figure 2.9.

The retroreflected spot is nominally circular, but passes through small apertures in the Faraday isolator, which would tend to produce Airy patterns or possibly asymmetric clipping. The collimation of the reference beam for the reflected light interferometry discussed in Section 2.2.1 was chosen specifically to best match the observed spot size of the light retroreflected from typical $\sim 5\ \mu\text{m}$ -diameter silica microspheres, after having passed through the isolator and associated apertures.

In contrast to the ideal discussion and analytic treatment in Chapter 1, the wavefronts of the reflected light and the first reference beam are neither perfectly flat, nor symmetric in shape or overall intensity (although the reference beam is much closer to an ideal Gaussian). The interference contrast is optimized empirically, by adjusting reference beam position, orientation, collimation, and power. Adjusting the position and orientation ensures the superposed optical fields are propagating in the same direction with maximal overlap of the spatial modes. This is accomplished by deflecting the reflected beam with a pair of protected-silver mirrors on adjustable mounts, as well as rotating the fiber-optical pigtail (a polarization-maintaining fiber) to match the reference beam polarization to that of the reflected beam. Adjusting the collimation allows the reference beam to match any curvature of the reflected trapping beam, achieved simply by changing the separation between the reference beam fiber-optic pigtail and the collimating lens. As a final step, the optical power of the reference beam is increased so that the photodiode is operating at approximately half of its saturation limit.

A simple AC-coupled transimpedance amplifier (See Appendix E) converts the oscillating photocurrent generated from the optical interference to an oscillating voltage, which is then sent directly to an ADC for simultaneous digitization and digital demodulation. The phase of the demodulated signal yields the z position of the microsphere, where the phase is estimated in real-time, as discussed in Appendix D. The real-time estimate of this phase can be used to generate active feedback in order to stabilize the vertical (axial) motion of the microsphere by modulating the beam power (see Section 2.7). In contrast, the amplitude of the demodulated signal is used mostly as an auxiliary measurement to ensure alignment and optimal mode-matching. Although a higher-order effect, a large signal amplitude relative to

any noise ensures a robust determination of the signal phase.

2.4.2 Transmitted Light Interferometry - Implementation

As discussed previously, the light that is transmitted through a trapped microsphere and collimated by the upper trap lens exits the chamber through an AR-coated window, is demagnified to reduce the spot size, and subsequently combined with a reference beam in a 50:50 non-polarizing beamsplitter cube. The resulting superposition impinges on a QPD, capable of imaging the beam displacements along two orthogonal axes. The reference beam is kept stationary, while the transmitted trapping beam is displaced due to the motion of the microsphere. As the trapping beam moves about, the mode-matching between the two interfering beams changes, and thus the amplitude of the resulting interference signal on each quadrant of the QPD changes, which is the basis for the measurement of the horizontal degrees of freedom, as discussed in Chapter 1. A schematic depiction of this optical system, together with an annotated photograph of the same components, is shown in Figure 2.10.

Much like for the wavefronts discussed in the previous section, the wavefronts of the transmitted light and the second reference beam are neither flat, nor symmetric in shape or intensity. As before, the interference contrast is optimized empirically, with nominally the same guiding principles. The zoom-lens collimator and six-axis alignment stage for the second reference beam, discussed in Section 2.2.1, make this task relatively straightforward. The ideal alignment has equal power (from both beams) impinging on each quadrant of the QPD.

The QPD used has a “common-cathode” configuration, so that one can apply the same positive bias voltage to all quadrants, whilst measuring the resulting photocurrent sourced by each of the four anodes separately. Each quadrant is AC-coupled to its own transimpedance amplifier, with a second adjustable gain stage following, while the DC photocurrent is sunk to ground through an inductor in a “tee” configuration with the coupling capacitor. The entire biasing and readout circuit is of custom origin and detailed in Appendix E.

Following amplification, the oscillating signals from all four quadrants are sent to an ADC for simultaneous digitization and demodulation. For the transmitted light interferometry, the distribution of light across the QPD is the quantity of interest, which is obtained by comparing the amplitudes of the interference photocurrent in each of the quadrants, after the signals have been demodulated. Using standard Cartesian coordinate quadrants (i.e. ‘I’ for $x, y > 0$, ‘II’ for $x < 0 \wedge y > 0$, ‘III’ for $x, y < 0$, and ‘IV’ for $x > 0 \wedge y < 0$), an estimate

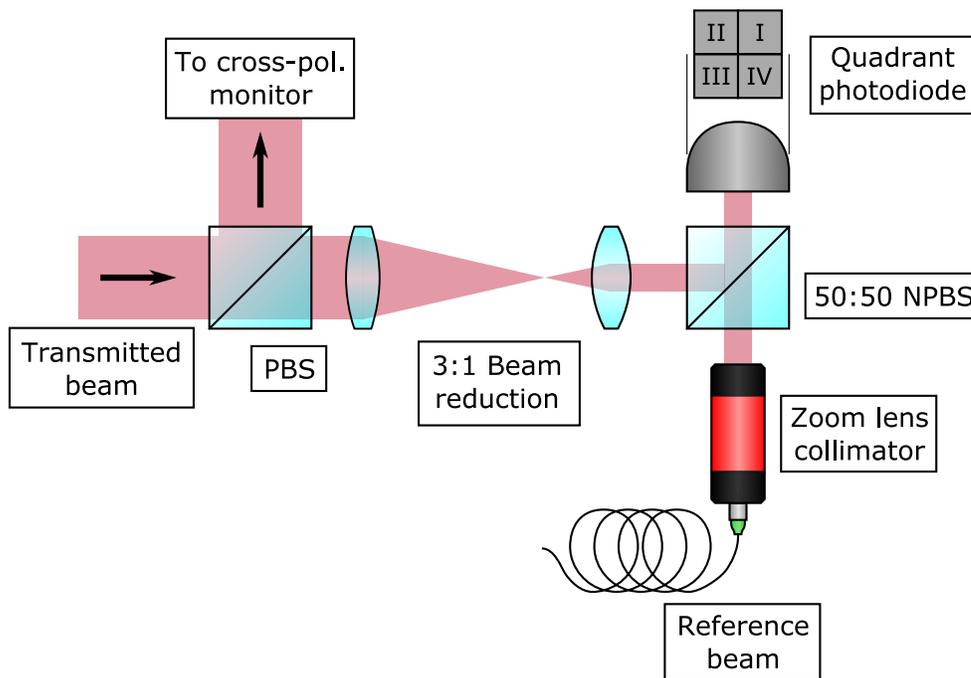


Figure 2.10: (left) A schematic depiction of the output optical system which handles the light transmitted through a trapped microsphere, de-magnifying the beam to reduce spot size, superposing with a reference beam, and finally, converting the optical interference to an electronic signal. (right) A photo of the same system, with components annotated.

of the beam displacement is given by

$$x = \frac{A_I + A_{IV} - A_{II} - A_{III}}{\sum_i A_i} \quad \text{and} \quad y = \frac{A_I + A_{II} - A_{III} - A_{IV}}{\sum_i A_i}, \quad (2.4)$$

with A_α being the demodulated amplitude of the digitized signal in each of the quadrants α . The horizontal and vertical differences are normalized by the sum so that power fluctuations are suppressed in the estimates of beam displacement. These estimates are calculated in real-time within the FPGA that performs the demodulation, as is discussed in Appendix D, so that they can be used to generate active feedback in order to stabilize the horizontal motion of the microsphere by deflecting the beam with the piezoelectrically-driven deflection mirror discussed in Section 2.2.2.

The phases of the oscillating signals from each quadrant are not explicitly used in estimating x or y , or in generating the horizontal feedback, however, they are extremely useful during the optimization of interference contrast and alignment. For two ideal Gaussian beams interfering in the manner described, the phase of the interference in each quadrant should be the same. Angular misalignments between the transmitted trapping beam and the reference beam generate phase shifts between quadrants. Thus, “optimal” alignment for this interference measurement is achieved by closely matching both the amplitude and phase of the interference signals in all of the quadrants, via adjustments of the reference beam, while a microsphere is within the trap. In fact, even with active feedback turned on and the system at high vacuum, adiabatic adjustments to the optical alignment are possible without jeopardizing the stability of a trapped microsphere.

2.4.3 Stray Light Rejection

Immunity to extraneous sources of light is critical for short-range force sensing where mechanical devices that can scatter light are brought in proximity to a trapped microsphere. Heterodyne systems, such as the one described here, provide substantial rejection of light propagating along a path that is different from the desired optical axis. Consider a point source of scattered light that is displaced orthogonal to the optical axis, but in the focal plane of the trap, as this will be the nominal position of mechanical devices during force sensing measurements. The collimated wavefront of scattered light from this displaced source will arrive at the QPD at an angle relative to the wavefront propagating along the main optical axis, and thus relative to the reference wavefront. The relative angle between two

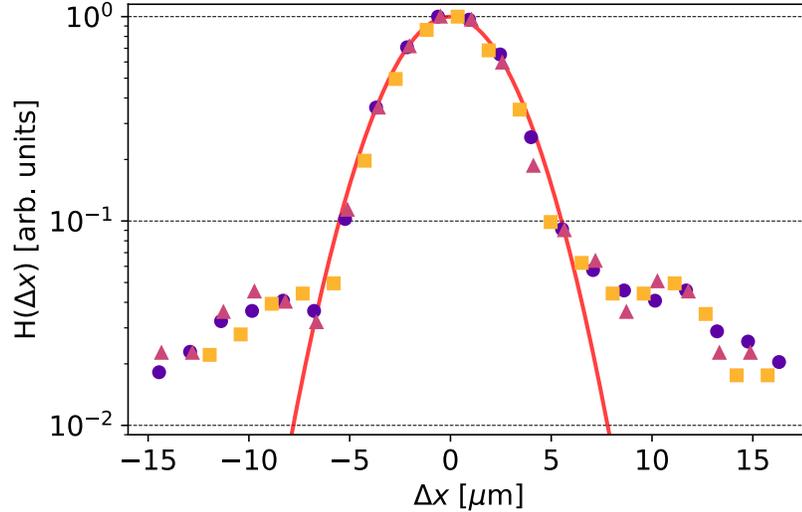


Figure 2.11: Interference contrast vs radial position in the focal plane of the trap, for three distinct sets of measurements, taken for consistency, and shown with differing marker shapes. The solid curve represents the prediction of Equation. (2.5) calculated with measured beam waists and known focal lengths. Data is normalized to a maximum of one and centered. The width of the predicted profile is not fit to data. The non-Gaussian tails in the data are likely the result of the halo in the trapping beam, as seen in Figure 2.8.

wavefronts causes them to interfere both constructively and destructively simultaneously, reducing their combined contribution to the total interference signal and rejecting angled wavefronts. Then, for a detector positioned in the Fourier plane of the trap, angular rejection corresponds to displacement rejection in the focal plane of the trap.

The angular rejection of the heterodyne system described is estimated by considering the interference of two Gaussian beams at their focus and separated by an angle α between their wavefronts. The profile of angular rejection $H(\alpha)$ can be computed from the normalized integral $\iint |\vec{E}_1 + \vec{E}_2|^2 dA$ with \vec{E}_1 and \vec{E}_2 the electric fields associated with the appropriately tilted Gaussian beams. The integral is decidedly difficult, and is thus expanded in a power series assuming α is small. The resulting profile of scattered light rejection is given by,

$$H(\Delta x) \simeq \exp \left[\frac{-(2\pi/\lambda)^2 w_1^2 w_s^2}{4(w_1^2 + w_s^2)} \left(\frac{\Delta x}{d} \right)^2 \right], \quad (2.5)$$

where w_s is the waist associated with the source of scattered light imaged onto the detector,

and $\Delta x = d\alpha$. This result can be compared with data collected by using the trapping beam as a test source of light and angling the reference beam, with a single-channel photodiode placed in the detector focal plane, in the place of the QPD. The response, calibrated in terms of position at radial distances Δx from the center of the trap, is shown in Figure 2.11, along with the prediction of Equation (2.5). The tails of the distribution present in the data, but not the calculation, are likely due to interference of the reference beam with the halo of the trapping beam, shown in Figure 2.8.

2.5 Auxiliary Optics

In addition to monitoring the center-of-mass motion of a trapped microsphere, it is useful to characterize the microsphere's rotation, as well as the region of space surrounding the microsphere's nominal position within the trap, as will become apparent in later chapters. The latter can be accomplished in a straightforward manner with simple camera and lens systems, both to monitor light scattered by the trapped microsphere, as well as light generated by the reflection of illumination beams from objects in close proximity to the microsphere. The former, however, does not have a generic solution, and instead depends on the specifics of the trapped particle and the trapping light.

2.5.1 Cross-polarized Light Monitor

Silica microsphere's grown via the Stöber are nominally amorphous, although there exist crystalline grains whose size depends significantly on the growth conditions. As a result of the crystalline grains and non-spherical geometry, silica microsphere's exhibit varying levels of residual birefringence (See Chapter 5), as well as permanent electric dipole moments (See Chapters 3, 5, and 7), both manifestations of an anisotropic polarizability tensor.

Now, consider such a microsphere within an optical trap whose light has a fixed linear polarization. As the microsphere rotates within the optical trap, the residual birefringence couples some of the trapping light into the polarization orthogonal to the trapping beam, which will be referred to as cross-polarized light. The power of the cross-polarized light modulates as the orientation of the microsphere changes, which can be used to characterize the orientation and rotation speed. This will be discussed in significant detail in Chapter 5, but a rudimentary explanation is necessary to understand the purpose of the following components.

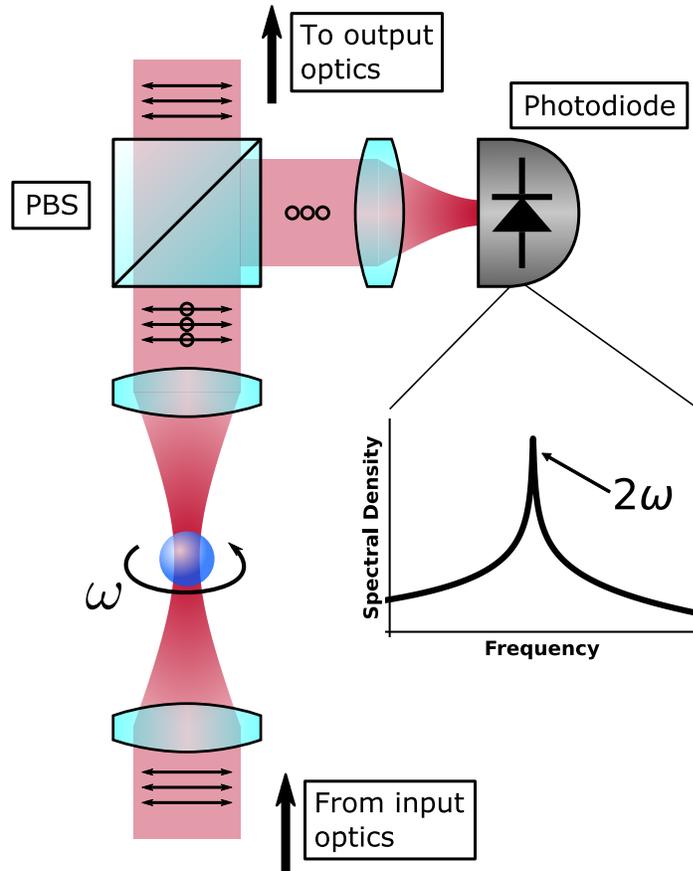


Figure 2.12: A schematic depiction of the output optical system to separate and characterize the cross-polarized light generated by rotations of a trapped microsphere. Linearly polarized light enters the chamber and interacts with a trapped microsphere. Rotations of the microsphere in the horizontal plane couple some of the linearly polarized light into the orthogonal polarization, which can then be separated by a polarizing beamsplitter (PBS)

To separate the cross-polarized light, there is a polarizing beamsplitter cube (Thorlabs, CCM1-PBS25-1064) immediately following the AR-coated window through which the transmitted trapping light exits the chamber. The cross-polarized light from the trapped microsphere is then directed toward a single photodiode (Thorlabs, DET10A) using a protected silver mirror, with a converging lens to ensure all of the light is incident within the active area of the photodiode. This optical system is detailed in Figure 2.12.

A transimpedance amplifier (Thorlabs, AMP120) converts the photocurrent to a voltage, which is further amplified and filtered by low-noise voltage preamplifier (Stanford Research Systems, SRS560). The SRS560 has two outputs, one of which is sent directly to an ADC as part of the main data acquisition system, while the other is sent to a secondary control and acquisition system dedicated entirely to driving and characterizing rotations of a trapped microsphere. The details of both electronics systems are detail in Appendix D.

Despite the use of polarization-maintaining fiber-optic cables, and various polarizing optics, there is always some excess, but roughly constant, amount of cross-polarized light, even in the absence of a trapped microsphere. Generally speaking, this is attributed to imperfect reflections and scattering, especially from mirrors used to align some of the optical elements, as well as finite polarization extinction ratios in practical polarizing optics, roughly 20 dB to 30 dB without cavity-based mode cleaners. Thus, only power modulations of the cross-polarized light are used to infer rotational motion of the microsphere, and there is no information at DC.

2.5.2 Side-View Microscope

In addition to the high-bandwidth, photodiode-based imaging systems dedicated to characterizing motions of the microsphere, there are two microscopes used to view the region of space surrounding the optical trap. Both make use of CMOS cameras and relatively simple lens systems.

The first of the two, which will be referred to as the side-view microscope, is dedicated to light scattered at high-angle from the microsphere itself. The silica microspheres commonly used in this apparatus are amorphous and not AR-coated, so there is significant reflection and scattering. The scattered light can be approximated as coming from a point-source, due to the microsphere's dimensions, with some angular distribution about the optical axis. Some of the light that is scattered horizontally is collected and converged by a 25 mm focal length lens mounted in the side of the electrode cube nearest the input flange, with optical

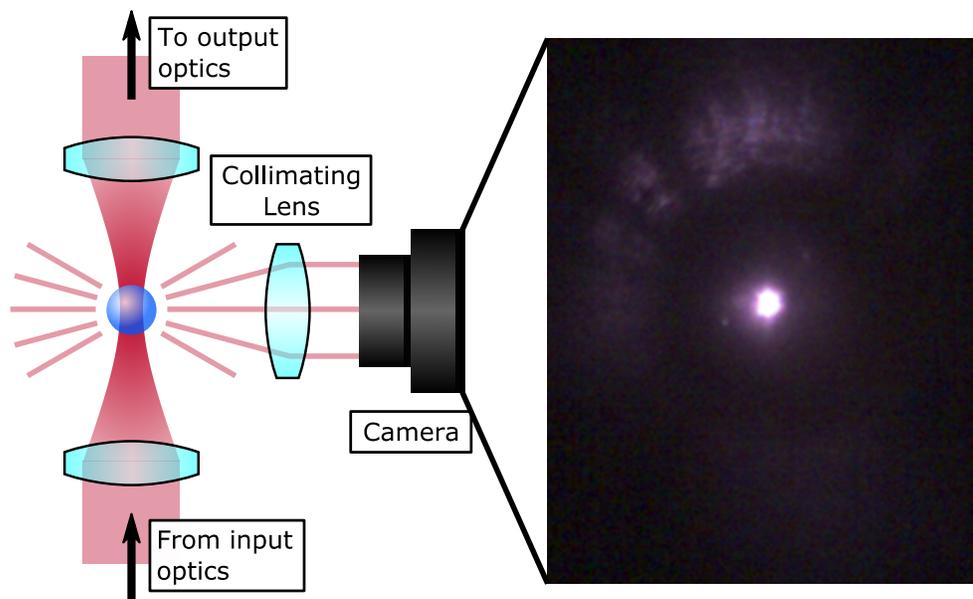


Figure 2.13: (left) A schematic depiction of the side-view microscope dedicated to light scattered at high-angle from a trapped microsphere. A single lens collimates the scattered light which is collected in a CMOS camera. (right) An example image captured with this microscope while a microsphere is within the trap. The circular feature visible is the opening of one of the truncated pyramidal electrodes.

access provided by one of the electrode pyramids and its conical bore. This light exits the chamber through a secondary AR-coated window mounted just above the main input window for the trapping light, and impinges directly on the surface of a CMOS camera (Edmund Optics, EO-0513).

A schematic depiction of the side-view microscope with an example image is shown in Figure 2.13. In the example image, the spot corresponding to the microsphere is saturated, and appears much larger than the microsphere itself. This is due, in part, to the use of a only a single converging lens which itself is designed for 633 nm light², so that the image of the microsphere is not properly focused, and the expected size of the image is small relative to the active area of the camera.

The primary use of the side-view microscope is to observe whether a trapped microsphere is present or not, so the clarity and precision of the image are unimportant. If a microsphere

²A previous iteration of the apparatus without interferometry made use of a 633 nm imaging beam incident from a horizontal direction, so that vertical motions of the microsphere could be observed. The focusing lens for this previous imaging system was left in place, as a properly focused image was unnecessary.

has significant defects or is otherwise abnormal (e.g. two fused into a dumbbell, a cleaved portion of a sphere), it will often behave strangely within the trap, oscillating rapidly up and down, or flashing a large amount of light into the side-view microscope. Thus the microscope is also used as a qualitative filter for well-behaved microspheres. During preparation of a trapped microsphere (discussed in Section 2.6), this simple microscope can also be used to assess, roughly, the axial position of the microsphere.

2.5.3 Top-View Microscope

The second microscope is a reflected light microscope dedicated to characterizing the position of mechanical devices within close proximity to the microsphere. One such device was the silicon cantilever discussed in Section 2.3.2 and used to profile the trapping beam with a knife-edge technique. The microscope itself consists primarily of an illumination beam, a single additional lens, and a CMOS camera (Allied Vision, Mako U-29B). The illumination beam is generated with a 1064 nm, fiber-coupled diode laser (Axcel Photonics, M9-A64-0300), which is projected to free-space with an FC/APC bulkhead, and collimated with a 25 mm focal length aspheric lens housed within a lens tube mount that allows adjustment of the lens position relative to the fiber tip. The large linewidth inherent to a diode laser, ~ 1 nm, ensures there is negligible optical interference with the trapping or reference beams (i.e. the beam is incoherent).

The illumination light is injected into the system first via a 50:50 non-polarizing beam-splitter cube (Thorlabs, CCM1-BS014), which eventually separates some of the reflected light for imaging with a camera. Following the first cube, there is a second polarizing beam-splitter cube which separates the cross-polarized light for the auxiliary monitor discussed in the previous section. Although the illumination light is initially unpolarized, the second beamsplitter cube injects only the polarization orthogonal to the trapping beam, further preventing interference between the microscope illumination beam and the trapping beam. The injected light is not well collimated by design (de-focus induced via adjustments of the collimating lens), so that after passing through the upper trap lens, the illumination light has a large spatial mode at the axial position of the trap, ensuring objects in the vicinity of the trap can be seen. Reflected illumination light propagates backward along its incident path (i.e. the output path for the transmitted trapping light), eventually passing through the initial 50:50 non-polarizing beamsplitter cube and impinging on a 100 mm focal length lens, which projects the image onto the active area of the camera.

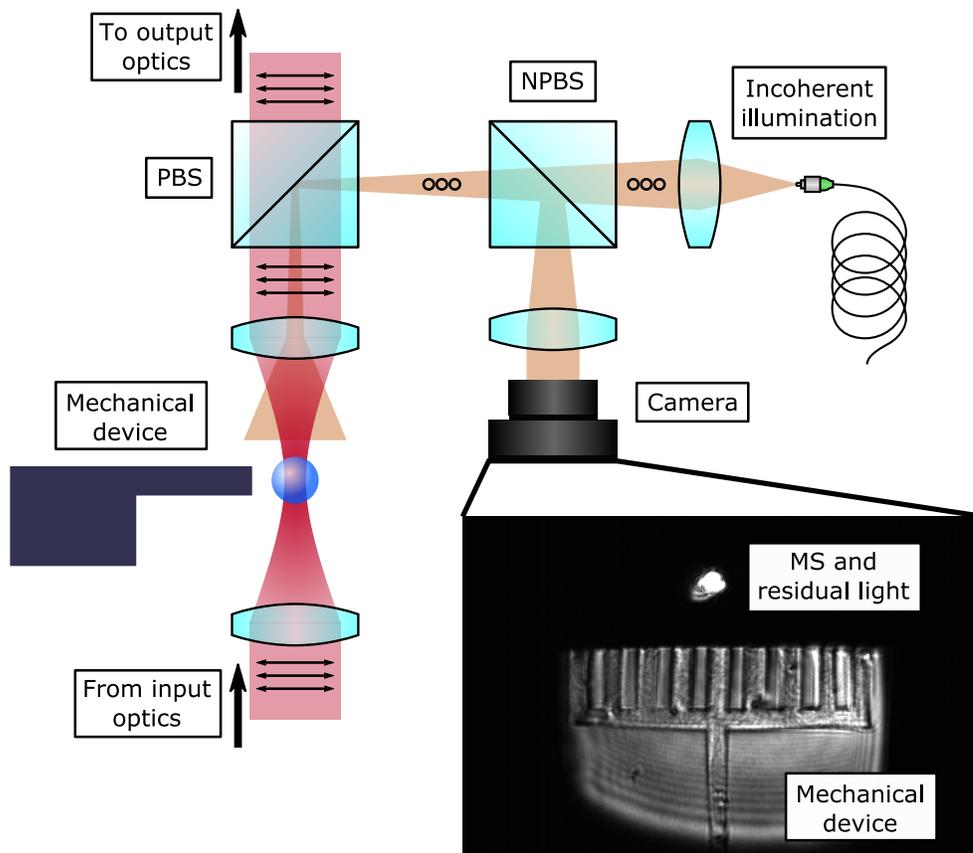


Figure 2.14: A schematic depiction of the top-view microscope dedicated to light transmitted through a trapped microsphere, as well as the imaging of devices within the trapping region. (inset) An example image captured with this microscope while a microsphere is within the trap and a mechanical device is in proximity. Due to imperfect polarization, some amount of the trapping beam is visible and indicative of the position of the microsphere. The image is significantly aberrated by the aspheric trapping lens through which the light must pass.

A schematic depiction of the top-view microscope and some of the surrounding optics, together with an example image is shown in Figure 2.14. Within the example image, a trapped microsphere is joined by another mechanical device in close proximity, where the light reflected from the device is visible together with the cross-polarized light from the trapping beam. The device shown in the image is known to have sharp rectangular cross-sections and internal features (See Chapters 3, 7, 8 and Appendix G), so the significant aberrations induced by the aspheric trap lens are clearly visible. However, these aberrations can be calibrated to an extent, so that images captured by the top-view microscope can be used to assess the relative position between the trap and any mechanical device.

The mechanical devices brought in close proximity to the microsphere have their position precisely controlled by the nanopositioning stage discussed in Section 2.3.1. By moving a device over a grid of known positions and saving an image of the device at each grid position as a calibration, image correlation and registration algorithms can be used to determine an arbitrary device position within the grid, based on the microscope image at that arbitrary position. In this way, it is possible to calibrate for the significant aberrations, without actually correcting them optically. This technique will be discussed in more detail in Chapter 3.

The next generation of the apparatus, which has been built and is under operation [37] at the time of writing, has a pair of much more sophisticated microscopes in place of the side- and top-view microscopes discussed here, with significant magnification that can generate proper images of trapped microspheres.

2.6 Microsphere Preparation

Obtaining a trapped microsphere that is useful for precision measurements is a long and potentially arduous process, with a number of low-yield steps, as well as necessary actions that can prove dangerous for a microsphere's livelihood. While the process detailed here pertains purely to the heuristic and technical aspects of preparing a microsphere, the author cannot discount the importance of repeated honorary rituals such as dancing and chanting, which appease the deities in charge of microsphere trapping (colloquially referred to as "the bead gods"). The author recommends establishing a good rapport, achieved by consistent observance of the established group rituals.

There are two generally accepted methods to get a microsphere within an optical trap,

although a variety of others exist. The most common methods make use of the law of large numbers, wherein a veritable cloud of microspheres are thrown into the vicinity of the trap, with the hope that one member of the cloud will be incident directly into the focus of the trapping beam, where optical confinement can keep the microsphere in place. In one method, spheres are suspended in a vacuum-compatible solvent like isopropanol or ethanol, which is then nebulized inside of the experimental chamber. Eventually, a droplet will encounter the optical trap, vaporizing the solvent and retaining the microsphere within, assuming the concentration of the suspension has been tuned such that there are 0 – 1 microspheres per droplet. The obvious downside of this method is the continual and intentional coating of the chamber interior with the solvent used to suspend the microspheres, especially problematic if there are non-trivial impurities in the solvent.

The second common method involves mechanically “throwing” a cloud of microspheres at the trap. This is usually accomplished by piezoelectrically exciting a surface that has been covered with microspheres, driving mechanical modes of the surface, which then ejects some of the spheres away from the surface. This is the preferred method for this apparatus, and will be described below.

2.6.1 The “Bead Dropper”

Following manufacture and sale, the microspheres begin their journey as a dry powder of roughly $N = 10^{10}$ spheres, the bulk of which remains in storage. A small portion of these spheres ($\sim 1 - 2\%$ of the total) are removed from the bulk supply and stored on a glass microscope slide. The microscope slide is placed on a hotplate and heated to 200°C in order to dry the spheres and remove any residual water or other solvents. A long and narrow quartz pane with dimensions $26\text{ mm} \times 2.5\text{ mm} \times 0.5\text{ mm}$, referred to as the “bead dropper”, is then pressed into the pile of dried microspheres on the slide, such that the microspheres remain attached to the bottom surface of the bead dropper via Van der Waals forces. While the bead dropper and the slide are pressed together, the dropper is moved in a circular motion to distribute the microspheres as even as possible on the surface, ideally creating a monolayer. The small amount of microspheres placed on the microscope slide can last many months and are repeatedly heated to ensure the microspheres are dry before coating the bead dropper.

Although not obvious, this “bead smearing” step is quite important to successful trapping. If there are too few microspheres, the low-yield steps later in the process can present

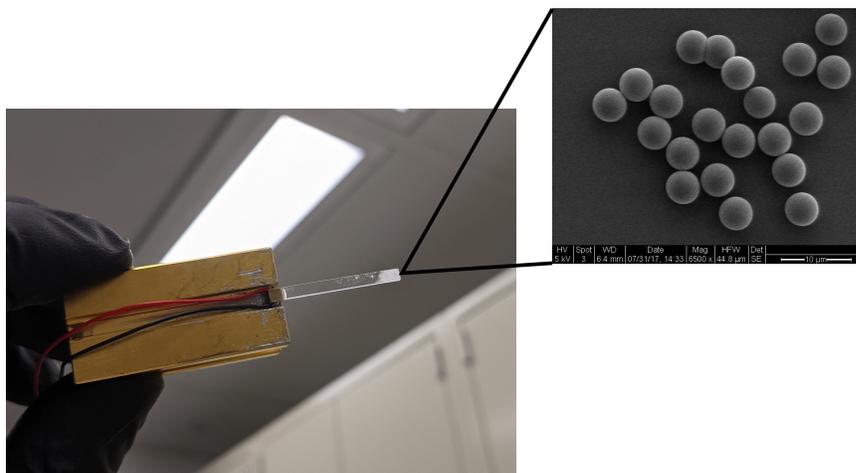


Figure 2.15: Photograph of the bead dropper, the piezoelectric actuator, and the supporting structure. The bead dropper has been prepared as described in the text, and a layer of microspheres is visible near the far end. An SEM image of the same lot of microspheres is included.

a significant challenge, while if there are too many microspheres (i.e. a thick layer, many spheres deep), it's difficult to reliably trap a single sphere, and the high density can interfere with the critical steps. For the bead dropper and piezoelectric actuator used here, a 'good' bead smearing is determined by holding the dropper up to the room lights, and looking for an optical density of about 0.75 to 1.5 in the covered region.³ The bead dropper and a stereotypical smearing are showing in Figure 2.15.

The end of the bead dropper opposite from the layer of microspheres is attached to the actuated surface of a single piezoelectric chip (Thorlabs, PA4DGW), which is used to excite multiple mechanical modes within the bead dropper, and thus eject microspheres from the surface to which they are adhered. In particular, the piezoelectric actuator is driven with a fixed DC bias, usually \sim half of the maximally allowed voltage to ensure there is never a negative driving voltage, plus a chirp signal from a direct digital synthesis waveform generator (Hewlett Packard, 3325b). The chirp passes through an RF amplifier with a gain of 30 dB, as well as the requisite slew rate in order to drive the capacitive load that is the piezoelectric chip. A typical chirp lasts 1 s, sweeping linearly across the frequency range 150 kHz to 300 kHz so that the drive has a relatively rapid dwell time on each mechanical mode of the dropper. The amplitude of the chirp is determined empirically for a particular

³This optical density has never been measured and is an off-the-cuff estimate from experience.

type and/or size of microsphere. For example, larger silica microspheres ($r_{\text{MS}} \gtrsim 3.5 \mu\text{m}$) appear to require less drive, as compared to their smaller counterparts, in order to be detached from the same quartz bead dropper.

This method of delivery is discussed in great technical detail in the PhD thesis of Tongcang Li [31] for the interested reader. For our purposes, the description here is both complete and reproducible. Although only conjecture, this method seems robust against changes of the bead dropper geometry or composition, as well as the use of different piezoelectric actuators and their accompanying fixtures, assuming of course the experimenter takes care to drive the piezoelectric actuator at frequencies appropriate for the mechanical modes of the bead dropper used, often informed by finite element analysis.

In this apparatus, the piezoelectric chip itself is mounted to a custom metal fixture on a linear translation stage (Newport, AG-LS25V6). Part of the fixture also supports a wider quartz pane, referred to as the “bead catcher”, which sits directly below the bead dropper and extends past it. The bead dropper and bead catcher are simultaneously inserted above and below the trap, respectively, making use of the linear translation stage and entering through dedicated slots in the electrode structure, as detailed in Figure 2.16. Once in place, the piezoelectric actuator can be driven as described above in order to eject microspheres from the bead dropper, toward the optical trap.

2.6.2 Trapping

In order to trap a microsphere using the delivery mechanism described, both a high laser power and some ambient residual gas within the chamber are required. The high laser power ensures a deep potential well for the microsphere within the trap, which can encourage radial confinement of the microspheres as they fall, as well as providing sufficient radiation pressure so that a microsphere can be trapped over a wide range of axial positions above the focus, where the irradiance is greatest. Recall from Chapter 1 that for a vertically-oriented optical trap with low numerical aperture and stabilized by gravity, the equilibrium position of a microsphere is determined by the radiation pressure, which scales with the optical power. Thus, higher optical power means that the equilibrium position resides above the focus and has a much larger lateral extent, both of which are due to the divergence of the beam. Let \mathcal{P}_0 be the power of the beam for which the microsphere’s equilibrium is at (or slightly above) the focus of the trap. A rule-of-thumb that has proved reliable for silica microspheres in this type of configurations is that one needs approximately $10 \times \mathcal{P}_0$ in order to initially catch the

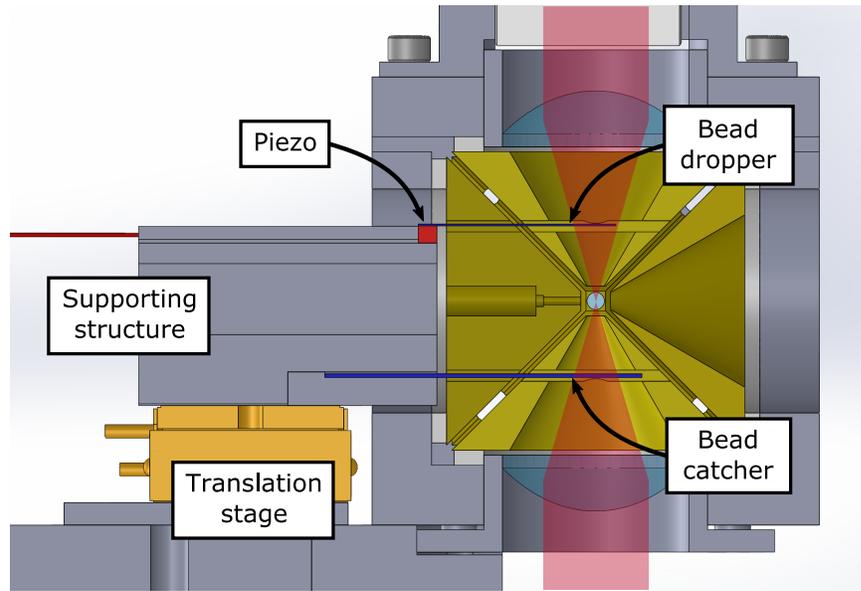


Figure 2.16: Schematic cross-section of the electrode cube and trapping lenses, highlighting the bead dropper, bead catcher, and the slots through which they enter the beam path above and below the optical trap. The components are shown near to their full extension, and can be retracted by 12 mm using the translation stage shown.

microspheres as they fall.

The ambient residual gas in the chamber helps to partially arrest the motion of the microspheres ejected from the dropper, as well as provides mechanical damping for the microsphere's center of mass motion within the trap thus encouraging confinement. It has been found, for silica microspheres in the size range $r_{\text{MS}} = 2.3 \mu\text{m}$ to $5 \mu\text{m}$, that filling the chamber with N_2 at pressures in the range 1 mbar to 10 mbar provides both sufficient drag to slow the microsphere's fall, as well as acceptable damping of the center of mass motion. Higher pressures are possible, however they vastly increase the time required to pump the chamber down to a high-vacuum environment, as the action of pumping creates wind which can violently remove a microsphere from the optical trap if the pumping is not exquisitely controlled with the system described in Section 2.1.1. This wind effect is observed at higher pressures where viscous flow dominates. For this apparatus, the transition to free molecular flow of N_2 in the vicinity of the optical trap is around a pressure of ~ 0.01 mbar.

With a coated bead dropper and the chamber fully closed, the procedure to trap a microsphere proceeds as follows:

1. Increase optical power and ambient pressure.
2. Insert bead dropper and bead catcher above and below the trap, respectively.
3. Drive piezoelectric actuator with a 1 s chirp as described.
4. Observe side-view microscope to check for a trapped microsphere.
5. If a microsphere is present, proceed. Otherwise, repeat steps 3 and 4, slightly increasing the drive amplitude at each iteration.
6. Retract bead dropper and catcher from the trap.
7. Incrementally lower optical power to bring the microsphere closer to the focus.

The drive amplitude for the piezoelectric actuator is increased at each iteration of driving, presumably because there is some distribution of contact forces for the microspheres attached to the dropper. Indeed, if the drive amplitude is not increased, initially some microspheres will fall, but they quickly stop. Often, Steps 3 and 4 have to be repeated 10 to 20 times before a particle gets trapped, and each repetition removes more microspheres, making this a relatively low-yield process. Additionally, trapped particles can be lost during step 6, due to mechanical vibrations associated to the linear stage which inserts and retracts the dropper/catcher, as well as the scattering of the trapping beam that occurs as the edge of the bead catcher passes through the beam and perturbs the optical mode. The process is quite reliable despite the low-yield, and has been tuned only tiny amounts during the course of this work (roughly ~ 6 years).

Although reliable, a non-negligible portion of the objects that get trapped are “weird”. One type of weirdness that is commonly observed is that the axial position of the trapped particle (as seen by the side-view microscope) for a particular optical power is inconsistent with the known size of the microspheres, suggesting the trapped particle is either too light or too massive. The second commonly observed weirdness is that the image of the trapped particle in the side-view microscope is flashing wildly, either in intensity, or position of the particle within the image. When any such weirdness is observed, the trapped particle is usually abandoned by temporarily turning off the optical power, allowing the particle to fall under the influence of gravity. This selection step most often occurs before the bead dropper is retracted, so that the abandoned particle is caught by the bead catcher.

It remains somewhat mysterious why such a large proportion of the initially trapped objects exhibit this weird behavior. As will be discussed in Chapter 4, large populations of microspheres characterized with scanning electron microscopy, and from the same source

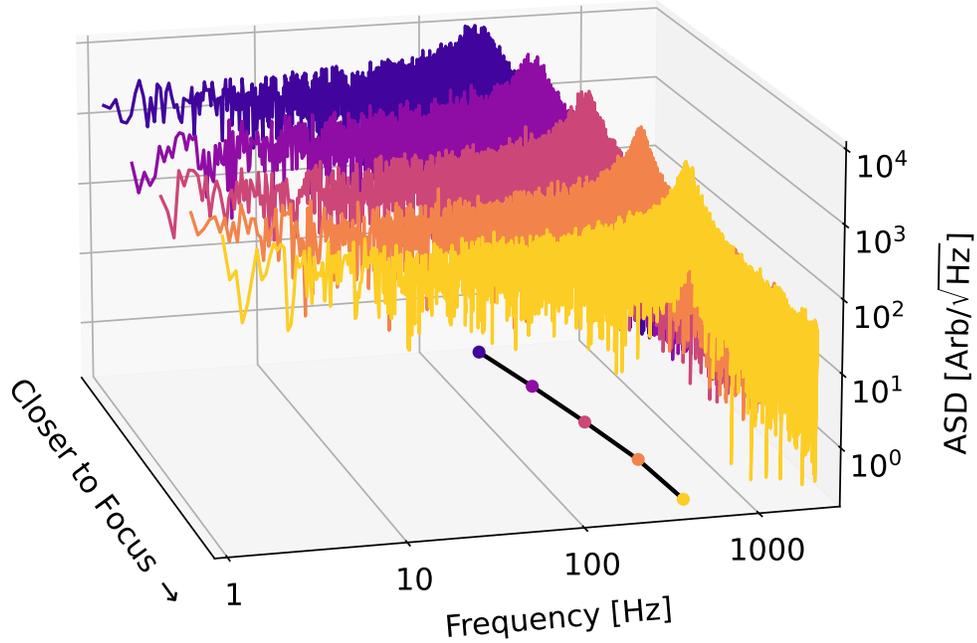


Figure 2.17: Amplitude spectral density (ASD) of microsphere motion along the radial axis x , for various different microsphere heights within the trap. The signal is uncalibrated. The resonant frequency from a fit of the ASD to Equation (B.4) is projected onto the horizontal plane below the ASDs. It is clear that as the height is lowered (warmer colors) the radial resonant frequency is increasing.

stock used for trapping, appear to not only be very consistent in size, but reasonably defect-free. It's possible that there are variations in internal structure, not visible to electron microscopy, that alter the interaction with the optical trap, but this is speculation.

Once a well-behaved microsphere has been trapped, and the dropper/catcher has been retracted, the optical power is lowered to bring the microsphere closer to the focus. With the 1.5 mbar of residual N_2 gas inside the chamber, the trapped microsphere is subject to collisions with gas particles which in turn impart momentum and drive the microsphere isotropically with white noise. Within the trap, the microsphere behaves as a damped harmonic oscillator, and thus has a very characteristic response spectrum, given by Equation (1.9). The frequency of the resonance within this spectrum is directly related to the optical spring constant, which is maximized at the focus, as was found in Chapter 1. Thus,

as the microsphere is lowered within the trap by reducing the optical power, the radial resonant frequencies increase and can be maximized, as shown in Figure 2.17. It is assumed that the microsphere resides at the focus of the trap as a result of this procedure.

2.7 Feedback and Calibration

Prior to pumping down the chamber for any type of precision measurement, a microsphere trapped by the above procedure must be stabilized with active feedback in order to retain it within the trap for long periods of time. Then, once the chamber has been brought to high-vacuum to reduce thermal effects, the response of the system to external forces needs to be calibrated, and the microsphere needs to be electrically neutralized to help minimize background forces from electrostatic effects. These topics are described below.

One of the advantages of optical tweezers is that the suspension and confinement of a trapped particle is determined almost entirely by the optical interaction between the trapping beam and the particle. However, this also means that fluctuations in this interaction directly affect the position and acceleration of the trapped particle. Fluctuations of the trapping beam's optical power, pointing, and polarization all apply forces and torques to the trapped particle, which can cause instabilities, as well as drive the microsphere through non-conservative orbits [30]. Additionally, at high-vacuum, there is very little damping from the residual gas, and the resonant response of the microsphere within the trap has a high quality factor. Fluctuations and transient excitations (i.e. acoustic noise, or impact-driven vibration of optical components) can drive this resonance and eject the trapped microsphere. Radiometric forces can also play a role at moderate pressures, although this hasn't been studied extensively in this work.

It should be noted that in some of Ashkin and Dziedzic's original work they were able to trap both silicone oil droplets and microspheres made from fused silica at high vacuum (1×10^{-6} mbar), without the use of feedback [8]. The authors attributed this achievement to passive stabilization of their beam steering optics, as well as the materials having low loss (as compared to Stöber microspheres), which reduce sources of technical noise that cause instabilities, and minimize radiometric forces during pumpdown, respectively. Recent work has demonstrated the stable trapping of microspheres produced via the Stöber process, without feedback and at high vacuum [34], a success due to the vast reduction of pointing fluctuations in their apparatus from a combination of different efforts.

2.7.1 Stabilizing Feedback

In any case, active feedback is used in this system in order to stabilize both the center-of-mass motion of optically trapped microspheres, as well as intensity fluctuations in the trapping beam. As will be discussed in detail Appendix D, estimates of a trapped microsphere’s position in x and y are calculated in real time by an FPGA, making use of the measured differences in interference photocurrent across a quadrant photodiode. Similarly, the FPGA computes an estimate of the z position from the measured phase of the retroreflected light, relative to a stationary reference beam. In contrast, the amplified photocurrent from the power-monitoring photodiode is digitized directly, and assumed to be proportional to the injected optical power \mathcal{P} . Control signals can then be computed by the same FPGA based on these estimates of x , y , z , and \mathcal{P} , together with user-defined parameters, which are ultimately used to exert forces on a trapped microsphere.

There are a multitude of different feedback algorithms one might consider using. Here, we adopt the Proportional-Derivative-Integral (PID) formalism to generate stabilizing feedback. Essentially, for a given coordinate $x_i \in \{x, y, z\}$ (or the power \mathcal{P}), an error signal $e_i(t) = x_i(t) - x_{0,i}$ is defined, where $x_{0,i}$ is the desired setpoint for the microsphere’s position along this coordinate. The resulting control signal is computed as,

$$PID_i = K_{P,i} e_i(t) + K_{D,i} \frac{de_i}{dt} + K_{I,i} \int_{t-T}^t e_i(t') dt', \quad (2.6)$$

where $K_{P,i}$, $K_{D,i}$, $K_{I,i}$, and T are user-defined quantities that control the scale of each of the separate feedback terms. The $K_{P,i}$ term, ‘proportional’, can be used to increase the apparent spring constant and stiffen the trap about the desired setpoint (note the similarity to a Hooke’s law potential). The $K_{D,i}$ term, ‘derivative’, is velocity dependent so it can be used to apply drag forces and increase the damping experienced within the trap. The $K_{I,i}$ term, ‘integral’, helps to offset slowly fluctuating forces that might push a trapped microsphere preferentially in one direction for a long period of time, for example due to a drift in alignment or collimation. In practice, the integral is performed with $T \rightarrow \infty$, with the caveat that the entire integral can be set to zero arbitrarily, essentially to ‘clear the memory’ of the integral gain term, sometimes necessary after changing parameters.

The first level of feedback implemented is independent of a trapped microsphere and is intended to stabilize intensity fluctuations of the trapping laser. The fiber-optic laser amplifier itself has an internally stabilized output which performs quite well across a range

of frequencies. However, thermal fluctuations in the fiber-optic components downstream of the amplifier add low-frequency noise that requires correction. It has been found that an integral gain term from the error signal for the measured power \mathcal{P} is sufficient. Feedback is accomplished by modulating the RF power sent to the trapping beam AOM, via an external modulation port on the signal generator. Note that scaling the value of $K_{I,i}$ is akin to changing the step response time of the feedback, as the integral will continue accumulating until the quantity $K_{I,i} \int (\dots)$ is large enough to offset the applied step.

With the power stabilized, error signals derived from the coordinates x , y , and z can be used to actuate components of the system in order to stabilize the center-of-mass motion of a trapped microsphere. Feedback forces can be applied radially, in the x and y coordinates of the horizontal plane, by manipulating the position of the optical trap with deflection optics. Recalling the input optics discussed in Section 2.2.2, there is a high-bandwidth, piezoelectrically driven deflection mirror (Thorlabs, ASM003), located along the waist of the collimated beam prior to its injection into the chamber. This focal plane (the beam waist) is a Fourier plane of the optical trap's focal plane, such that deflections of the collimated trapping beam here produce radial displacements of the optical trap itself, exerting a force on the trapped particle toward the new equilibrium position. Control signals for x and y calculated within the FPGA are the inputs to a multi-channel, high-voltage piezoelectric driver (Thorlabs, MDT693), which in turn drives the axes of the deflection mirror to produce displacement of the trap proportional to the applied voltage.

In the axial direction, feedback forces can be applied by modulating the optical power and thus the radiation pressure force applied to the trapped microsphere. As discussed in Chapter 1, an optical tweezer with a low numerical aperture essentially has no axial gradient force, and thus the axial equilibrium of a trapped microsphere within such a vertically-oriented tweezer is determined by a balance between gravity and radiation pressure. Modulation of the radiation pressure force is accomplished with the fiber-coupled AOM within the input optics (see Section 2.2.1), which tunes the optical power of the trapping beam injected into the chamber.

An immediate concern one might (rightfully) have is that the feedback to stabilize \mathcal{P} and z may interfere with one another. The two error signals are derived from different measurements that are not entirely independent (e.g. power fluctuations induce real forces on the microsphere, which can be observed in the z measurement), however there is only one actuator to stabilize both \mathcal{P} and z . Essentially, the combined feedback can ‘double-count’,

resulting in over-corrections and unstable behavior. To mitigate this effect, the error signal for the \mathcal{P} feedback is modified slightly $e_{\mathcal{P}}(t) \rightarrow e'_{\mathcal{P}}(t)$ where

$$e'_{\mathcal{P}}(t) = \mathcal{P}(t) - \mathcal{P}_0 - g[\text{PID}_z(t)], \quad (2.7)$$

where the function $g[\text{PID}_z(t)]$ is an empirically determined proportional scaling, although it can include phase shifts to offset the time delay from the nested feedback structure. The ideal solution to this double-counting would be to have two AOMs in series along the trapping beam, with a beam-sampler and power measurement photodiode nestled in between, such that the upstream AOM can be used to correct solely power fluctuations measured by the sampler and photodiode, and the downstream AOM can be used correct for actual axial motions of the microsphere measured by the retroreflected interferometer.

A simple block diagram of the feedback system as it exists is shown in Figure 2.18, included the nested structure with the laser power feedback.

The microsphere position along the axial coordinate is critical to control due to the geometry of the optical trap itself, as well as a more technical limitation. With a low numerical aperture and a vanishing axial gradient force near the focus, the optical spring constant in the absence of active feedback approaches 0 as the microsphere gets near the focus. Thus, the associated potential well is very shallow, and a trapped microsphere can easily be ejected by fluctuations in laser power or any of the aforementioned noise sources. Active feedback, specifically a proportional gain term, can provide a replacement restoring force to maintain confinement about a desired setpoint, usually near the focus. The value of $K_{P,i}$ can be increased arbitrarily to tune the spring constant, although signal-processing time delays inherent to the feedback eventually cause instability when $K_{P,i}$ becomes too large.

From a more technical standpoint, minor misalignment of the input beam can result in a tilted tweezer, where the direction of propagation is not truly vertical. This necessarily couples the imaging and feedback axes, as the z position interferometer is sensitive to axial motion within the trap, which now includes a horizontal component (visible to the $x - y$ interferometer) due to the tilt of the trap. Given the non-ideality of the input beam profile, in particular the mildly asymmetric tails (see Section 2.3.2), as well as the slow drift in optical alignment, some amount of this coupling is unavoidable. By reducing axial motions with active feedback, sensitivity can be maintained while also reducing the cross-talk from

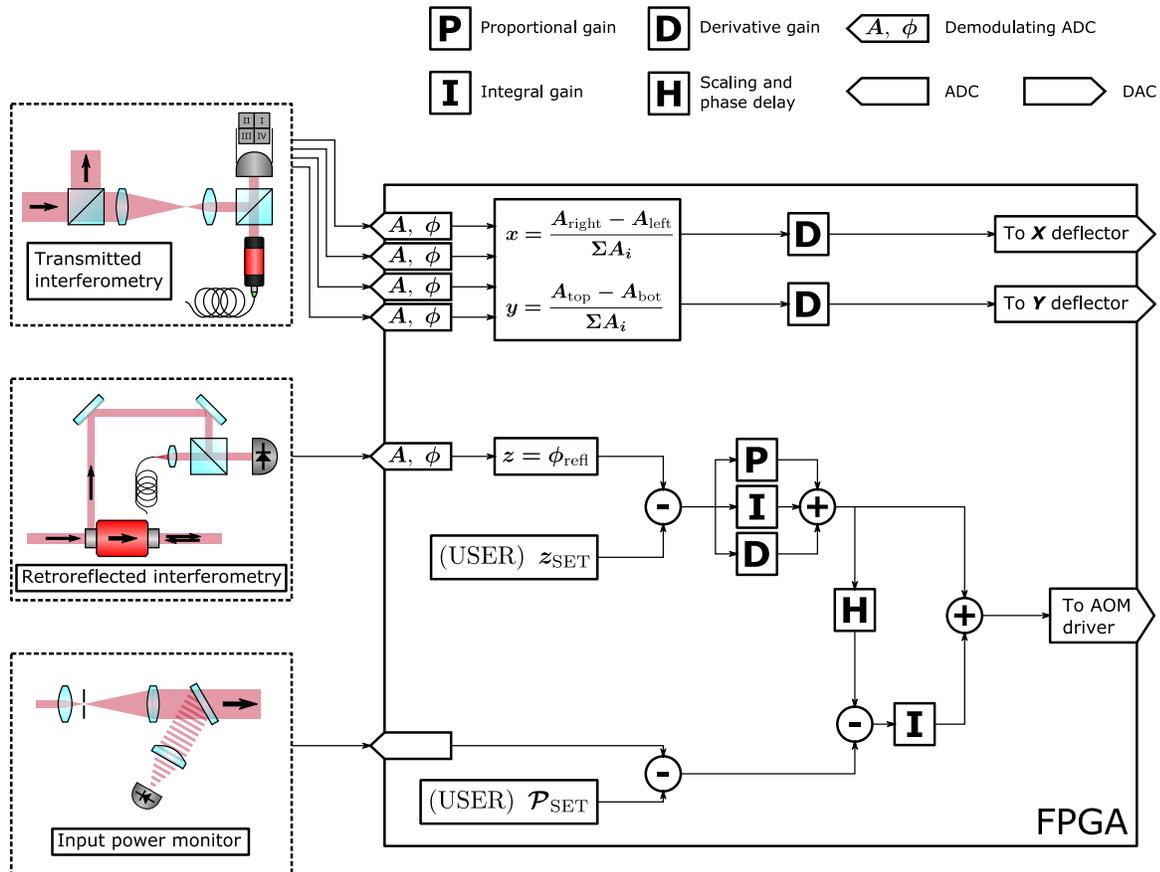


Figure 2.18: Block diagram of the feedback loop to stabilize a trapped microsphere’s center of mass motion. The x and y (radial) degrees of freedom only require damping and thus derivative feedback, while the z (axial) degree of freedom is stabilized with proportional, integral, and derivative gain. Both the measured power and retroreflected phase are used to construct a feedback signal for the axial degree of freedom, as there is only one actuator to control the trapping beam power. The **H** block assures that the two axial feedbacks don’t “see” each other, by removing the effect of the z -phase feedback signal from the power feedback error signal. The digital signal processing happens entirely within the FPGA.

z into either or both of the radial directions.

In practice, only derivative feedback is applied to the radial directions ($K_{P,i} = K_{I,i} = 0$ for $i \in \{x, y\}$) in order to damp the resonant motion of a trapped microsphere, as well as other fluctuations, while leaving the response at low frequency unfettered. For the axial direction, all three PID terms are used to tightly constrain the microsphere to a particular value of the retroreflected phase, and thus a fixed axial position. For each degree of freedom and feedback loop, the values of the gain terms are determined empirically, and can differ slightly between unique microspheres. In the radial directions, the value of $K_{D,i}$ is gradually increased until the observed damping just barely exceeds the natural damping from the ~ 1.5 mbar of residual N_2 inside the chamber, although slight adjustments may be made after the chamber is pumped to high vacuum. In the axial direction, the values of all $K_{I,i}$, $K_{P,i}$, and $K_{D,i}$ are individually increased (in that specific order) until the feedback loop exhibits ringing, at which point they are reduced by a factor of 5 – 10, with the intention of maximizing the action of the feedback whilst avoiding instabilities.

With the feedback initialized as described, the microsphere can continue to be lowered nearer to the focus by adjusting the setpoint of the z position feedback loop. By fully maximizing the radial resonant frequencies, we can ensure the microsphere is trapped at the focus, where the lateral extent of the trapping beam is minimized. Once the microsphere is safely at the focus with feedback engaged, the chamber can be opened to the vacuum system and pumped down to base pressures of $\sim 1 \times 10^{-6}$ mbar.

2.7.2 Microsphere Neutralization

Unlike many other micro- and nano-mechanical oscillators, an optically trapped microsphere has no electrical connections with the surrounding environment, by virtue of the optical suspension. During long-term measurements, changes in the overall charge state of the microsphere can thus be completely neglected, and the electrical isolation allows for a direct calibration of the system response to external electrostatic forces if the specific charge state is known.

Why, you might ask, are we suddenly talking about the charge of a macroscopic object composed of nominally neutral matter? Consider a candidate microsphere being ejected from the bead dropper via piezoelectric excitation. The fantastically unpredictable and poorly understood triboelectric effect [45] causes electrification of the microsphere during this ejection, such that when trapped, the microsphere is not electrically neutral, but instead

charged. Essentially, when two solids are in physical contact, there is some overlap between the electron clouds of neighboring atoms at the boundary layer between the two solids. As they are separated from one another, electrons can be redistributed between the overlapping electron clouds, resulting in electrification (a change in overall charge state) of both objects. If either or both objects are electrically isolated, such as an optically levitated microsphere, then they can retain the charge acquired in this electrification process.

Initially after trapping a microsphere, the exact charge state, q_{MS} , is unknown and can range from roughly $-2000e$ to $2000e$, with e the elementary charge. The sign and approximate scale of q_{MS} can be determined by applying an electric field to the charged microsphere and observing the response. This is usually accomplished by biasing two opposing pyramidal electrodes with oscillating voltages of the same frequency, but opposite sign, while keeping the remaining four grounded. This induces a nominally uniform, but time-varying electric field at the location of the trap, which applies a force to a charged microsphere within the trap along one of the coordinate axes defined by the choice of electrodes. The amplitude of the microsphere response to the oscillating field is indicative of the charge state, as under the assumption of quasi-static equilibrium, the applied electrostatic force $F_{\text{es},x}$, assumed to be along the x direction for this example, is equal in magnitude to the optical restoring force $F_{\text{opt},x}$,

$$F_{\text{es},x}(t) = q_{\text{MS}} E_x(t) = k_{\text{opt},x} \Delta x(t) = F_{\text{opt},x}(t), \quad (2.8)$$

where $E(t)$ is the time-varying electric field at the location of the trap, and $k_{\text{opt},x}$ is the optical spring constant associated to microsphere displacements along the x direction. With a finite element method, it is possible to compute the expected value of the electric field $E_x(t)$ for a given set of biases applied to the electrodes. On the other hand, despite the existence of numerical methods to produce solutions to the equations of Mie scattering theory [13], impurities, poorly understood bulk properties (e.g. index of refraction), non-spherical geometries, and non-ideal optical modes (both beam profile and divergence), all of which can fluctuate between successive microspheres and/or optical alignments, make it difficult for these solutions to be consistent with observed behavior.

Instead, an empirical calibration is used to determine k_{opt} , wherein the charge state is modified and the change in response is measured. The charge state of a trapped particle can be modulated with a variety of methods [24, 34, 46–49] (including corona discharge and

the generation of localized plasmas near the particle), and we make use of the technique where a charged microsphere is flashed with UV photons in order to eject electrons from the microsphere with a photoelectric mechanism. This works exceedingly well when the microsphere has a net negative charge, as it removes excess electrons. When the microsphere has a net positive charge, and thus a deficiency of electrons, ejecting more electrons further exacerbates the charge state. By introducing a gold-coated, silicon mechanical device in proximity to the microsphere, it is possible to reverse the direction of charging. When the surface of this device is in close proximity to the microsphere, flashes of UV photons eject electrons from the gold surface (much more readily than from the silica microsphere), many of which are then incident onto the microsphere changing the charge state from more positive, to more negative.

The presence of this device perturbs the electric field quite significantly, so it must eventually be removed to accurately characterize the relation in Equation (2.8). For this reason, if a trapped microsphere is found to have $q_{\text{MS}} > 0$, the gold-coated device is introduced to the trapping region, and UV photons are repeatedly flashed at the microsphere and device until the overall charge state of the microsphere becomes negative, i.e. with an excess of electrons. Then, the device is removed from the trapping region so that the electric field in the vicinity of the trapped microsphere can be well-modeled with finite element analysis.

Now, with a negatively charged microsphere, the rate of UV flashes is slowed and the amplitude of the microsphere response to an applied electric field is carefully characterized. Then, if the flashing rate is sufficiently slow and the measurement of the response sufficiently precise, it is possible to see individual electrons being ejected from the microsphere, characterized by a discrete step in the response amplitude. Flashing can then be continued until the microsphere reaches overall neutrality. An exemplary measurement of microsphere response amplitude during this discharging procedure is shown in Figure 2.19, together with a fit to a series of successive step functions. The step functions are all constrained to have a step size that is an integer multiple of a single free parameter, i.e. the microsphere response with a single excess electron. Often, steps are observed ranging in size from $1e$ to $5e$, which are all fit simultaneously, although only $1e$ and $2e$ steps are visible for the discharge of this particular microsphere. The presence of multiple step sizes within a single discharge further supports the supposition that the unit step size indeed corresponds to a single electron. For example, if the unit step size determined from the fit was actually $2e$, it is difficult to imagine a process that only ejects $2e$, $4e$, $6e$, etc., and not $3e$ or $5e$.

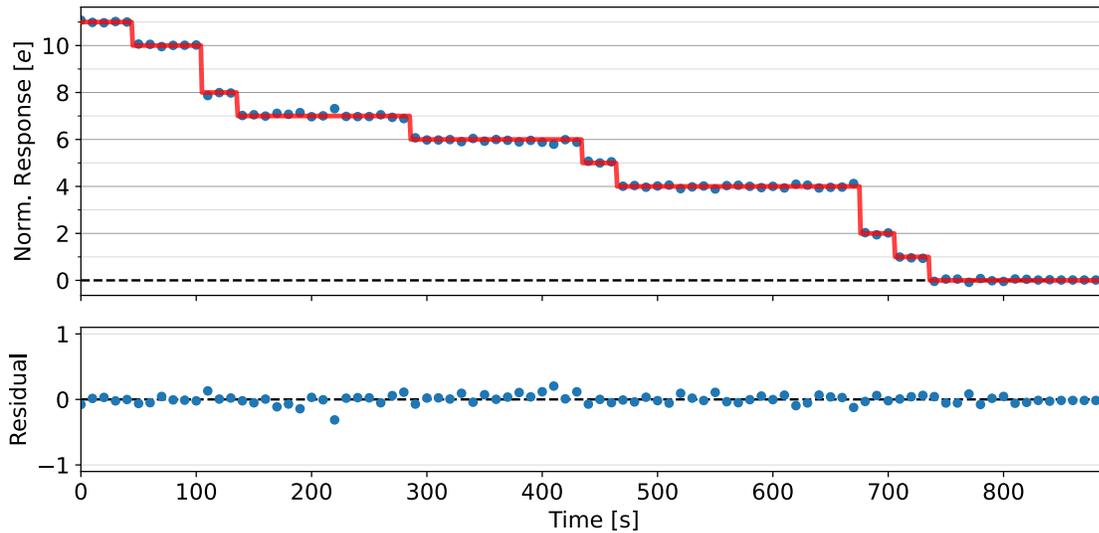


Figure 2.19: Amplitude of microsphere response to an applied electric field as a function of time during discharge. The magnitude of the electric field remains constant throughout the measurement, while flashes of UV photons induce discrete changes in the response amplitude. The changes in amplitude are found to be integer multiples of a unit step, corresponding to the expected response of the microsphere with a single electron charge. The response amplitude here has been normalized into units of electrons, so that the discrete behavior is exceedingly obvious. Additionally, the response amplitude has been fit to a series of successive step functions, where the step sizes are constrained to be integer multiples of a single free parameter: the single-electron response. In this way, it is possible to measure the optical spring constant (along the direction of the applied field), and determine an empirical calibration of the microsphere response to external forces.

The neutral state of a microsphere can be confirmed by massively increasing the applied electric field with high-voltage amplifiers, and looking for any residual response. In fact, carefully characterizing the residual response of a nominally neutral microsphere can yield constraints on the neutrality of matter, or the existence of milli-charged particles. Just such an investigation was the focus of some work with a previous iteration of the apparatus, completed prior to the beginning of the author's time at Stanford [24]. Since that time, work by colleagues has continued the investigation of the neutrality of matter making use of neutral microspheres and large electric fields [].

Let $R_i(ne)$ be the amplitude of the uncalibrated response in the $i \in x, y, z$ direction (where R_i is the result of the digital demodulation algorithm, as well as digital scaling), of a microsphere with n excess electrons, subject to an oscillating electric field with amplitude E_i and some frequency f_0 . Assuming the microsphere response is linear with the applied force, the measured step size, $R_i[n e] - R_i[(n - 1)e]$, is equivalent to the applied electric field multiplied by the fundamental charge, a fact that can be used to construct a calibration constant, κ_i , for the measured data,

$$\kappa_i = \frac{e E_i}{R_i[n e] - R_i[(n - 1)e]} \quad \text{with} \quad [\kappa_i] = \frac{\text{N}}{\text{Arb}}, \quad (2.9)$$

where the step size is in arbitrary units, although the absolute scale and precision of these arbitrary units are fixed by the signal-processing FPGA. This allows a direct calibration of the measured microsphere response to applied force, a critical component when using a trapped microsphere as a precision force sensor. The calibration method is direct, based on a fundamental constant e , and only makes two basic assumptions: that the response is linear for the applied forces used during calibration, which can be verified with a linearity measurement as in Reference [27], and that the electric field is well modeled by the finite element analysis discussed in Section 2.3. The only technical concern is a robust measurement of the applied voltages.

In contrast, one of the more common calibration techniques makes use of the thermal forces applied by collisions between a trapped microsphere and residual gas particles within the experimental chamber. Essentially, a moderate gas pressure is maintained within the chamber, and the response spectrum of the microsphere is measured. The power spectral density can then be fit to an analytic expression like Equation (1.9), so that the response can be interpreted in physical units. However, this requires knowledge of the microsphere

mass (often assumed from the manufacturer), knowledge of the microsphere shape (often assumed to be a perfect sphere), the assumption that the microsphere is thermalized with the gas, as well as the more fundamental assumption that a trapped microsphere behaves as an ideal Brownian particle within a confining potential. Clearly, the calibration based on the discharge of a microsphere is far superior.

2.7.3 Frequency Response of Trapped Microsphere

The discharge based calibration discussed in the previous section is powerful for calibrating the response of a trapped microsphere at the single frequency used to generate the oscillating electric field drive. However, from Chapter 1 we know that a microsphere confined to an optical potential has a non-negligible frequency response, usually assumed to be that of a damped harmonic oscillator. It is important to characterize this frequency response to fully utilize the force sensing capabilities of an optically trapped microsphere. Simultaneously, any cross-talk between imaging channels can be measured in a robust manner.

During the discharge, but prior to full neutralization, the UV flashing is ceased and the electrodes are configured to generate a frequency comb electrostatic drive of the form,

$$E_j(t) = E_{0,j} \sum_{n=1}^N \sin(2\pi f_n t + \phi_n), \quad (2.10)$$

where $E_{0,j}$ is the common amplitude of each component, $\{f_n\}$ is a set of nearly evenly-spaced frequencies ranging from 1 Hz to 600 Hz chosen such that they lie exactly centered within a single bin of a fast Fourier transform (where it is implied that the full integration time is both fixed and known) and none are integer multiples of another within the set, and ϕ_n is randomly sampled from a uniform distribution over the interval $[0, 2\pi)$ in order to avoid generating aggressive impulses from a superposition of in-phase sinusoids (consider the Fourier series of a delta function). This allows the microsphere response to be probed at multiple frequencies simultaneously.

A single electrode is driven while all others are kept grounded, generating an electric field like that described in Equation. (2.10) along one of the three Cartesian coordinate axes defined by the electrode cube. While this electric field is present, the response of the microsphere along all three degrees of freedom is measured. This process is repeated for each drive axis in order to fully characterize the response of a trapped microsphere to an external

force. Let $F_j(t) = q_{\text{MS}}E_j(t)$ be the applied force along axis $j \in \{x, y, z\}$, and let $R_i(t)$ be the response of the microsphere along axis $i \in \{x, y, z\}$. Writing their Fourier transforms as $\tilde{F}_j(\omega)$ and $\tilde{R}_i(\omega)$, respectively, the frequency-dependent transfer function of the microsphere response to external forces can be written as,

$$\mathbf{H}_{ij}(\omega) = \frac{\tilde{R}_i(\omega)}{\tilde{F}_j(\omega)}, \quad (2.11)$$

where \mathbf{H}_{ij} is clearly complex-valued, including both amplitude ratios and phase shifts for each drive frequency present in the comb. An example of the amplitude (phase) of a typical $\mathbf{H}_{ij}(\omega)$ is shown in Figure 2.20 (Figure 2.21). The on-diagonal components of this matrix, where $i = j$, are often well-approximated by the frequency response of a damped harmonic oscillator, and Figures 2.20 and 2.21 include the best-fit regressions of Equations (B.4) and (B.4). In contrast, the off-diagonal components are not well-approximated by these analytic functions, especially for drive frequencies near any of the resonances. For these terms, a functional expression for calibration is constructed by interpolating the result over the measured frequencies, and extrapolating at high and low frequencies with simple power laws. Depending on the value of the feedback coefficients for the z direction, and how well the nested feedback has been tuned, the $i = j = z$ component sometimes requires the same interpolation, as with strong feedback, this degree of freedom deviates from the expected behavior of a damped harmonic oscillator.

The values of the resonant frequency and damping coefficients from the fits shown in Figures 2.20 and 2.21 are enumerated in Table 2.1. For the x and y degrees of freedom, the resonant frequency is fixed by the interaction between the trapping beam and microsphere (dependent on the geometry of both, as well as the axial position of the microsphere within the optical field), whereas the damping is tunable by changing the value of the derivative gain, as discussed in Section 2.7. For the z degree of freedom, recall that near the focus the resonant frequency is expected to approach zero, so the measured frequency response is determined almost entirely by the active feedback, and can thus be tuned somewhat arbitrarily. Usually, the feedback for z is tuned to be the strongest possible while maintaining some degree of stability to large impulses, which have a tendency to induce ringing in systems operating near unstable points.

During this measurement, the charge state of the microsphere is not well-known, as the full discharge has not been performed, and thus the precise scale of the applied force is not

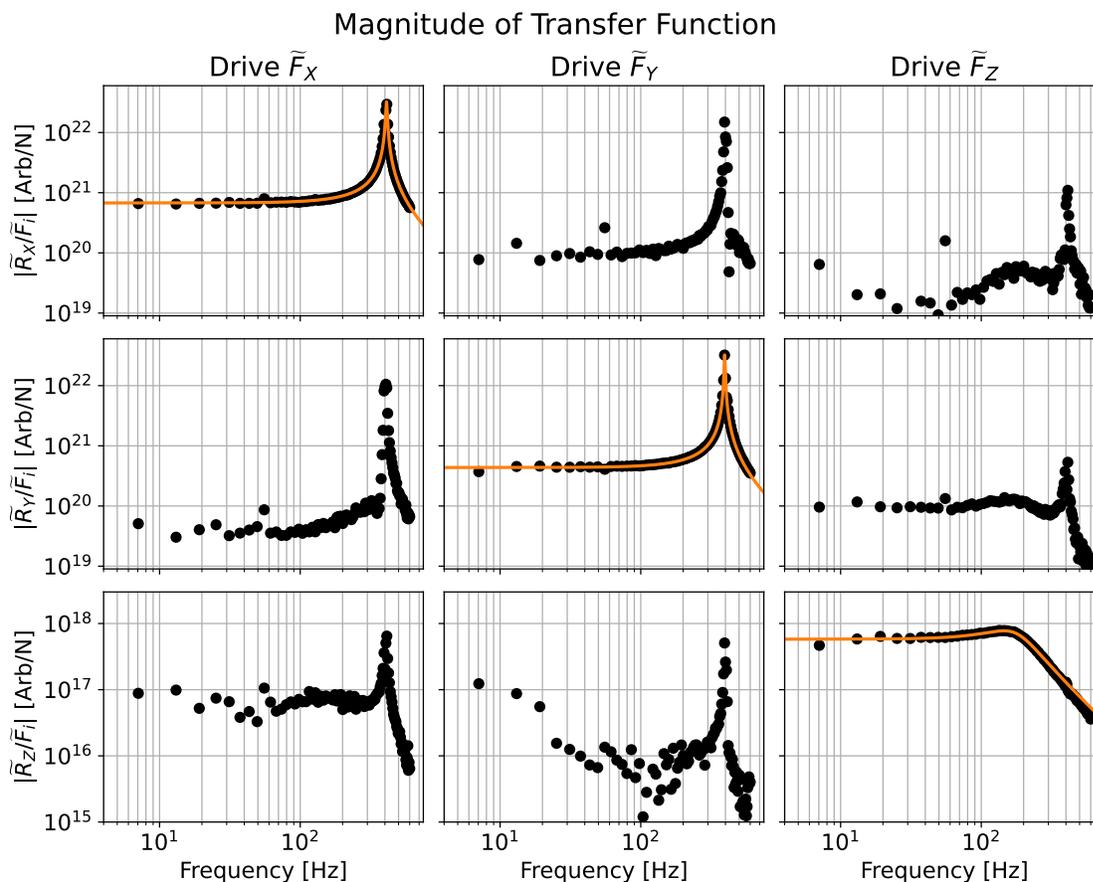


Figure 2.20: Amplitude a typical complex-valued, frequency-dependent transfer function, \mathbf{H}_{ij} , of an optically trapped microsphere. The amplitude scale is in units of [Arb/N], since the applied force is known in physical units, whereas the x , y , and z variables are in some units arbitrarily (but consistently) scaled by the FPGA. The rows don't necessarily have the same vertical extent, as each response axis can be scaled independently. The on-diagonal components include a fit to the expected response of a damped harmonic oscillator.

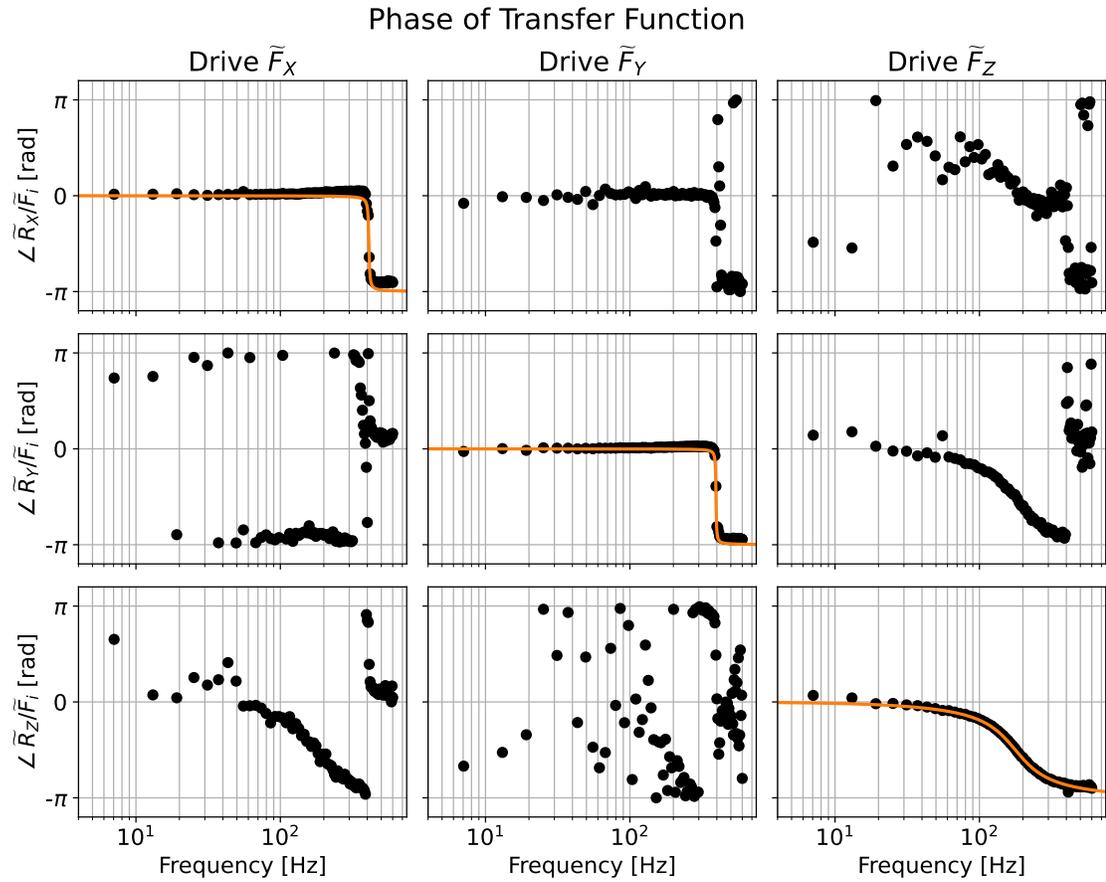


Figure 2.21: Phase of a typical complex-valued, frequency-dependent transfer function, \mathbf{H}_{ij} , of an optically trapped microsphere. The relative phase of the off-diagonal response has not been ‘unwrapped’ for ease of representation, although an interpolation (discussed in the text) would require this. The on-diagonal components include a fit to the expected response of a damped harmonic oscillator.

Table 2.1: Properties of the center of mass degrees of freedom, derived from fitting the amplitude and phase of a particular trapped microsphere’s response to external forces, seen in Figures 2.20 and 2.21. The coefficients in the equation of motion (Equation (B.2)) are in units of rad/s, whereas the coefficients are reported in units of Hz: $f_0 = \omega_0/2\pi$ and $\Gamma = \gamma/2\pi$. The damping coefficients for x and y are determined almost entirely by the derivative gain terms in the feedback, whereas both the resonant frequency and damping for z are set by the feedback. Integral gain terms primarily affect stability below the frequencies used in the measurement in Figures 2.20 and 2.21.

Axis	Resonant Frequency f_0 [Hz]	Damping Coefficient Γ [Hz]
x	410.3	8.6
y	394.7	5.3
z	177.1	144.6

well-known. Additionally, given the relatively small amplitude of each component of the frequency comb, a large charge ($q_{\text{MS}} \geq 50e$) is usually maintained during this measurement so that there is sufficient response amplitude to adequately characterize the full frequency dependence of the microsphere in a reasonable amount of time (measurement uncertainty scales as $1/\sqrt{T}$, with T in the integration time, but signal scales as q_{MS}). The discharge step size κ_i from Equation (2.9) is then used to estimate the charge on the microsphere during the frequency response measurement,

$$q_{\text{MS}} = \frac{\mathbf{H}_{xx}(\omega = 2\pi f_0)}{\kappa_x}, \quad (2.12)$$

where f_0 is the frequency of the applied electric field during the discharge measurement (usually 41.0 Hz), and where the single axis, x , that is used during the discharge measurement has been specified. With q_{MS} known during the frequency response measurement, $\mathbf{H}_{ij}(\omega)$ can be properly calibrated (see Equation (2.11)). Then, for any given measurement of the microsphere response, $R_i(t)$, along all axes i , the applied force along the geometric axes defined by the electrode cube can be reconstructed as

$$F_j(t) = \mathcal{F}^{-1} \left\{ \mathbf{H}_{ij}^{-1}(\omega) \mathcal{F}[R_i(t)] \right\}, \quad (2.13)$$

where \mathcal{F} is the Fourier transform operator. Although the formalism has been shown for continuous signals, it is implied that all measurements and drive signals are discretely sampled and thus the Fourier transform operator is implemented with the Fast Fourier Transform

(FFT) algorithm, usually limited to positive frequencies, since the signal is known to be real-valued. The details of the discrete mathematics are left to the reader. If they so wish, an interested reader may attempt navigating to a Python code repository designed exactly for these measurements, although they venture at their own risk into a sea of poorly commented, cryptically structured, spaghetti code [50].

Chapter 3

Three-dimensional Force-field Microscopy

The content of this chapter has been adapted from Reference [27], for which the author was a primary contributor. Apparatus and generic experimental method descriptions have been removed, appealing to Chapter 2 of this thesis instead.

3.1 Introduction

The ability to make measurements at ever smaller length scales has had profound implications for both fundamental science and technology. In particular, atomic force microscopy has enabled the measurement and manipulation of surfaces at atomic length scales. Traditionally, atomic force microscopes have sensed the interaction of a tip, suspended by a cantilever, with a surface, by measuring the displacements of the cantilever in the direction perpendicular to its surface [51]. The mechanical suspension of the force-sensing element limits electrical, thermal, and mechanical isolation from the outside world.

Here we present a technique for measuring three-dimensional forces over a three-dimensional volume by levitating a 4.7 μm diameter dielectric microsphere at the focus of a Gaussian laser beam, with a 10 μm minimum distance between the center of the microsphere and another object. This results in the full mapping of a vector field anywhere in space and close to mechanical objects. The optical levitation enables three important features: the absence of dissipation associated with the cantilever support allows measurements with substantially lower force noise at room temperature; the electric isolation provided by the optical support

makes electrostatic measurements at fixed charge possible; and force vectors measured in three dimensions are characterized by similar spring constants in each of the three trapping degrees of freedom.

Since the pioneering work of Ashkin [1, 7, 9], a number of experiments with optically levitated microspheres have been demonstrated, especially in recent years [21, 24–26, 31–33, 52–66]. Some authors have proposed short-range force detection experiments using their optically levitated microspheres [46, 67], but only a few have actually positioned free-standing objects micrometers away from a trapped microsphere [25, 66, 68] or other mesoscopic objects. Positioning a Si beam or other attractor close to a trapped microsphere and measuring its effect on the trapped microsphere are crucial for any short-range force-sensing application.

As a demonstration of this technique, we measure the electrostatic forces between a charged microsphere and an Au-coated Si beam (ACSB) structure. A force calibration, obtained by charging the microsphere with a unity charge and applying an external electric field (see Chapter 2), allows us to precisely determine the electric field due to a overall bias on the ACSB as well as infer the distribution of “patch potentials” on the Au surface in absolute terms.

3.2 Experimental Setup

Microspheres are prepared within the optical trap as outlined in Chapter 2. For all of the measurements detailed within this chapter, the experimental chamber is maintained at the system base pressure. Feedback is tuned for minimal damping in the x and y (radial) degrees of freedom, while strong proportional, integral, and derivative gain is used for the z (axial) degree of freedom. The system’s response to known electrostatic forces is calibrated with the feedback active, and at the system base pressure.

The ACSB is designed to have contrast in density (and baryon number) for a future experiment detailed in Chapter 8. This is achieved by alternating Au and Si fingers with 25 μm distance between the centers of contiguous fingers. A 3 μm thick Si fence bridges all Si fingers together, so that the Au fingers are entirely surrounded by Si. A 200 nm Au coating is then applied to all surfaces. A more complete description of the ACSB is provided in Reference [69] and Appendix G, while, for the purpose of the work described here, only this external layer of Au is relevant, given that only electrostatic forces are applied. The ACSB is 10 μm thick to minimize its interference with the tails of the Gaussian trapping beam.

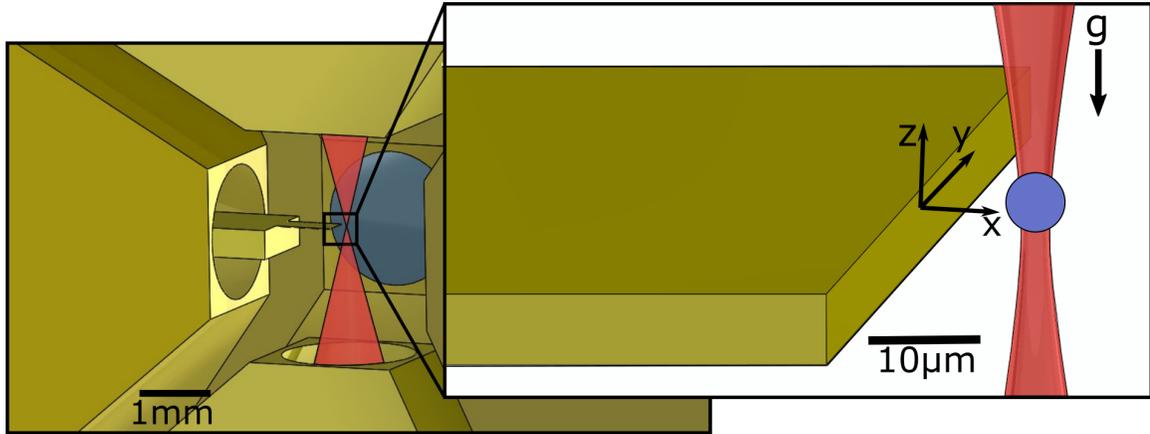


Figure 3.1: Trap region: A drawing of the trap region is shown in the left panel, illustrating the 2-mm-diameter holes in the Au-coated pyramidal shielding electrodes through which the ACSB and the trapping beam are brought in and extracted. The trapping light is represented by the red conical feature. The panel to the right illustrates a detail of the end of the ACSB and trapped microsphere, along with the coordinate frame used in the data analysis. A separate arrow shows the direction of Earth’s gravity.

More work on both shaping of the trapping beam and the microsphere imaging system are likely to make thicker field-generating components possible; nevertheless, constraints of this type are likely to remain the main limitation of this new technique.

After trap loading and force calibration, the ACSB is brought close to the trap region through one of the holes in the pyramidal electrodes, as shown in Figure 3.1, where the coordinate system employed here is defined. Sufficient clearance is provided for the full $80\ \mu\text{m}$ range of the translation stage in the y and z directions. Force measurements are carried out in three configurations: to evaluate the noise and linearity of the measurement, to demonstrate the ability of the technique to map the field produced by an external bias on the ACSB, and to map the field produced by patch potentials on the ACSB surface. In all measurements involving the ACSB, the trap is held at a fixed location, while the ACSB is displaced and/or actively driven.

For the electric field mapping discussed, the microsphere is charged to approximately $q \approx -400e$, with an excess of electrons. This large charge is intended to overwhelm possible multipolar (mainly dipolar) effects, clearly visible for neutral microspheres. A dipole moment of $d \approx 500e \times 5\ \mu\text{m}$ (aligned for maximal coupling to the field gradient) produces a force that is an order of magnitude smaller than the force on a charge of $q \approx 500e$ in the electric field

configuration employed here (discussed below).

The single-beam configuration of the apparatus is ideal for the application described here, as it allows optimal access to the microsphere from all directions in the horizontal plane. The current noise level on the motion of the microspheres is dominated by effects [26] other than the residual vacuum, which is limited at $\sim 10^{-6}$ mbar by the out-gassing of a number of translation stages.

3.3 Sensor linearity and noise

There are two features that define the performance of a force sensor: noise and linearity. The linearity of the force measurement is obtained, along the three degrees of freedom, by the same process used for the force calibration, scanning a 41-Hz AC drive signal with the shielding electrodes, from (0 ± 0.1) to 500 fN. The linearity has frequency dependence, which was measured as part of the calibration, as discussed in Sec. II. The result, for forces along the x axis, is shown in Figure 3.2, along with the residuals to a perfectly linear behavior. At driving forces over 300 fN, a $\sim 10\%$ non-linearity is observed.

The noise is characterized by collecting data with the ACSB placed over a three-dimensional grid of positions in front of a trapped and charged microsphere. The six trapping electrodes are nominally grounded and the ACSB is driven with a single-frequency AC voltage at 41 Hz. The response of the microsphere at a frequency far from this applied tone is then examined. Similar noise conditions are observed when the ACSB is nominally grounded. The resulting force micrographs are shown in Figure 3.3 for 10-s integrations at each position and with a single microsphere. The root-mean-square (RMS) noise forces over the whole $80 \times 80 \times 80 \mu\text{m}^3$ measurement volume are 5.5, 17.2, and 8.2 aN, for the x , y , and z directions, respectively. Histograms of the noise over the whole measurement volume are plotted in Figure 3.4. The anisotropy in the measured noise may be the result of residual astigmatism of the trapping beam, inconsistencies in x and y feedback, or, possibly, effects related to the finite geometry of the ACSB [70, 71]. This force noise is comparable to a cryogenic Si cantilever [72, 73], but is obtained at room temperature and on all three degrees of freedom simultaneously. This noise performance, along with the linearity up to $\sim 10^{-13}$ N, enables force measurements spanning over four orders of magnitude in amplitude.

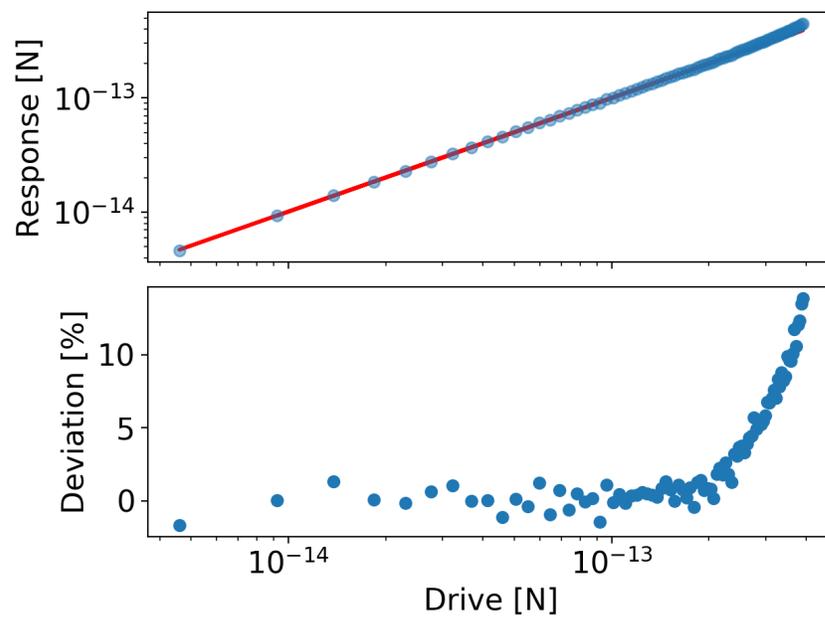


Figure 3.2: Linearity of the force sensor in the x direction, with residuals from perfectly linear behavior. Data are collected for a 41-Hz drive signal applied with a shielding electrode. The other two degrees of freedom have comparable linearity.

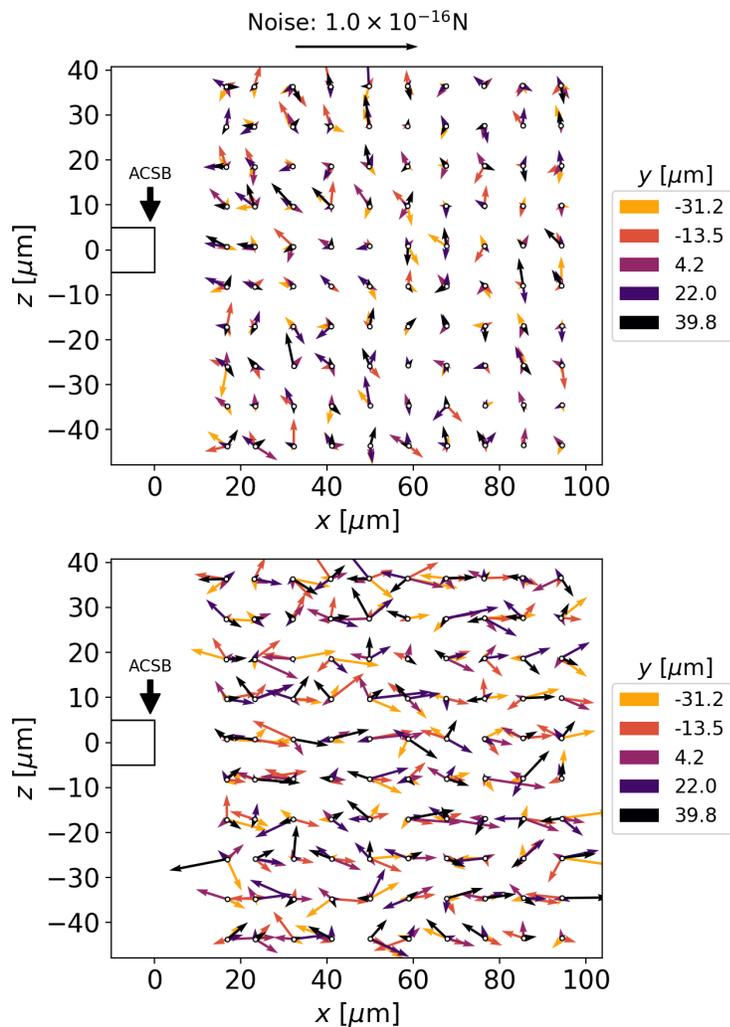


Figure 3.3: Force noise micrographs with a charged microsphere, represented with respect to a stationary ACSB at various relative positions. Practically, changes in the relative position are produced by displacing the ACSB, yet for ease of representation, the figure shows the ACSB as stationary. The force vectors represent the microsphere response at a frequency far from a single-frequency AC voltage drive applied to the ACSB during these measurements. For the data in the top (bottom) panel, the horizontal component of the force vector represents F_x , the force in the x direction (F_y , the force in the y direction), while the z component of the vector is always the force in the z direction. In both cases, force vectors measured at five y positions (with a different color indicating each y position shown in the legend) in a regularly spaced grid ranging from -40 to 40 μm are plotted at each (x, z) location.

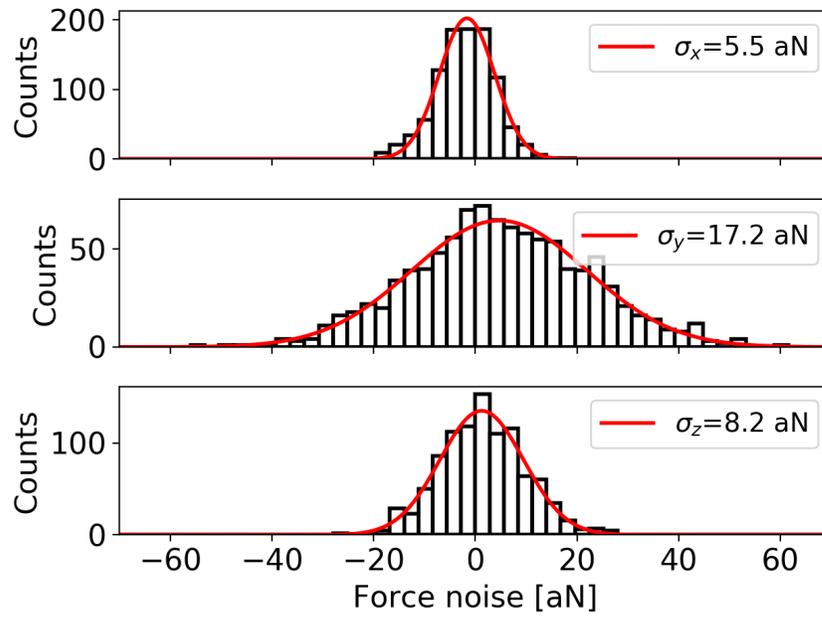


Figure 3.4: From top to bottom: histograms of the measured force noise in the x , y , and z directions at different grid locations. Also shown are the fits to normal distributions. The low noise and high linearity of the apparatus enable measurements with a dynamic range of over four orders of magnitude.

3.4 Electric field from an overall bias voltage

The overall electric field from a bias voltage applied to the ACSB is measured. This procedure is used here to validate the technique, but also to register the relative position of the ACSB to the trap and its orientation in space, with a fit to a model produced by FEA. In a new version of the trap, currently under construction, high-quality metrology will be possible through auxiliary optics, something only available in a rudimentary fashion in the current system.

The three-dimensional electric field is mapped over a $10 \times 10 \times 10$ grid of points, spanning the full $80 \mu\text{m}$ of closed-loop travel in the translation stage, along each of its three orthogonal axes. The relative displacement between points in the three-dimensional scan is known with an uncertainty of $\sim 10 \text{ nm}$, set by the accuracy of the translation stage used. To perform the measurement at each point on this grid, the microsphere is driven for 10 s by an AC voltage on the ACSB at 41 Hz, with a 100 mV peak-peak amplitude. The force at each grid point is represented by a three-dimensional vector. A slice of this vector field along an xz plane at $y = 0$ (centered along the y axis) is shown in Figure 3.5, together with the results of the FEA. Shown are both F_x and F_z (top) and F_y and F_z (bottom) in the same xz plane.

The data are fit to the FEA model by constructing a least-squared cost function from the difference between them, normalized by the error in the data, and summed over all 1000 grid locations. In the minimization, the charge of the microsphere, q is allowed to float (which is equivalent to applying an overall scaling of the \vec{E} field produced by ACSB), as well as three translations of the coordinates reported by the stage, and six independent rotation angles. Three of these angles floated in the fit represent angular misalignment between the axes of the translation stage and the axes of the trap, which are used to construct a rotation matrix applied to the measurement grid points. The remaining three angles account for a possible angular misalignment of the ACSB itself to the axes of the translation stage, and are used to construct a rotation matrix applied to the measured vector field. Constant offsets in the measured force, as would arise from contact potentials, are excluded since an oscillating electric field is applied and the amplitude and phase of the microsphere's response at the oscillation frequency is measured.

With this procedure, we find $q = -459e$ (consistent with the estimate made during the charging process), and a closest separation of the ACSB face of $15 \mu\text{m}$ from the center of the microsphere, centered in the z -axis with an uncertainty $2.5 \mu\text{m}$. The coordinate axes of

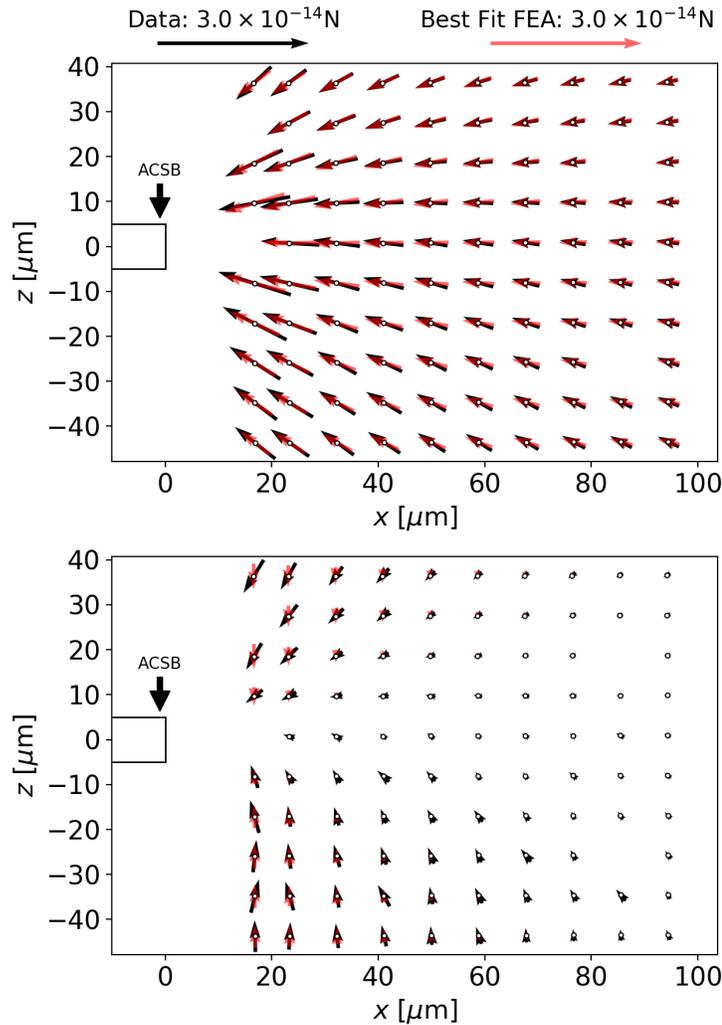


Figure 3.5: Vector field plots of (F_x, F_z) , top, and (F_y, F_z) , bottom, in an xz plane of relative positions. As before, changes in the relative position are produced by displacing the ACSB, yet for ease of representation, the figure shows the ACSB as stationary. The black arrows represent the measured force, while the red arrows represent the best fit from the FEA of the \vec{E} field produced by an ACSB with the same overall bias voltage as in the experiment. A few grid points are missing, due to data corruption introduced by the data acquisition instrumentation.

the ACSB are found to be tilted relative to the coordinate axes of the translation stage by no more than $\pm 5^\circ$ for all three rotation angles, while the coordinate axes of the translation stage are tilted relative to the physical axes of the trap by no more than $\pm 2^\circ$ for all three rotation angles.

Residual deviations from the best-fit FEA, particularly apparent at short separations, are likely the result of non-uniformity of the voltage on the ACSB, as our FEA assumes a perfect conductor with ideal geometry. Small dielectric particles (e.g., silica dust) on the surface of the ACSB, and/or metallic grains in the Au coating, may contribute to the small discrepancies observed at shorter distances [74, 75]. Permanent dipole moments in the microsphere, estimated from Reference [25], may produce forces more than an order of magnitude smaller than the residual force from the fit described above. Similarly, induced dipoles, estimated from the microsphere index of refraction [76] and the contact potential on the ACSB assumed to be < 100 mV, may produce a force more than two orders of magnitude smaller than the residuals.

The closest approach of the ACSB, as described above, is $15\ \mu\text{m}$ for the data set shown in Figure 3.5. Force sensing at smaller distances, down to $\sim 8\ \mu\text{m}$, is possible and has been achieved with this technique, but with an increase of the system noise for which a satisfactory explanation has not been found.

3.5 Patch potential measurements

Patch potentials on the surface of the 200 nm thick evaporated Au surface of the ACSB create an electric field which can be measured as a force on a charged microsphere. To perform this measurement, the ACSB is mechanically driven sinusoidally along the y axis with the translation stage, over a regular grid of x positions (separations), and z positions. At each point within the xz grid, $\vec{F}(t)$ and $y(t)$ are measured, so that the three-dimensional force field $\vec{F}(x, y, z)$ can be obtained. The relative registration of the microsphere with respect to the ACSB in terms of the three translations and six rotations is taken from the fit to the case of the biased ACSB, discussed in the previous section. The three-dimensional electric field can then be extracted as $\vec{E}(x, y, z) = \vec{F}(x, y, z)/q$, where q is the charge of the microsphere, also determined in the previous section.

The electric field due to patch potentials on the ACSB is numerically modeled following results obtained by Kelvin probe atomic force microscopy [75, 78]. Those authors define

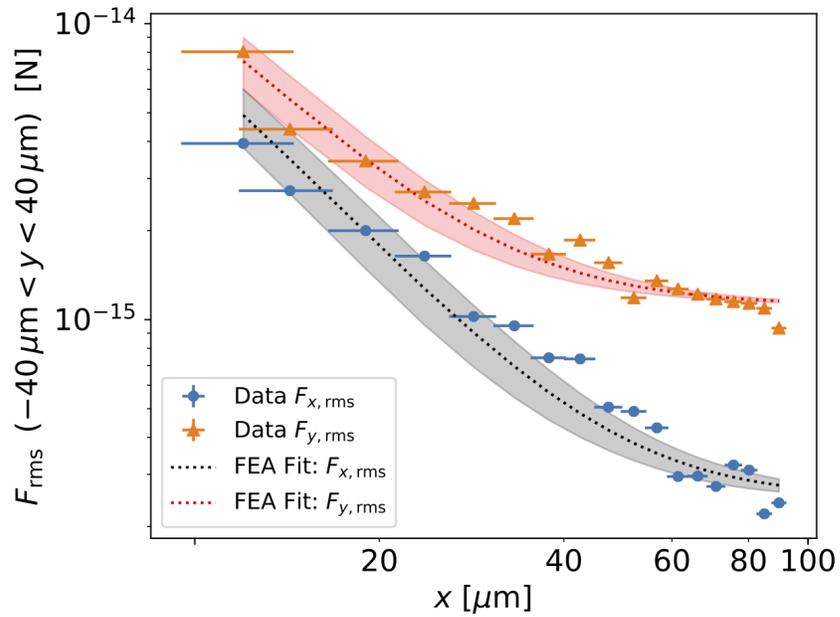


Figure 3.6: The RMS of the x and y components of the measured force on the microsphere in the plane of the ACSB, as a function of separation from the ACSB. The data are compared to a model described in the text. The bands represent the standard deviation of the model, using different random implementations of the patch potentials. The fit of $F_{x,RMS}$ ($F_{y,RMS}$) implies voltage patches of size $\sim 0.8 \mu\text{m}$ ($\sim 0.7 \mu\text{m}$), assuming $V_{\text{patch}} = 100 \text{ mV}$ [77, 78].

a voltage autocorrelation function, $R(\vec{s}) = \iint d^2 \vec{s}' V(\vec{s}) V(\vec{s}' + \vec{s})$, where \vec{s} and \vec{s}' are positions on the surface. They then find that $R(\vec{s})$ is approximately constant for short length scales, followed by a sharp knee at a length scale consistent with the expected patch size, l_{patch} . As an approximation to this autocorrelation function, a toy model is constructed with triangular patches measuring l_{patch} on a side, with voltages randomly sampled from a normal distribution, $N(0, V_{\text{patch}})$. FEA is used to determine the electric field due to these patches (as the geometry and boundary conditions do not permit an analytic solution), which is then compared to measured data. Since the patches and resulting electric fields are random, many realizations of patches are simulated in order to sample the mean and variance of the RMS force, F_{RMS} .

Measurements of F_{RMS} along x and y for different separations x are shown in Figure 3.6. The RMS is computed over all grid points along the y axis, at $z = 0$. In the figure, the data are fit to the results of the FEA model described above using a least-squared cost function, with a constant added variance to account for noise. Misalignments between the physical axes of the ACSB and the axes of the driven motion, together with an overall contact potential on the ACSB, would produce both a constant offset in $\vec{F}(t)$, as well as a term at the ACSB's driving frequency. Patch potentials with length scale $l_{\text{patch}} \leq 40 \mu\text{m}$ would only produce higher harmonics of the ACSB's driving frequency, as its maximum displacement is $(\Delta y)_{\text{max}} = 80 \mu\text{m}$. Thus, when computing $\vec{F}(y)$ from $\vec{F}(t)$ and $y(t)$ via Fourier techniques, the dc and fundamental terms of the ACSB's driving frequency were set to zero in the FFT of $\vec{F}(t)$. At distances much greater than the patch size, the RMS \vec{E} field is proportional to the product of the patch length scale, l_{patch} , and the RMS patch voltage, V_{patch} . Although the microsphere is not close enough to the ACSB to resolve the underlying patches, the patch length-scale voltage product can be extracted directly.

To compare to other measurements, we assume $V_{\text{patch}} = 100 \text{ mV}$ [77, 78] and obtain l_{patch} . The fit of the x RMS (y RMS) force implies $l_{\text{patch}} = 0.8 \mu\text{m}$ ($l_{\text{patch}} = 0.7 \mu\text{m}$). The length scales inferred from the fit are somewhat larger than the $\sim 200 \text{ nm}$ grain size expected for a 200 nm thick evaporated Au film [77, 78]. Others have observed sample contamination by dielectric adsorbates (such as silica dust) that become embedded in metallic surfaces, which can collect charge and affect measurements of patch potentials [74, 75]. Finite electrode geometry may also play a role [71].

3.6 Conclusion

We have demonstrated a technique capable of mapping, with sensitivity competitive to that of cryogenic atomic force microscope, a three-dimensional vector field over a volume of $\sim 10^6 \mu\text{m}^3$ in free space. The center of this sensor, a $4.7 \mu\text{m}$ diameter optically levitated dielectric microsphere, can be brought as close as $10 \mu\text{m}$ to a metallic surface, while performing the measurement. This instrument is used to map the patch potentials on an Au surface.

While this result is exclusively sensitive to electric fields, three-dimensional mapping of other physical fields appears possible. Magnetized microspheres are commercially available, and possibly neutral microspheres or microrods with large electric dipole moments could be used with enhanced sensitivity to electric field gradients. Our group is actively pursuing the technique to search for new long-range interaction coupling to mass or other intrinsic properties of matter at the micrometer scale.

Chapter 4

Direct Mass and Radius Measurements of Individual Microspheres

The content of this chapter has been adapted from Reference [19], for which the author was a primary contributor. Apparatus and generic experimental method descriptions have been removed, appealing to Chapter 2 of this thesis instead, while some other sections have been expanded for completeness.

4.1 Introduction

Optical trapping and manipulation of micron-sized dielectric particles in vacuum has been applied to optomechanics [1, 7, 21, 53, 59, 65] and cavity cooling [55, 57, 62], fundamental forces and interactions [24, 25, 58, 60, 64, 67, 79], quantum mechanics [80, 81], quantum information [32], and surface science [27]. In many of these applications, knowing the size, mass, and other characteristics of the trapped particles is critical to drawing conclusions about moments of inertia, optical spring constants, and force sensitivity.

We present a technique to measure the mass of individually trapped microspheres, by balancing a known electrostatic force, the optical levitation force, and Earth's gravity. The electrostatic force is extrapolated to the condition of no optical power to determine the gravitational force on the microsphere, and thus infer its mass. This measurement requires

fewer assumptions than other techniques [48, 82] and is found to be independent of environmental conditions. The method is applicable to particles of any size, in any scattering regime, provided that a component of the optical power opposes gravity, and the direction of the gravitational field can be controlled. The technique is demonstrated for two distinct sizes of microsphere within the same optical trap.

Similar electrodynamic balances have been used to stably trap and levitate aerosol particles [83–87] as a platform for studying such things as droplet evaporation. It is possible to estimate charge-to-mass ratios for micron-sized aerosol particles thus levitated, but practical constraints severely limit both the precision and accuracy of these estimations, as discussed in References [86, 87].

It may be possible to derive a direct relation between the optical power required to levitate a microsphere at the center of the trap and the mass of the microsphere using numerical methods to develop solutions to the wave equations of Mie scattering theory [13], as was shown in Chapter 1 for an ideal Gaussian beam and a perfectly homogeneous and spherical microsphere. However, performing the same calculating with quantitatively accurate results (beyond order of magnitude) requires a detailed understanding of the microsphere radius, non-sphericity, and index of refraction, as well as a full description of the optical potential in three dimensions. The technique described here bypasses these complications and their associated systematics, resulting in increased accuracy.

A subset of microspheres are also individually collected from the optical trap with use of a mechanical probe and imaged via scanning electron microscopy (SEM) to determine their radii. Knowing both the mass and the radius of individual microspheres, their density can be calculated. The radii determined from SEM images of those specific microspheres are compared with the radii determined from SEM images of large populations of approximately 1000 microspheres that have never been in the optical trap.

4.2 Experimental Setup

The optical trap used here is described in Chapter 2. The critical features of the apparatus for the measurement itself are the electrodes above and below the trap and the power monitoring photodiode. Active stabilizing feedback is used for all measurements, and thus the standard imaging system is required, although the x , y , and z responses of the microsphere

are not part of the final measurement. Importantly, the axial degree of freedom of the microsphere, stabilized by modulating the power of the trapping beam with an acousto-optic modulator, is held at a fixed position by strong proportional and integral feedback terms.

The masses and radii of two types of silica microspheres are measured, both of which produced via Stöber process [16], but of different sizes and sourced from two distinct manufacturers: $r = 2.35 \mu\text{m}$ from Bangs Laboratories [76], and $r = 3.76 \mu\text{m}$ from microParticles GmbH [88]. The two electrodes directly above and below a trapped microsphere are used to generate a uniform, slowly varying electric field at the trap location, exerting an axial force on a charged microsphere. The relation between the applied voltage and the electric field within the trapping region is modeled by finite-element analysis, with an uncertainty that is much smaller than any other systematic uncertainty.

After measuring their mass with the technique to be described, three microspheres are collected on the end of a polymer-coated silicon beam, described in References [27, 69], where they remain attached via van der Waals forces. Individual microspheres are addressed to particular locations, recognizable from features on the silicon beam. The silicon beam is then removed from the chamber, and the three microspheres are imaged by SEM to determine their individual radii. A population of microspheres of the same variety and lot as those used in the trap are also measured by SEM. For this purpose, a monolayer of microspheres is spread onto a silicon wafer and subsequently imaged by SEM. Various diffraction gratings [89] are used to calibrate the instrument at each of the magnifications.

4.3 Electrostatic Co-Levitation

Once a constant, known charge is obtained for a trapped microsphere, its axial position is fixed near the focus of the optical trap using the feedback. The slowly varying (0.5 Hz) electric field is applied in the vertical direction, while the power of the trapping beam injected into the chamber, controlled by the feedback, is monitored with a beam pick-off and a photodiode. As the applied electrostatic force increases, the axial feedback reduces the optical power required to maintain a net force of zero, counteracting gravity. The electric field can then be extrapolated to zero optical power, which allows a determination of the microsphere mass. This process is shown schematically in Figure 4.1. The case of zero optical power cannot be directly measured, as there is a minimum power necessary both to constrain the microsphere to the optical axis, and to generate sufficient back-reflected

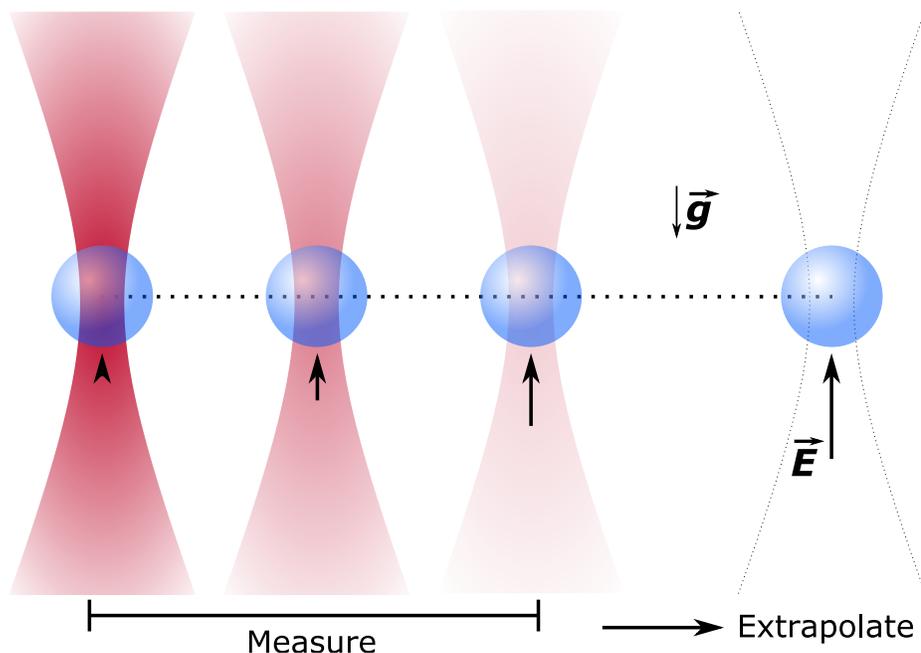


Figure 4.1: Illustration of the measurement technique. A charged microsphere is trapped by a Gaussian laser beam and held at fixed axial position with active feedback. A slowly varying electric field is applied, depicted with a black arrow. The active feedback reduces the optical power, indicated by the intensity of the trapping beam, such that the sum of the optical and electrostatic levitation forces opposing gravity is constant. The relation between optical power and applied field is then extrapolated to zero optical power, allowing a determination of mass from the implied electrostatic levitation field and the known charge.

light to measure the axial position via the methods described in References [26, 27]. The technique described is applicable only to single-beam traps [26], as its extension to systems with more than one beam requires care to account for the contributions of auxiliary beams to the total optical levitation force.

The equilibrium of axial forces F_z is expressed as,

$$\sum F_z = qE(t) - mg + F_{\text{opt},z}(t) = 0, \quad (4.1)$$

where q and m are the charge and mass of the microsphere, respectively, $g = 9.806 \text{ m/s}^2$ is the local gravitational field strength [90], $F_{\text{opt},z}(t)$ is the optical levitation force, assumed to be proportional to the trapping beam power, and $E(t)$ is the applied electric field strength. For each microsphere and charge-state combination, the slowly varying electric field and power are measured at least 50 times, each with a 50-s integration. Exemplary datasets for both the $r = 2.35 \mu\text{m}$ [76] and the $r = 3.76 \mu\text{m}$ [88] microspheres are shown in Figure 4.2, with the calculated masses from the extrapolation to zero optical power.

The extrapolation of this linear regression is performed over ≤ 1 order of magnitude and relies on only a few simple assumptions: firstly, the superposition principle, whereby the total force on the microsphere is the sum of gravity, the optical levitation force, and the electrostatic force, all of which are applied independently to the microsphere; secondly, proportionality between radiation pressure (the optical force) and photon flux, and thus the optical power [4]; and finally, the linearity of the photodetection system, which we operate at a factor of 1000 below the manufacturer’s reported saturation level [91].

The mass measurement is performed on 13 microspheres, of the $2.35 \mu\text{m}$ radius variety [76], in various charge states around $|q| = 20e$, with both signs of charge, as well as in two vacuum-pressure regimes: trapping pressure, approximately 1 mbar, and chamber base pressure, 10^{-5} mbar or less. The use of different pressures tests whether microsphere mass is lost due to heating, as reported for larger microspheres in Reference [63]. Cooling via residual gas decreases significantly with decreasing pressure, while absorption and scattering of laser light, the dominant heating mechanisms, remain constant. The results of mass measurements for all experimental conditions are shown in Figure 4.3, while the results from the final three microspheres later imaged by SEM are provided in Table 4.1.

To collect the final three microspheres, the polymer-coated silicon beam is rapidly inserted between the trapping laser and the microsphere, allowing the microsphere to fall under

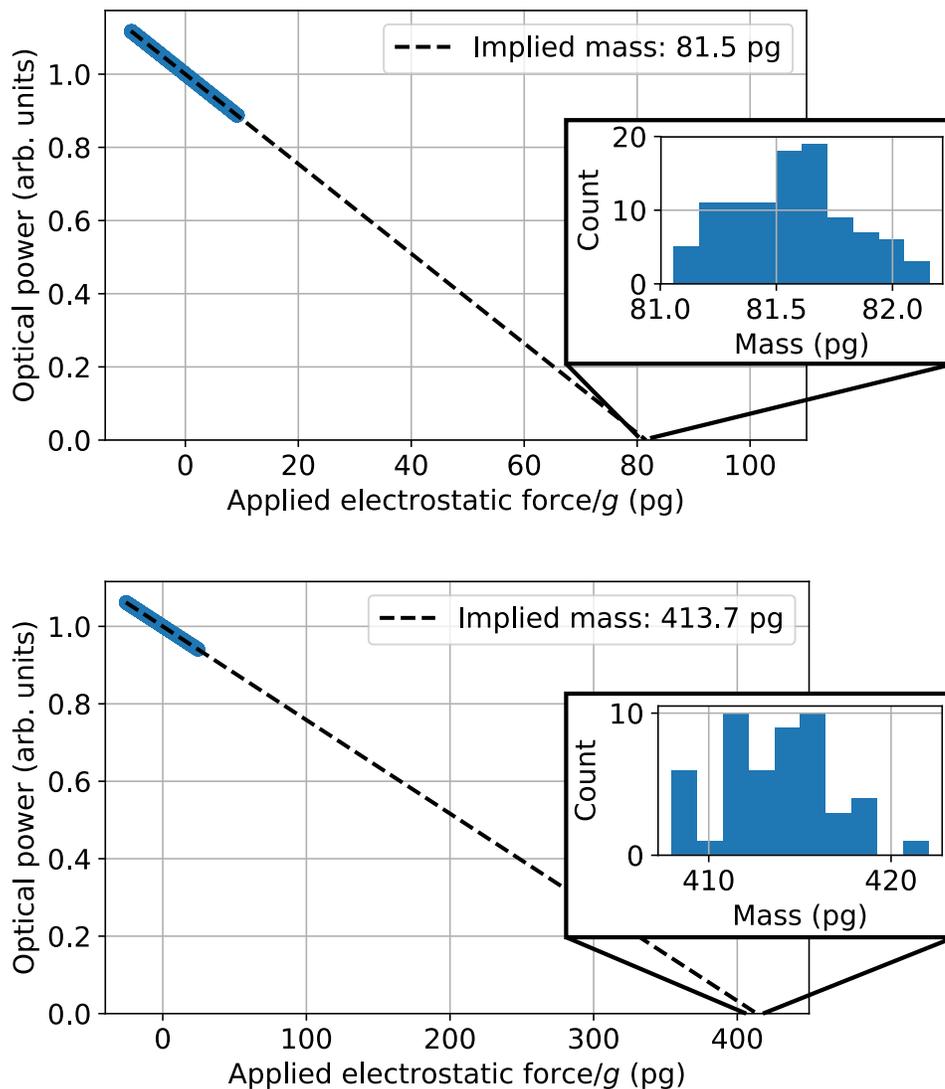


Figure 4.2: (upper) Normalized optical power versus applied electric field for 100 50-s integrations with a single, $2.35\ \mu\text{m}$ radius microsphere [76]. The extrapolation is performed separately for each integration. The mean of all extrapolations is shown with a dashed black line. The inset shows the distribution of the 100 extrapolated masses. (lower) The same style of measurement for 50 50-s integrations with a single, $3.76\ \mu\text{m}$ radius microsphere [88], demonstrating broader applicability of the measurement technique.

the influence of gravity. Each distinct microsphere can be associated with the respective mass measurement given its position relative to the internal structure of the silicon beam. Van der Waals forces, enhanced by the polymer, serve to keep the microspheres attached, whereas doped silicon and gold were both found to produce insufficient adhesion during previous attempts. The fluorocarbon polymer coating is made with a plasma-deposition technique inherent to the Bosch process [92], with use of C_4F_8 and SF_6 gases in a 1.5-kW inductively coupled plasma.

The same rigorous data collection was not performed with the $r = 3.76 \mu\text{m}$ microspheres, as these were only stably trapped following the initial publication of this work in Reference [19]. For this variety of microsphere, mass measurements of only two microspheres were made, both with a negative overall charge, and with vacuum maintained at the chamber base pressure.¹ The extrapolated masses for these two measurements are also included in Table 4.1. Neither of the two larger microspheres were collected for dedicated imaging, instead making use of SEM micrographs of a collection of microspheres sourced from the same batch in order to determine their radius.

Table 4.1: Microsphere masses, m , averaged over all experimental conditions; radii, r , averaged from two distinct high magnification images of the individual spheres (for the $r = 2.35 \mu\text{m}$) or the population average (for the $r = 3.76 \mu\text{m}$); and the derived density, ρ , for the three microspheres caught on the silicon beam, as well the two microspheres of the larger variety. All measurements include statistical and systematic uncertainties, and the relative contributions are shown explicitly for the measured masses. Numbers 4 and 5 were not individually imaged with SEM, and have significantly different statistical uncertainties due to an improvement in the electrostatic drive system.

Microsphere	m (pg)	r (μm)	ρ (g/cm^3)
1	84.0 ± 0.8 (stat.) ± 1.5 (sys.)	2.348 ± 0.038	1.550 ± 0.080
2	83.9 ± 1.1 (stat.) ± 1.5 (sys.)	2.345 ± 0.037	1.554 ± 0.079
3	85.5 ± 0.2 (stat.) ± 1.5 (sys.)	2.355 ± 0.038	1.562 ± 0.081
4	427 ± 20 (stat.) ± 8 (sys.)	3.76 ± 0.09^a	1.92 ± 0.16
5	413.7 ± 2.5 (stat.) ± 7 (sys.)	-	1.86 ± 0.13

^aDerived from the population measurements to be discussed in the following section, *not* from an individual measurement of the microsphere in question

¹The apparatus was optimally configured for measurements of the rotational dynamics of the smaller $r = 2.35 \mu\text{m}$, as will be discussed in Chapter 5.

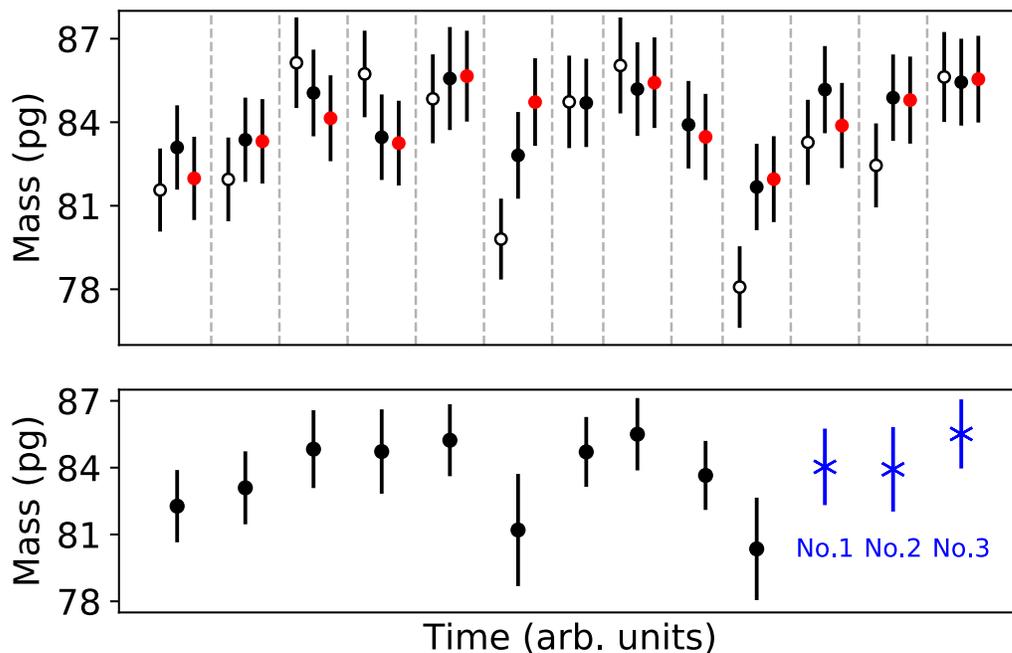


Figure 4.3: Measured microsphere masses in chronological order are shown in the top panel. Unfilled markers indicate a low-vacuum environment, $P = 1.5$ mbar, while filled markers indicate high-vacuum environment, $P = 10^{-6} - 10^{-5}$ mbar. Black markers correspond to measurements with a negatively charged microsphere, while red markers correspond to measurements with a positively charged microsphere. Different microspheres are separated by vertical dashed lines. The mean mass for each microsphere weighted over all experimental conditions is shown in the bottom panel. The blue data points with cross-shaped markers indicate the three microspheres imaged by SEM following their mass measurement.

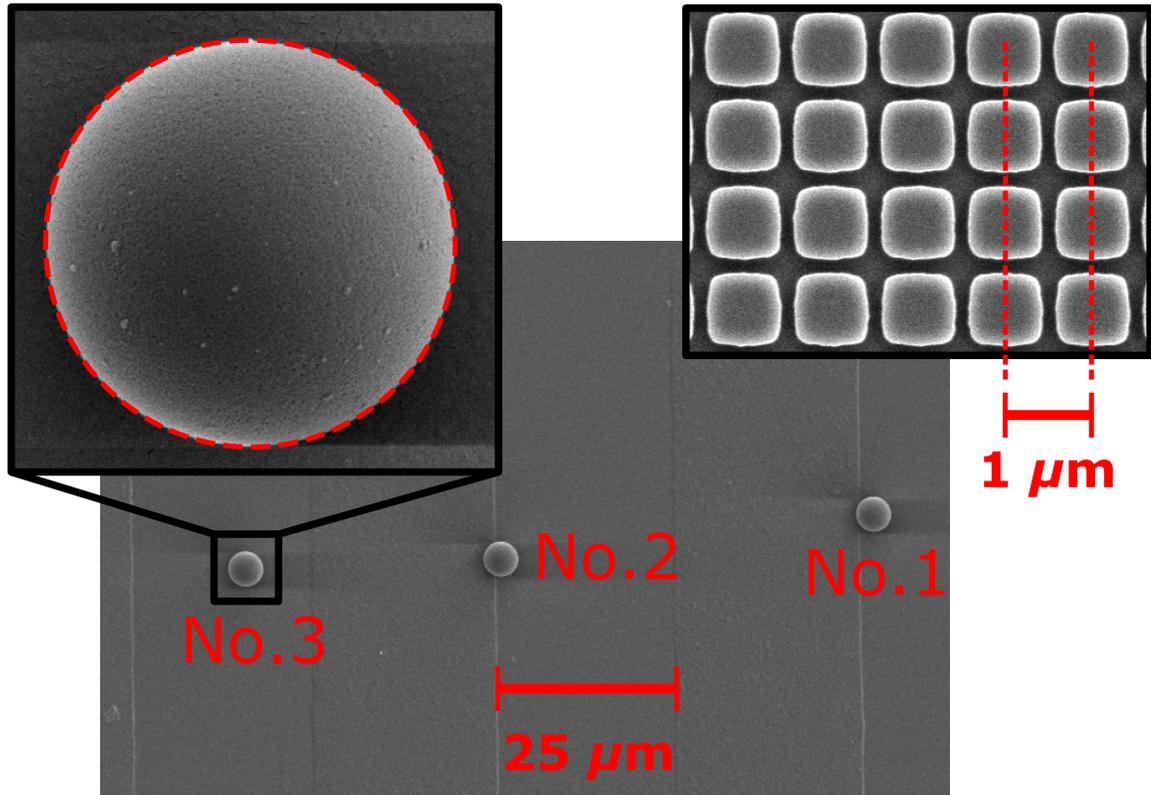


Figure 4.4: SEM images of the three microspheres collected on the silicon beam, at $\times 2500$ magnification. The left inset shows one microsphere at $\times 35000$ magnification, overlaid with the best-fit ellipse, and the right inset shows the $(1.000 \pm 0.005) \mu\text{m}$ diffraction grating [89], also seen at $\times 35000$ magnification. The diffraction grating serves as a calibration length scale for the high-magnification images of individual microspheres.

4.4 Radius Measurements from SEM

For the SEM measurements, the silicon beam with three microspheres is first sputter-coated with $(100 \pm 50) \text{ nm}$ of a Au/Pd alloy to prevent charging and the resulting microsphere ejection from the silicon beam. Charging effects from the scanning electron microscope are significantly exacerbated by the non-conductive polymer, necessitating the relatively thick metal coating. A diffraction grating with a pitch of $(1.000 \pm 0.005) \mu\text{m}$ [89] is used to calibrate the SEM images of individual microspheres at high magnification, as seen in Figure 4.4.

The microsphere diameter is first determined in terms of raw pixels. This is done by edge detection and contour tracing to outline the microspheres. The contour is then fit with an ellipse to account for real ellipticity in the microspheres, as well as astigmatism in the

electron microscope. The radius is taken as the average of the semi-major and semi-minor axes, which differ by less than 1%. A systematic uncertainty of ± 1 pixel in the determination of the semi-major and semi-minor axes is included.

At the same level of magnification, images of the calibration grating are used for conversion from pixels to physical distances. This is done by locating the centroids of the grating's repeated structure in the image, and averaging the pixel distance between neighboring centroids across the entire image. The ratio of grating pitch in microns to observed grating pitch in pixels serves to calibrate the images. The (100 ± 50) nm thickness of the conductive coating is subtracted from the final radius.

To characterize the distribution of radii by SEM, sparse monolayers of $r = 2.35 \mu\text{m}$ microspheres are prepared on two heavily doped silicon wafers. Charging effects are reduced without the polymer, so a conductive coating is not strictly necessary. However, the two wafers are imaged with and without a conductive coating to study possible systematics. The coated wafer is sputtered with (40 ± 10) nm of the same Au/Pd alloy mentioned previously. Each wafer includes a diffraction grating with $(9.98 \pm 0.02) \mu\text{m}$ pitch [89] to serve as a length calibration. Images of approximately 1000 distinct microspheres, both conductively coated and uncoated, are collected at a range of magnifications, together with images of the calibration grating. The same ellipse identification and calibration procedures used for images of individual microspheres are used to characterize the large microsphere populations.

The radii of the final three $r = 2.35 \mu\text{m}$ microspheres measured are compared with the distribution of radii from the microsphere population measurements, shown in the upper panel of Figure 4.5. The conductive coating reduces the apparent size of the microspheres by approximately 20 nm after accounting for the coating thickness. This may be the result of charging of the uncoated microspheres.

A similar procedure is used to characterize the distribution of radii of $r = 3.76 \mu\text{m}$ microspheres, although the sampled population is 127 microspheres instead of ~ 1000 . The distribution is shown in the lower panel of Figure 4.5. The microspheres are left uncoated, as a more advanced scanning electron microscope with magnetically-enhanced high resolution imaging, electron beam deceleration, and segmented backscatter detectors allows for robust imaging of dielectrics with limited surface charging. Images and thus radii are calibrated with the $(1.000 \pm 0.005) \mu\text{m}$ pitch diffraction grating.

It should be noted that the electron microscope has a built-in length scale for calibrating from pixels to physical distances at a particular magnification, displayed as a horizontal bar

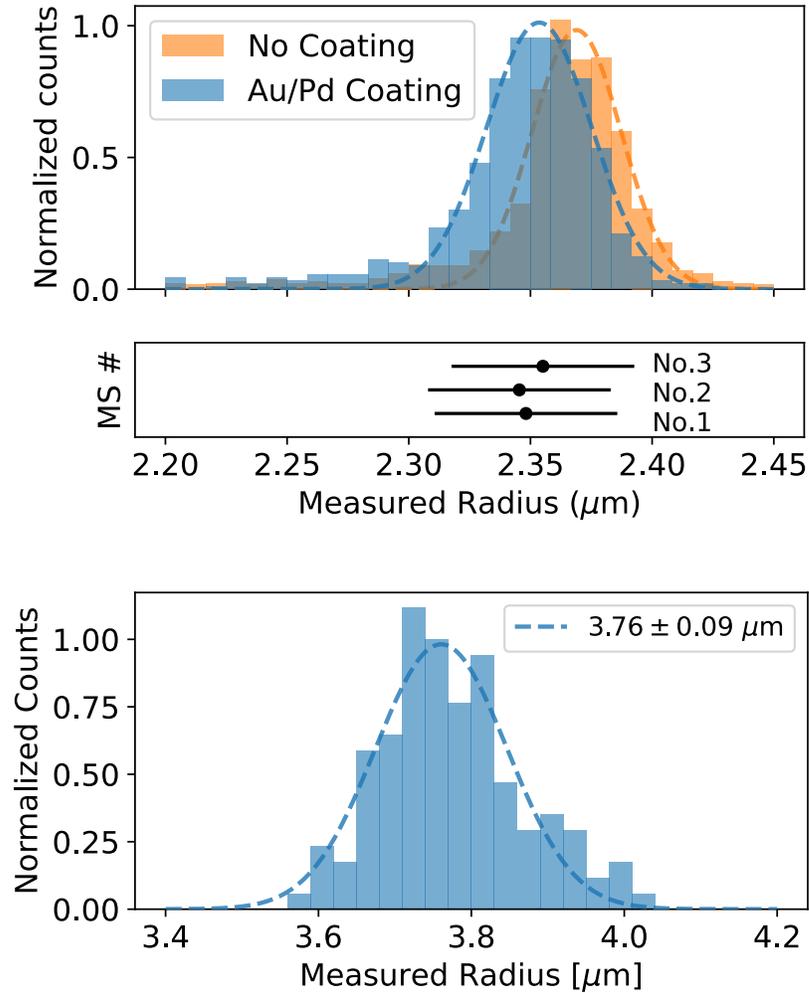


Figure 4.5: (upper) Distributions of microsphere radii measured by SEM for both conductively coated and uncoated microspheres of the smaller variety are shown in the top panel. Gaussian fits yield central values of $r_{\text{Au/Pd}} = (2.35 \pm 0.04) \mu\text{m}$ and $r_{\text{NC}} = (2.37 \pm 0.04) \mu\text{m}$ for the conductively coated and uncoated microspheres, respectively. Each distribution is generated from approximately 1000 distinct microspheres. Individual radius measurements from analysis of the microspheres seen in Figure 4.4 are shown in the panel immediately below. (lower) Distributions of microsphere radii uncoated microspheres of the larger variety, generated from only 127 distinct microspheres. A Gaussian fit yields a central value of $r_{\text{Au/Pd}} = (3.76 \pm 0.09) \mu\text{m}$. Individual microspheres of this size were not independently imaged, and thus only the population measurement is shown.

in the exported image. For both variety of SEM measurements (the individual and the large population) with the $r = 2.35 \mu\text{m}$ microspheres, the built in length scale was found to differ from the length scale derived from the calibration grating by less than 0.1%, suggesting the calibration grating was likely unnecessary given the magnitude of other systematic uncertainties. Further generic SEM imaging of different microspheres can then be performed without requiring a calibration. Regardless, all measurements shown here make use of either the $(1.000 \pm 0.005) \mu\text{m}$ or $(9.98 \pm 0.02) \mu\text{m}$ diffraction gratings.

4.5 Discussion

The mass measurement technique described avoids a number of systematic uncertainties inherent to derivations of the microsphere mass from optical properties and the trapping potential [13]. Importantly, the extrapolation to zero trapping-beam power is sensitive only to an offset in the power measurement, so an exact calibration of power is also unnecessary. Indeed, for the measurements reported here, the optical power is normalized to an arbitrary value of 1, as seen in Figure 4.2. The only requirement for the measurement is that the photodiode responds linearly to the incident optical power, which is easily achieved with a device operating well below saturation.

Fluctuations in the mass measurement over the 50 or more distinct 50-s integrations for a set of experimental conditions are normally distributed with a standard deviation on the order of 0.5 pg for the $r = 2.35 \mu\text{m}$ microspheres and 2.5 pg for the $r = 3.78 \mu\text{m}$, as seen in the insets in Figure 4.2. However, the total uncertainty of the measurement is dominated by common systematic effects that are enumerated in Table 4.2. Each effect listed is interpreted as an uncertainty on the applied electric field, the measured optical power, or the assumed value of g .

From Equation (4.1), these relative uncertainties directly propagate onto the extrapolated mass, whose uncertainty is computed as a quadrature sum of all contributions. The accuracy of the high-voltage amplifier's output monitor and the tolerance on the trapping lens focal length dominate the overall uncertainty. The second effect may offset the trap axially, thus sampling a different electric field strength. Each of the effects in Table 4.2 should result in a systematic shift common to all mass measurements. The total uncertainty obtained is 1.8%, which is included as a systematic uncertainty on the mean mass for each microsphere.

Table 4.2: Systematic effects on the mass measurement. The amplifier discussed produces the voltage driving the electrodes, and thus the electric field. Geometric misalignment, including optical tolerances, can change the value of the electric field at the location of the trap. The application and subsequent measurement of the electric field is also subject to systematic effects, and each measurement channel can experience electrical pickup. Effects are determined empirically where possible, or are obtained from instrument specifications.

Effect	Uncertainty ($\times 10^{-3}$)
Amplifier-monitor accuracy ^a	$\sigma_E/E \sim 15$
Lens focal length ^a	$\sigma_E/E \sim 10$
Amplifier-gain uncertainty ^a	$\sigma_E/E \sim 2$
Tilt of field axis	$\sigma_E/E \sim 2$
Tilt of trap (optical)	$\sigma_E/E \sim 1$
ADC offsets ^a	$\sigma_{\mathcal{P}}/\mathcal{P} \sim 1$
Electrode-voltage offset	$\sigma_E/E \sim 0.5$
dc-power offsets	$\sigma_{\mathcal{P}}/\mathcal{P} \sim 0.3$
Local g ^b	$\sigma_g/g \sim 0.1$
Fitting uncertainty	$\sigma_m/m \sim 0.1$
Electrical pickup	$\sigma_{\mathcal{P}}/\mathcal{P} \sim 0.02$

ADC, analog-to-digital converter.

^a From manufacturer datasheets.

^b Estimated from Reference [90].

We also observe scatter in the measured mass of a single microsphere between different experimental conditions, as seen in Figure 4.3. These variations could be due to a number of effects, including optical-path-length fluctuations in the axial feedback, electronic fluctuations in the axial feedback, and real changes in the mass of a microsphere. We do not observe any correlations between measured mass and experimental parameters such as the microsphere charge state or the vacuum pressure. The fluctuations are quantified by the standard deviation of measurements with different experimental conditions, which is included as part of the statistical uncertainty on the measured mass.

The measured masses and radii of the three individually imaged microspheres, Numbers 1-3, are shown in Table 4.1, together with their calculated individual densities. The densities of these first three microspheres are consistent and the average value, $\rho_{\text{small}} = (1.55 \pm 0.08) \text{ g/cm}^3$, is significantly smaller than that of amorphous fused silica, $\rho_{\text{SiO}_2} \approx 2.2 \text{ g/cm}^3$ [93], as well as the value provided by the manufacturer, $\rho_{\text{Bangs}} = 2.0 \text{ g/cm}^3$ [76]. Other, indirect measurements of the density of silica nanoparticles [48] also differ significantly from the values reported here. This could be the result of nonidentical synthesis conditions and post-synthesis treatment by different manufacturers, which can have a large effect on final particle porosity [94, 95].

Individual mass measurements and large population radius measurements of larger microspheres from a distinct manufacturer, Numbers 4 and 5 in Table 4.1, can also be used to estimate particle density. The calculated densities for these two microspheres are also consistent with one another, and have an average value of $\rho_{\text{large}} = (1.88 \pm 0.10) \text{ g/cm}^3$. This density also matches that provided by the manufacturer, $\rho_{\text{microParticles}} = 1.86 \text{ g/cm}^3$, which is based on a sedimentation test [96]. The simple fact that distinct manufacturers of nominally the same type of silica microspheres produce particles with clearly different densities, suggest that robust characterization is necessary for any application making use of these particles.

Spin-echo small-angle neutron-scattering measurements on particles synthesized via the Stöber have found an open-pore volume fraction of 32% and an inaccessible-pore volume fraction of 10% for particles with radius of approximately 80 nm [18]. It is distinctly possible that microspheres in solution absorb a nontrivial amount of water or other solvent, and that under low- to high-vacuum conditions, the liquid is removed, effectively lowering the mass and density. The classical electron oscillator model [4] implies that the reduced density

should result in a reduced refractive index

$$n_{\text{microsphere}}^2 - 1 = (n_{\text{SiO}_2}^2 - 1) \frac{\rho_{\text{microsphere}}}{\rho_{\text{SiO}_2}}, \quad (4.2)$$

leading to $n_{\text{small}} \approx 1.33$ and $n_{\text{large}} \approx 1.39$ (for Bangs Laboratories and microPartilcles GmbH, respectively) at 1064 nm, where $n_{\text{SiO}_2} = 1.45$ is the refractive index of fused silica [97]. These reduced densities necessarily change the optical interaction, requiring an updated simulation with the Optical Tweezers Toolbox (see Chapter 1) to better estimate the trapping forces and potential well associated to the trap.

4.6 Conclusion

We have presented a technique using an electrodynamic balance together with an optical tweezer to precisely measure the gravitational mass of optically levitated silica microspheres. The measurement is limited by systematic uncertainties of approximately 1.8% and is demonstrated to be independent of the (known) microsphere charge state, as well as the pressure of residual gas surrounding the microsphere.

After measuring their mass, three microspheres are collected from the trap with use of a mechanical probe and transferred to a scanning electron microscope, where their radii can be characterized. Together, the two precision characterizations allow direct calculation of the microsphere density. Large populations of microspheres used for trapping are also imaged. After accounting for the thickness of the coating, the individually measured radii of conductively coated microspheres are found to be consistent with the distribution of radii measured from the large population of conductively coated microspheres.

The apparent independence of the measured mass on the vacuum pressure, as well as the consistency between the measured radii of individual microspheres that were optically trapped and large populations of microspheres that were never trapped, both indicate a negligible loss of microsphere material by heating, under the environmental conditions tested for silica microspheres of radius approximately 2.35 μm . The simplicity and accuracy of the mass measurement, along with the reliable transfer of specific microspheres from the optical trap to air and subsequently to a different vacuum environment, opens the possibility for other correlated, precision measurements on microscopic objects.

Chapter 5

Electrically Driven Rotation and Residual Gas Sensing

The content of this chapter has been adapted from References [28, 38], for which the author was a primary contributor. Apparatus and generic experimental method descriptions have been removed, appealing to Chapter 2 of this thesis instead, while some other sections have been expanded for completeness.

5.1 Introduction

The ability to manipulate microscopic objects has found important applications in science and technology. The interest in optical levitation of dielectric objects in vacuum, pioneered by Ashkin and Dziedzic [1], has grown in recent times, with applications in precision measurements [24, 25, 46, 60, 79, 98], surface science [27] and quantum technology [21, 53, 57, 58, 62, 81, 99–101]. The suggestion was made in [1] that the rotational degrees of freedom (DOFs) of an optically levitated dielectric microsphere could be manipulated using the angular momentum in circularly polarized light. This has been realized in more recent times [52, 59, 61, 65, 102–111].

Here we present a novel technique for manipulating the rotational DOFs of an optically levitated microsphere by using electric fields to apply a torque to the $|\vec{d}| \sim 100 \text{ e}\mu\text{m}$ permanent electric dipole moment [25] found in $2.35 \mu\text{m}$ radius silica microspheres grown using the Stöber process [16]. The orientation of the dipole moment follows the orientation of the driving field so that the angular velocity of the microsphere, $\vec{\omega}_{\text{ms}}$, can be set in both

magnitude and direction [112]. Using the control afforded by electric torques, we observe the spin down of a microsphere suddenly released from a rotating electric field, measure the microsphere dipole moment by tuning the libration frequency, test the relationship between residual gas pressure and drag, and induce gyroscopic precession by rapidly changing the electric field axis of rotation. Because the electric field and dipole moment are both known in this system, precise quantitative measurements are possible. Similar techniques with an electric drive and optical readout have been demonstrated for graphene nanoplatelets in a Paul trap [113].

This technique has immediate applications in vacuum technology, which generally plays an integral role in science and technology. While a number of different gas-sensing technologies exist [40, 42, 43, 114, 115], under high vacuum conditions, absolutely calibrated pressure measurements are challenging. Some of the most sensitive pressure gauges ionize residual gas molecules and measure the resulting electrical current. Such ionization gauges require an empirical calibration, accounting for the efficiency of the ionization technique employed, which varies across molecular species and filament materials, and may change over time [116–118]. Additionally, the production of cations and their associated electrons, inherent to ionization gauges, can have a detrimental effect on conditions within an experimental chamber. At higher pressures, the measurement of heat transport in a gas does not require ionization but still requires empirical and species-dependent calibrations.

Gauges based on mechanical measurements and the kinetic theory of gases are known to provide absolute measurements of pressure. For example, capacitance manometers measure the force on a membrane exposed to the residual gas. Force sensitivity limits the pressure accessible to such gauges to approximately 10^{-5} mbar [40]. By comparison, spinning-rotor gauges measure the torsional drag induced by residual gas on a macroscopic rotor (an idea originally proposed by Maxwell [119]) which, in existing devices, is generally magnetically levitated. At high vacuum, in the molecular flow regime, such drag can be simply related to the pressure [104, 120–124], resulting in an absolute calibration. The minimum measurable pressure, approximately 5×10^{-7} mbar, is usually limited by systematic uncertainties or the required integration times [122, 125, 126].

With the rotational techniques developed here, we have devised a spinning-rotor vacuum gauge based on an optically levitated, electrically driven microsphere. The torsional drag on the microsphere can be measured by analyzing two distinct types of rotational dynamics of the trapped microsphere, where rotation around a single axis is measured optically via the

coupling of a linearly polarized trapping beam with a residually birefringent microsphere. In this scheme, the driving and interrogation mechanisms are not influenced by magnetic fields, and thus, the gauge can operate in conditions inaccessible to conventional spinning-rotor gauges. Another significant advantage of this technology, most applicable to experiments based on optically levitated particles, is that the measurement of pressure is relative to the environment in the immediate vicinity of the particle itself. A translation of the optical trap, or an array of optical traps each filled with a single rotating microsphere, would thus allow for mapping of pressure gradients with a spatial resolution limited only by the dimensions of the trapped particle(s).

5.2 Experimental Setup

The rotational response of a trapped microsphere, including an applied electric field, is described by the equation of motion:

$$\dot{\vec{L}} = \vec{T} = \vec{d} \times \vec{E} - \frac{\beta_{\text{rot}}}{I} \vec{L} + \vec{T}_{\text{opt}}, \quad (5.1)$$

where \vec{T} is the total torque, \vec{d} is the electric dipole moment, \vec{E} is the electric field, \vec{L} is the angular momentum related to the angular velocity by $\vec{\omega}_{\text{ms}} = \vec{L}/I$, β_{rot} is the rotational damping coefficient, I is the moment of inertia, and \vec{T}_{opt} is the optical torque. The part of the optical torque which does not average to zero over a rotation is generally negligible compared to the electric torques used here [127].

The angular velocity and the rotational phase of the microsphere are measured optically. As the microsphere rotates, it couples some of the incident linearly polarized optical power, \mathcal{P}_0 , into the cross-polarized optical power, \mathcal{P}_\perp , according to,

$$\mathcal{P}_\perp = \mathcal{P}_0 \sin(\eta/2)^2 \sin^2 \phi, \quad (5.2)$$

where η is the phase retardation between the fast and slow axes, and ϕ is the angular displacement of the microsphere relative to an origin in which the fast axis is aligned with the incident polarization [128]. The $\sin^2 \phi$ term implies that the phase of the microsphere is encoded as a modulation of \mathcal{P}_\perp , at twice the rotation frequency, ω_{ms} .

The optical trap used here is described in detail in Chapter 2. The critical features of the apparatus for the measurement itself are the six shielding electrodes used to apply arbitrary

electric fields and the cross-polarized light monitor to measure the microsphere response. Active stabilizing feedback is used for all measurements, and thus the standard imaging system is required, although the x , y , and z responses of the microsphere are not part of the final measurements presented. The dynamics of a rotating microsphere governed by Equation (5.1) will be studied with three distinct sets of initial conditions and steady-state drives. Two of these three cases will allow measurement of the rotational drag coefficient β_{rot} , and are used to implement a spinning-rotor vacuum gauge. The relevant equations will be derived prior to the presentation of the measurements themselves.

Each drive electrode is connected to a high-bandwidth, high-voltage amplifier driven by a digitally-synthesized analog signal, described in Appendix D. This apparatus is capable of producing arbitrary three-dimensional electric fields up to 100 kV/m in magnitude at frequencies as high as 1 Mrad/s, which limits the rotation frequencies achieved here. To produce a spinning electric field, a sinusoidal voltage is applied to a set of four electrodes in a plane, with a phase offset of $\pi/2$ between successive electrodes. A finite element analysis (FEA) is used to calculate the electric field based on the applied voltages.

Before performing a measurement, the microsphere is discharged as described in References [24, 27, 47] and detailed in Chapter 2. In addition, the microsphere is prepared in a state of known angular momentum and rotational phase by first dissipating any initial angular velocity using 0.1 mbar of N_2 gas. An electric field rotating at $\omega_0 \approx 2\pi$ rad/s with $E = 41$ kV/m is then turned on to align \vec{d} with \vec{E} . The chamber pressure is reduced to the base pressure of the vacuum system and ω_0 is increased at a rate of 300 rad/s² to the desired rotation frequency. Changes in the rotational dynamics in response to changes in the electric field magnitude and direction, as well as the damping coefficient, β_{rot} , can then be observed. The initial demonstration of the rotational dynamics presented here were collected with one specific microsphere, while three distinct microspheres were used for the spinning-rotor vacuum gauge. Similar qualitative behavior is observed across microspheres, and the measured dipole moments are representative of the population.

A typical amplitude spectrum of \mathcal{P}_\perp , in the region around twice the drive frequency, with $\omega_0 = 100\pi$ krad/s is shown in the top panel of Figure 5.1. A clear peak which follows the frequency, ω_0 , of the electric driving signal is observed at $2\omega_0$. The ~ 10 ppm amplitude modulation caused by rotation of the microsphere implies that the relative phase retardation of this microsphere is $\eta \sim 10^{-2}$. The prominent sidebands are caused by harmonic oscillation of the dipole about the electric field, which can be seen by demodulating the phase of \mathcal{P}_\perp

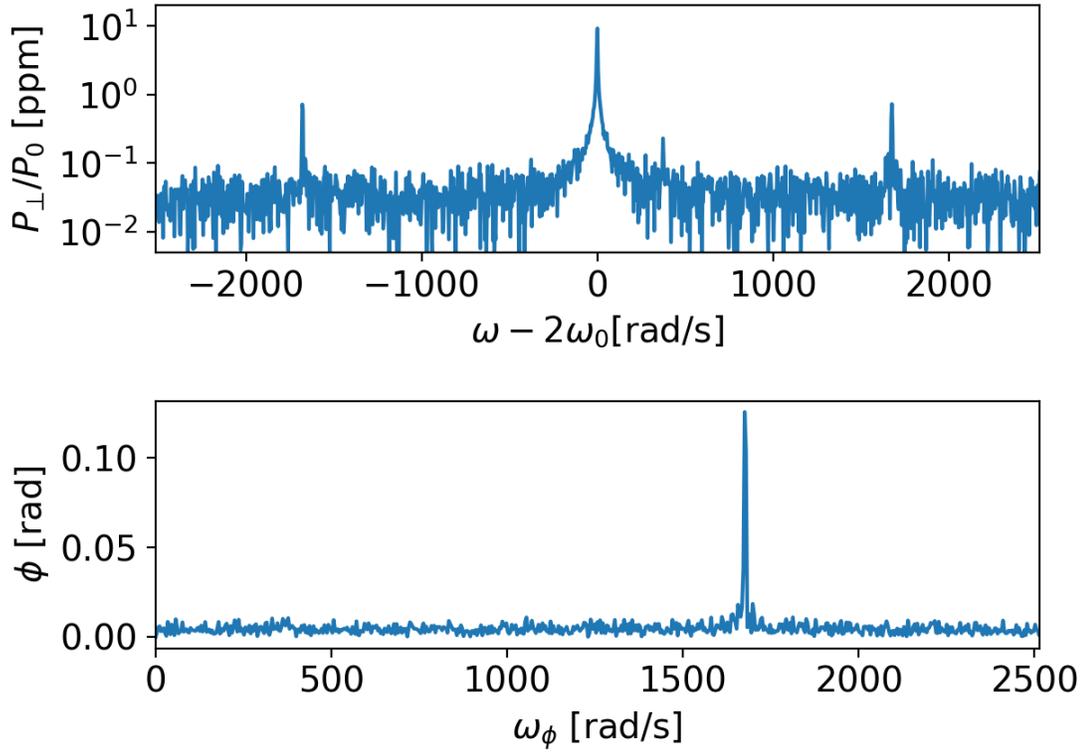


Figure 5.1: Top: typical amplitude spectrum of P_{\perp} for a microsphere prepared in a state of angular momentum pointing along the \hat{z} direction. The microsphere is spinning with $\omega_{\text{ms}} = 100\pi$ krad/s, driven by an electric field with $E = 27$ kV/m. The prominent line with sidebands are signatures of the microsphere rotation, with the position of the central line at twice the drive frequency. The sidebands are phase modulation of the rotation frequency, as shown in the bottom panel.

relative to the electric field carrier signal as shown in the bottom panel of Figure 5.1.

5.3 Libration

The simplest dynamics with the electric dipole interacting with a rotating electric field can be analyzed in the case where the electric field is rotating about a fixed axis with the dipole lying in the same plane as the electric field. In the frame co-rotating with \vec{E} , Equation 5.1 reduces to an equation of motion for ϕ , the angle between the electric field and the dipole,

$$\ddot{\phi} = -\omega_0\Omega \sin(\phi) - \frac{1}{\tau}(\omega_0 + \dot{\phi}), \quad (5.3)$$

where,

$$\Omega \equiv dE/(I\omega_0). \quad (5.4)$$

For sufficiently low damping, $\tau\Omega > 1$, this equation has an equilibrium solution,

$$\phi_{eq} = -\arcsin\left(\frac{1}{\tau\Omega}\right) = -\arcsin\left(\frac{\beta_{rot}\omega_0}{dE}\right), \quad (5.5)$$

and can be linearized to give harmonic oscillation at the frequency

$$\omega_\phi = \sqrt{\cos(\phi_{eq})\omega_0\Omega} = \sqrt{\cos(\phi_{eq})\frac{Ed}{I}}. \quad (5.6)$$

This results in the sidebands shown in Figure 5.1. The dependence of ω_ϕ on the magnitude of the driving electric field, E , is well-modeled by Equation (5.6), as shown in Figure 5.2. The equilibrium phase lag ϕ_{eq} may be neglected because $\tau\Omega \gg 1$ at the base pressure of the vacuum system. The fit shown in Figure 5.2 extracts the ratio d/I , which can be used to determine the dipole moment, d , if the microsphere is assumed to be a uniform sphere with the radius, $r_{ms} = (2.35 \pm 0.02) \mu\text{m}$ and mass $M_{ms} = 85 \pm 9 \text{ pg}$, measured for this lot of microsphere in Reference [19] (see Chapter 4). This procedure gives $I = (1.9 \pm 0.2) \times 10^{-25} \text{ kg m}^2$, which implies $d = 127 \pm 14 \text{ e } \mu\text{m}$, in agreement with Reference [25].

The ability to measure the permanent electric dipole moment from the frequency of harmonic oscillation enables precise measurements of torques on an optically levitated particle in high vacuum by balancing an unknown torque against a known electric torque. Importantly, this measurement technique also yields another, previously unknown, physical

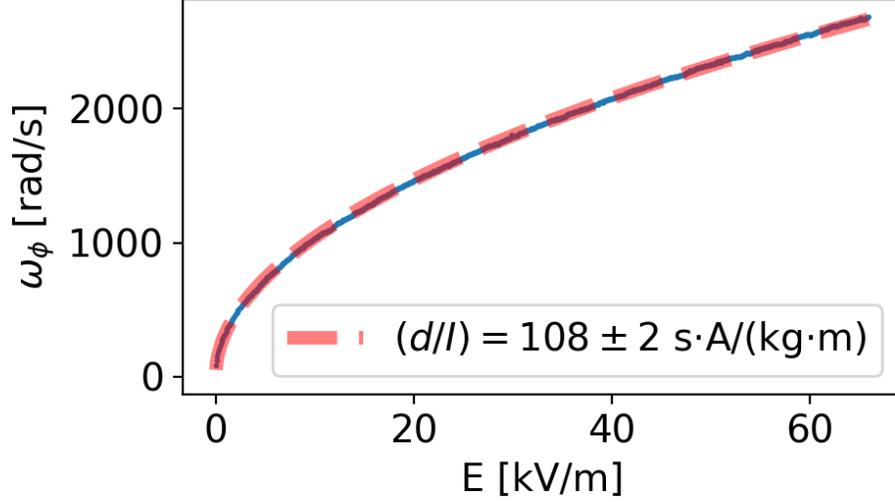


Figure 5.2: Harmonic oscillation frequency, ω_ϕ , versus driving electric field amplitude, E , for a microsphere spinning at $\omega_{\text{ms}} = 100\pi$ krad/s at a pressure of 2×10^{-6} mbar. The data is fit to Equation (5.6), obtaining $(d/I) = (108 \pm 2)$ A s/kg m and $d = 127 \pm 14$ e μm .

property of the test masses, which can be used to better understand background forces in short-distance force sensing applications and will also serve to calibrate some of the measurements to be described below.

5.4 Gyroscopic precession

When the electric field rotation axis and the angular momentum are not aligned, the dynamics are complex and depend sensitively on the initial conditions. For $\Omega \ll \omega_0$, and a field $\vec{E}(t) = E (\cos(\omega_0 t) \hat{z} - \sin(\omega_0 t) \hat{y})$, rotating about the \hat{x} axis, an approximate solution is given by,

$$\vec{L}(t) = L \{ \cos [(\Omega/2)t] \hat{z} + \sin [(\Omega/2)t] \hat{y} \}, \quad (5.7)$$

and,

$$\begin{aligned} \vec{d}(t) = -d \left\{ \right. & \cos [(\omega_0 + \Omega/2)t] \hat{x} \\ & + \sin [(\omega_0 + \Omega/2)t] \cos [(\Omega/2)t] \hat{y} \\ & \left. - \sin [(\omega_0 + \Omega/2)t] \sin [(\Omega/2)t] \hat{z} \right\}, \end{aligned} \quad (5.8)$$

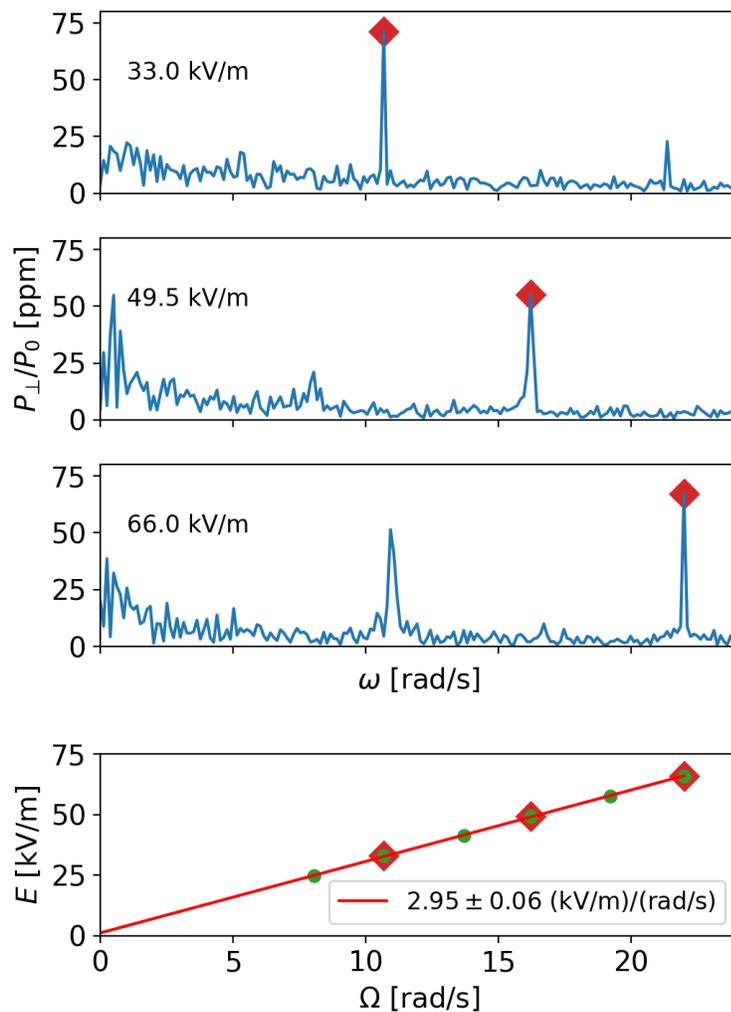


Figure 5.3: Top three panels: Spectra of the cross-polarized light intensity, \mathcal{P}_{\perp} , for a microsphere precessing about the \hat{x} axis while spinning at $\omega_{\text{ms}} = 100\pi$ krad/s). The modulation of the cross-polarized light occurs predominantly at twice the precession frequency, denoted by red diamonds. Bottom panel: Ω for different E . The slope of the fit (red line) provides $(d/I) = \omega_0\Omega/E = (106 \pm 2) \text{ A s/kg m}$, which is consistent with the measurement of d/I from the frequency of small oscillations.

in the absence of dissipation.

The angular momentum vector \vec{L} , undergoes left-handed precession about the electric field axis of rotation at an angular frequency $\Omega/2$. The factor of 1/2 is due to the fact that the torque is averaged over a rotation of the microsphere. This behavior is observed as a low-frequency modulation of the cross-polarized light most prominently at twice the precession frequency. The amplitude spectra for a precessing microsphere driven by a range of electric fields magnitudes are shown in Figure 5.3, in which the microsphere is spinning at $\omega_0 = 100\pi$ krad/s. The low-frequency amplitude modulations of the cross-polarized light are larger than those for rotation about the trapping beam axis because of the coupling with the feedback used to stabilize the microsphere at low frequencies. As expected, the precession frequency is proportional to the magnitude of the driving electric field and the slope of Ω vs E , which implies $(d/I) = (106 \pm 2)$ A s/kg m, is consistent with the value of $(d/I) = (106 \pm 2)$ A s/kg m from the measurement of libration.

5.5 Torque Noise and Torsional Drag

The previously presented measurements demonstrate that the dynamics of a rotating microsphere are indeed well-modeled by Equation 5.1, so that metrological applications can be easily understood, such as the implementation of a vacuum gauge. Spinning-rotor vacuum gauges generally operate in the molecular flow regime, where they make use of the proportionality between gas pressure and the torsional drag induced by the gas in order to measure the chamber pressure. From the fluctuation dissipation theorem, the torsional drag coefficient β_{rot} is proportional to the single-sided torque noise power spectral density S_N ,

$$S_N = 4k_B T \beta_{\text{rot}}, \quad (5.9)$$

where $k_B = 1.381 \times 10^{-23}$ m² kg s⁻² K⁻¹ is the Boltzmann constant and T is the temperature of the residual gas. S_N can also be computed by considering the force noise spectral density, $S_{F_{\parallel}}$, imparted parallel to a surface element of the microsphere by successive gas molecule collisions and adsorption, followed by thermalization and re-emission, the common way to treat diffusive scattering. $S_{F_{\parallel}}$ can be integrated across the surface of the microsphere to

compute S_N , which is then combined with Equation (5.9) to obtain the relation,

$$\beta_{\text{rot}} \times \frac{v_{\text{th}}}{P} = \pi r^4 \left(\frac{32}{9\pi} \right)^{1/2}, \quad (5.10)$$

where $v_{\text{th}} = \sqrt{k_B T / m_0}$ is the characteristic thermal velocity of residual gas molecules of mass m_0 at temperature T , P is the gas pressure, and r is the microsphere radius. This is identical to the results in References [123, 124], and allows an absolute measurement of the pressure P via β_{rot} . The right hand side of Equation (5.10) is a geometric factor that depends only on r , the radius of the microsphere itself. Thus, if the gas species (i.e. v_{th}) is known, a measurement of the torsional drag constant β_{rot} can be used to directly infer the gas pressure from Equation (5.10).

Two phenomena can be used to measure the torsional drag β_{rot} produced on a spinning microsphere. (1) When released from a driving field, the angular momentum of the microsphere will decay exponentially with a time constant inversely proportional to the torsional drag coefficient. This is the primary method used to measure pressure with a spinning-rotor gauge and depends only on knowledge of the spherical rotor's radius and moment of inertia [38]. (2) When driven with an electric field, the torque exerted on the microsphere is proportional to the sine of the angle between the rotating field and the microsphere's permanent dipole moment. Gas surrounding the microsphere produces a drag torque which determines the equilibrium angle between the driving field and the microsphere's dipole moment. This angle can be measured as a phase difference between the driving field and the power modulation produced in the cross-polarized light [38].

It is distinctly possible that translational damping on resonance could also be used to infer the vacuum pressure. However, optically trapped microspheres in the described apparatus exhibit instabilities at high vacuum, requiring active feedback in order to retain them for long periods [19, 24–27, 38, 60]. The active feedback is designed primarily with derivative gain to mimic the damping provided by residual gas in lower vacuum environments where the trap is stable. It is difficult to deconvolve this feedback with the damping provided by actual residual gas in the vicinity of the microsphere. At first order, the rotational degrees of freedom are unaffected by the stabilizing feedback.

5.6 Release from a Spinning Field

After initializing the microsphere with a definite angular momentum and phase, the conditions can be changed to observe different solutions of the equation of motion. The simplest solution occurs when the drive electric field initially rotating about the trapping beam axis is switched off, so that only the drag term, $-(\beta_{\text{rot}}/I)\vec{L}$, remains in Equation (5.1) and the initial angular momentum decays according to:

$$\vec{L}(t) = e^{-t/\tau}\vec{L}(0). \quad (5.11)$$

Here, the damping time τ , is related to the damping coefficient by $\tau = I/\beta_{\text{rot}}$. This decay is illustrated in Figure 5.4 at the base pressure of the vacuum system (2×10^{-6} mbar). For the first 1000 s, $\omega_{\text{ms}} \gtrsim 150$ krad/s, the drag torque dominates, and the data are well modeled by an exponential decay. The behavior beyond 1000 s can be attributed to an optical torque, $\vec{T}_{\text{opt}} \sim 10^{-23}$ N m.

The average optical torque on a birefringent particle is given approximately by

$$\vec{T}_{\text{opt}} \approx \frac{\mathcal{P}}{\omega_{\text{opt}}}(1 - \cos(kr\Delta n)\sin 2\phi), \quad (5.12)$$

where \mathcal{P} is the ~ 1 mW trapping beam power, $\omega_{\text{opt}} \sim 10^{15}$ rad/s is the optical frequency, $k \sim 2\pi/(1.0 \mu\text{m})$ is the wavenumber, $r = 2.35 \mu\text{m}$ is the radius of the microsphere, Δn is the birefringence, and ϕ is the degree of ellipticity [127, 129]. The terminal angular velocity of the microsphere can then be explained by a fluctuating ~ 100 ppm degree of ellipticity in the trapping beam.

From Equation (5.10), the pressure can be written in terms of a constant $\kappa = \kappa(T, r)$, the gas particle mass m_0 , and the torsional drag coefficient β_{rot} ,

$$P = \frac{\kappa}{\sqrt{m_0}}\beta_{\text{rot}}, \quad (5.13)$$

where $\kappa \equiv (1/r^4)\sqrt{9k_B T/32\pi}$ is determined by measuring the microsphere radius r and the experimental chamber temperature T . The microsphere radius is determined indirectly by measuring the mass of the microsphere via electrostatic co-levitation [19] (see Chapter 4) and assuming the same uniform density found in Reference [19], as the microspheres are derived from the same lot. The drag coefficient β_{rot} is measured by observing the spindown time

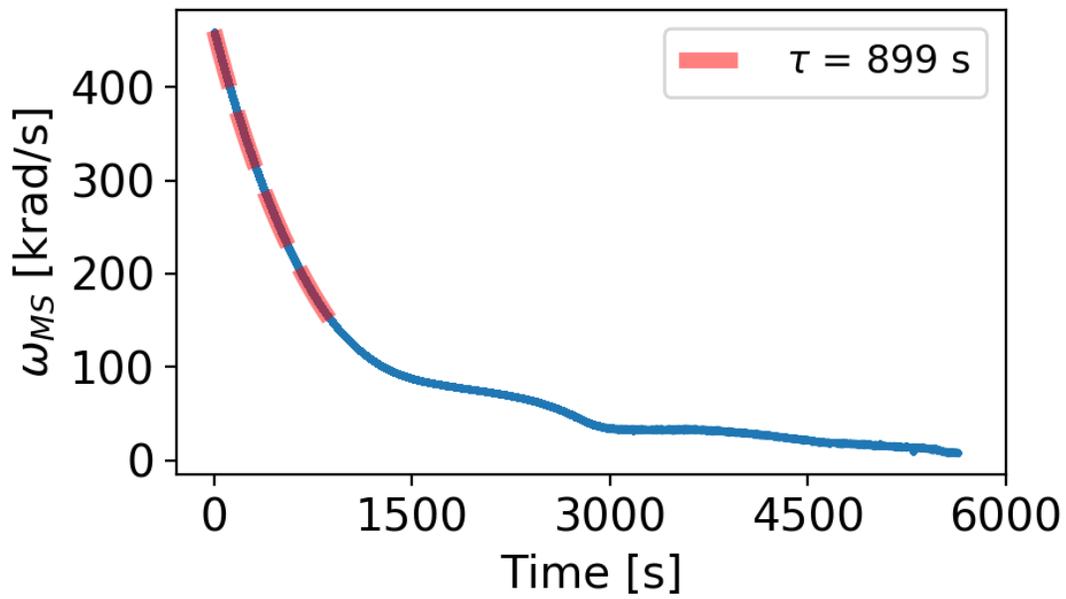


Figure 5.4: Time evolution of ω_{ms} after the driving electric field is switched off. For $\omega_{ms} \gtrsim 150$ krad/s the angular velocity exhibits an exponential decay. For $\omega_{ms} \lesssim 150$ krad/s the dynamics are modified by torque that could be explained by a ~ 100 ppm degree of ellipticity in the 1 mW trapping beam and the $\eta \sim 10^{-2}$ phase retardation of the microsphere.

of the microsphere released from a driving field. If torsional drag due to residual gas is the only torque present, then the angular velocity of the microsphere will decay exponentially with time constant $\tau = I/\beta_{\text{rot}}$, where I is the microsphere moment of inertia.

5.6.1 Spindown Pressure Measurement

As can be inferred from Figure 5.4, there is a roughly constant residual torque, likely optical in nature and generated from a small ellipticity in the trapping beam, which couples to the residual birefringence of the microsphere, as discussed in References [38, 65]. Under these conditions, the equation of motion for the microsphere angular momentum and the solution for the frequency of rotation are given by:

$$\frac{dL}{dt} = N_{\text{opt}} - \frac{\beta}{I}L(t), \quad (5.14)$$

$$\implies f(t) = f_0 e^{-(t-t_0)/\tau} + f_{\text{opt}}(1 - e^{-(t-t_0)/\tau}), \quad (5.15)$$

where $L(t)$ is the angular momentum of the microsphere, f_0 is the initial rotation frequency at time $t = t_0$, $\tau = I/\beta_{\text{rot}}$ is the decay time due to gas drag, and $f_{\text{opt}} = N_{\text{opt}}/(2\pi\beta_{\text{rot}})$ is the terminal rotation frequency of the microsphere due to a constant optical torque N_{opt} . The quantity f_{opt} can be measured by allowing the microsphere to reach steady state in the absence of any electrical driving torques and observing the terminal rotation frequency. Three different microspheres were used in this work, with generally consistent results. For microsphere No. 1 of those used for the spinning-rotor vacuum gauge demonstration, the terminal rotation frequency was measured to be $f_{\text{opt}} = (8315 \pm 559)$ Hz, where the uncertainty is the maximum observed deviation from the central value over approximately 6 hrs of successive measurements, where for each 2-s integration, the terminal rotation can be measured to within ± 1 Hz uncertainty. This type of systematic effect is similar to the “offset correction” necessary for magnetically levitated spinning-rotor gauges [121, 122, 125, 126].

To measure τ , the microsphere is released from an $f_0 = 110$ kHz driving field, and the frequency of power modulation of cross-polarized light produced by the rotation of the microsphere is observed. Successive 2-s integrations, separated by approximately one second, are first bandpass-filtered and then Hilbert-transformed to recover the instantaneous frequency of power modulation as a function of time. The median time of the i -th integration, $\langle t_i \rangle$, relative to the field turning off at t_0 , and the arithmetic mean of the instantaneous frequency measured during this integration, $\langle f_i \rangle$, are used to calculate a value for τ from

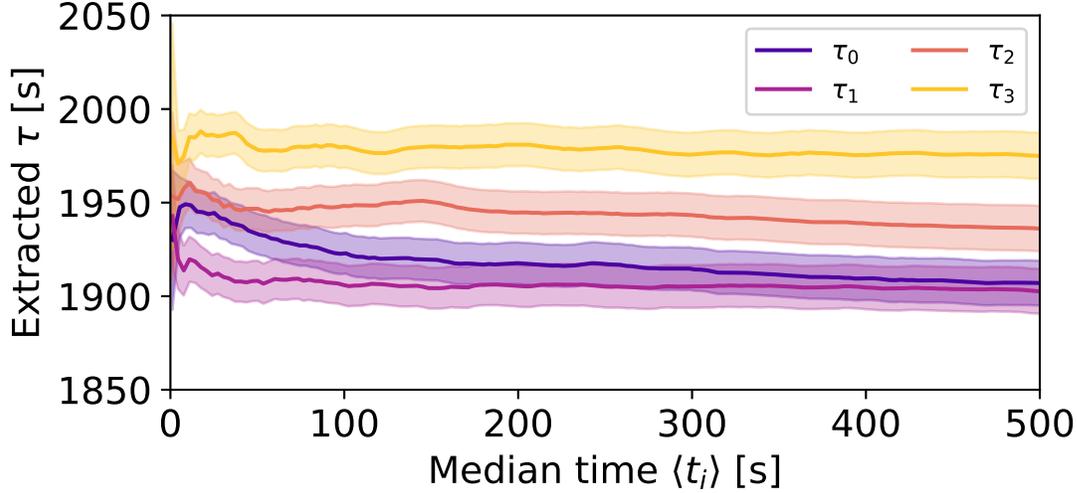


Figure 5.5: Damping time τ calculated from Equation (5.16) versus median integration time $\langle t_i \rangle$, for the four measurements with microsphere No. 1, indexed chronologically. The bands show the 1σ uncertainty propagated from the individual uncertainties of the values used to calculate τ .

Equation (5.15):

$$\tau_i = \frac{\langle t_i \rangle - t_0}{\log \left(\frac{f_0 - f_{\text{opt}}}{\langle f_i \rangle - f_{\text{opt}}} \right)}. \quad (5.16)$$

The calculated value of τ_i is plotted as a function of median integration time in Figure 5.5 for four example measurements performed with microsphere No. 1. The shaded bands represent the quadrature sum of systematic and statistical uncertainties, and are dominated by a systematic uncertainty in the assumed value of the terminal rotation frequency. For a single measurement, the estimation of τ from Equation (5.16) is self-consistent for approximately the first 500-s. For longer times, the effect of a slowly changing optical torque skews an estimation of τ by up to 10%.

The exponential decay of the rotation frequency for each of the four measurements with microsphere No. 1 are detailed in Figure 5.4, where data are shown for about 1 hour following the release of microsphere No. 1. Data in Figure 5.4 are overlaid with Equation (5.15), where the values of τ_j for measurements $j = 0, 1, 2, 3$ are the mean values of τ_i for $\langle t_i \rangle < 500$ s, and $f_{\text{opt}} = 8315$ Hz is fixed. The structure in the residuals, with an amplitude approximately 2% of f_0 , is likely the result of the slowly fluctuating optical torque.

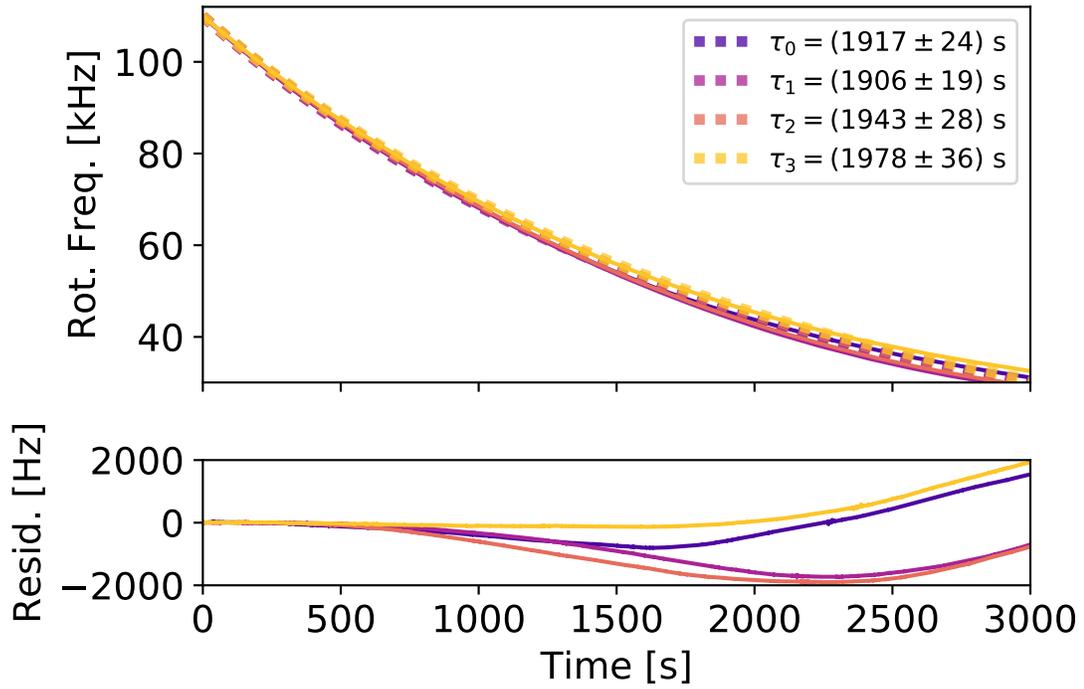


Figure 5.6: Exponential decay of microsphere No. 1 angular velocity due to torsional drag from residual gas after a 110 kHz driving field has been turned off. Data are shown for the same four measurements detailed in Figure 5.5. Dotted lines indicated realizations of Equation (5.15) with $\tau = \tau_j$ for measurements $j = 0, 1, 2, 3$, where τ_j is the average of the values plotted in Figure 5.5. The structure in the residuals, with amplitude $\sim 2\%$, is likely the result of a slowly fluctuating optical torque.

Table 5.1: Results of spindown pressure measurements, taken with microsphere No. 1. Statistical and systematic uncertainties are propagated independently from uncertainties on r , I , and τ .

	P [10^{-6} mbar]
P_0	3.53 ± 0.17 (stat.) ± 0.30 (sys.)
P_1	3.56 ± 0.17 (stat.) ± 0.31 (sys.)
P_2	3.49 ± 0.17 (stat.) ± 0.30 (sys.)
P_3	3.43 ± 0.16 (stat.) ± 0.29 (sys.)

The values of P computed with Equation (5.13) are shown in Table 5.1, assuming $m_0 = 18$ amu, as RGA analysis found the residual pressure to be dominated by water vapor. The four successive measurements, taken over the course of a few days with the same microsphere, are consistent within statistical uncertainties, and the systematic uncertainty should be common to all. Magnetically levitated spinning-rotor gauges typically have similar measurement uncertainties [120, 125, 130]. As a specific example, at a calibration pressure of 2.0×10^{-4} mbar, NIST calibration services for steel spinning-rotor gauges report a 0.02% relative uncertainty [130], i.e., $\sigma_P \approx 4 \times 10^{-8}$ mbar. This is within a factor of five of the uncertainty of our newly demonstrated silica microsphere spinning-rotor gauge (Table 5.1).

Importantly, these are measurements of the chamber pressure in the immediate vicinity of the microsphere, so that they can be used to estimate force and torque noise in precision measurement applications. The total pressure measured by the RGA is lower by approximately a factor of three, which could be attributed to a poor filament or response calibration, or may be indicative of a real pressure difference across the experimental chamber. With the current apparatus, there is a substantial pumping impedance between the trapping region and the RGA filament, as the trap is confined within the six electrodes which form a cage.

While the presence of anomalous sources of dissipation can't be excluded, two possible sources are shown to be too small to account for the pressure difference observed. The image charge induced on the 6 electrodes by the electric dipole moment would exert a torque roughly six orders of magnitudes smaller than the drag torque at the initial rotation velocity. Anomalous damping may also arise from electric field noise on the, nominally grounded, driving electrodes. Results in plasma physics suggest that fluctuating fields in the frequency band around the rotation frequency of the microsphere could also tend to increase

the angular velocity of the microsphere through stochastic acceleration [131], resulting in an apparent pressure that would be lower instead of higher.

To test the second phenomenon, the equation of motion given by Equation (5.14) was numerically integrated, including 100 distinct implementations of random torque fluctuations given by Equation (5.9), and distinct realizations of anomalous torque from the electric field noise produced by the driving electronics when the field is nominally off. As this anomalous torque is incoherent with the decaying angular momentum of the microsphere, its effect on the apparent damping time was found to be less than 0.1%, and can't explain the higher value of the measured pressure.

The apparent pressure difference may also be due to an elevated temperature of the microsphere, which is a notoriously difficult quantity to measure [36, 63, 132–134], often yielding only an upper bound. The derivation of Equation (5.10) assumes that the microsphere itself is in thermal equilibrium with the surrounding gas, so that incident and re-emitted gas particles have the same Maxwellian thermal velocity. If this assumption is removed, then the right-hand side of Equation (5.13) would have a multiplicative factor $[2T_{\text{gas}}/(T_{\text{gas}} + T_{\text{ms}})]$, with T_{gas} being the temperature of incident residual gas particles assumed to be in equilibrium with the chamber itself, and T_{ms} being the elevated temperature of the microsphere assumed to be in equilibrium with outgoing residual gas molecules. If the apparent pressure difference between the RGA and the location of the trap is due entirely to an elevated microsphere temperature, this implies $T_{\text{ms}} \sim 1500$ K.

5.7 Equilibrium Phase Lag

Once the dipole moment is known, it is possible to measure the drag due to residual gas and further verify models of dissipation for optically levitated particles. The equilibrium phase lag ϕ_{eq} between \vec{d} and \vec{E} is expected to increase with the pressure P , as the drag from the gas increases. In the molecular flow regime, the damping coefficient, β_{rot} , can be written as $\beta_{\text{rot}} = kP$, where $k = \sqrt{m_0}/\kappa$ is a constant that depends on the geometry of the microsphere, as well as the temperature and species of residual gas [123, 124]. The argument to the arcsin in Equation (5.5) can then be parameterized by

$$\frac{\omega_0 \beta_{\text{rot}}}{dE} = \frac{\omega_0 k}{dE} P = \frac{P}{P_{\text{max}}}, \quad (5.17)$$

where $P_{\max} \equiv dE/(\omega_0 k)$ is the pressure at which $\phi_{eq} \rightarrow -\pi/2$, where the microsphere rotation loses lock with the driving electric field (in practice fluctuations cause the microsphere to lose lock before ϕ_{eq} reaches $-\pi/2$).

This behavior is shown for one microsphere in the three top panels of Figure 5.7 for three different amplitudes of the driving field. It is evident that the unlocking pressure depends on the field amplitude and that after losing lock ϕ_{eq} becomes random. P_{\max} can be extracted from a fit for each field amplitude, as plotted (with additional values of the field) in the bottom panel of Figure 5.7. The linear relationship $k = dE/(\omega_0 P_{\max})$ indicates that the dissipation is proportional to the residual gas pressure and there are no significant additional sources of dissipation at pressures $\gtrsim 10^{-2}$ mbar.

The fit reports $k = (4.1 \pm 0.6) \times 10^{-25}$ m³s, assuming $d = 127 \pm 14$ e μm and a residual gas dominated by N₂. This is consistent with the value $k = 3.4 \times 10^{-25}$ m³s predicted in References [123, 124] for a $r = 2.35$ μm microsphere in thermal equilibrium with 300 K N₂ gas. No evidence for increased dissipation due to an elevated microsphere temperature or surface roughness is observed. This measurement of k can be used to infer the base vacuum pressure in the vicinity of the microsphere for the spin-down time measured in Figure 5.4 from $P = I/(k\tau)$. This give $P = (4.9 \pm 0.7) \times 10^{-6}$ mbar, which is roughly a factor of two greater than the pressure measured after the demonstration experiment with the cold cathode gauge. This discrepancy could be due to another source of dissipation which becomes significant at low pressures, inaccuracy of the cold cathode gauge, or a real pressure difference between the residual gas pressure in the trap and the cold cathode gauge at the end of the measurement.

With the technique in hand, we can consider further dedicated pressure measurements with other microspheres. This second method proposed to measure β_{rot} can be substantially faster, but in our apparatus has only sensitivity for moderate vacuum. Since this is also the regime in which absolutely calibrated capacitance manometers operate, this method is used for a cross-check of the technique. Consider a microsphere rotating at fixed frequency under the influence of a rotating electric field. In the rotating reference frame, the equation of motion for the phase ϕ between the driving field and the orientation of dipole moment has an equilibrium solution developed in Reference [38] and given by,

$$\phi_{\text{eq}} = -\arcsin\left(\frac{\beta_{\text{rot}}\omega_0}{Ed}\right) = -\arcsin\left(\frac{P}{P_{\max}}\right), \quad (5.18)$$

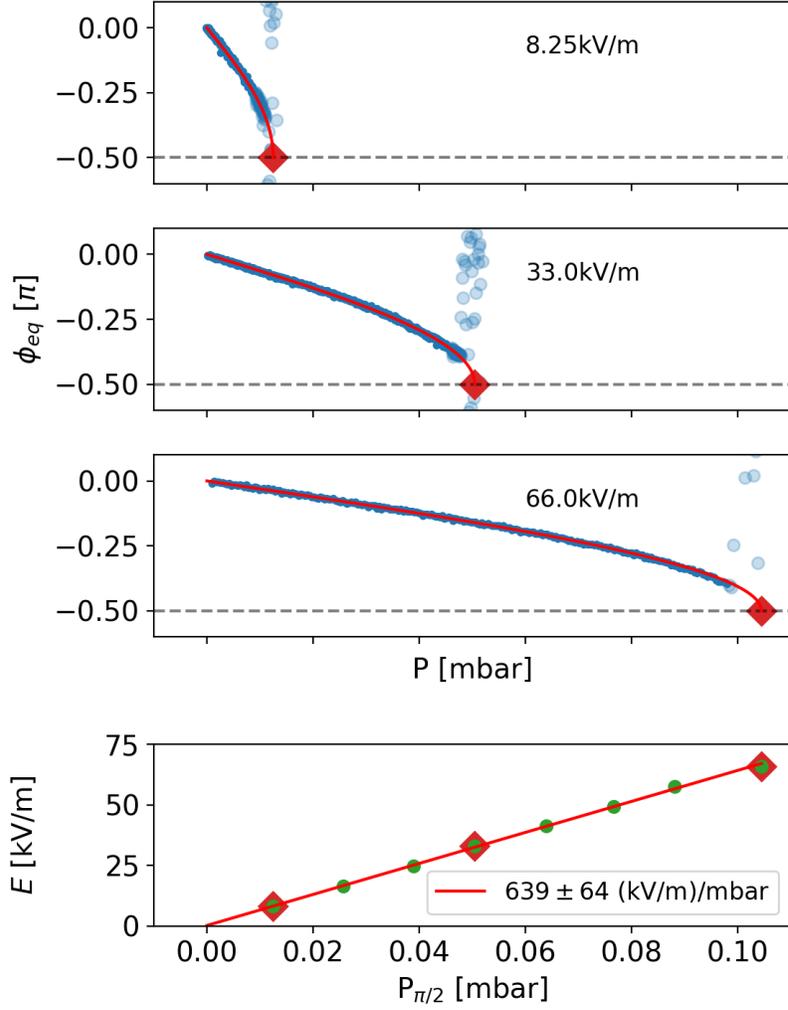


Figure 5.7: Top three panels: Equilibrium phase, ϕ_{eq} , versus chamber pressure for several magnitudes, E , of the driving field with $\omega_{ms} = 100\pi$ krad/s. For each value of E , ϕ_{eq} increases until the microsphere loses lock with the field, and the phase becomes random. For each E , a fit to Equation (5.5) (with the argument in Equation (5.17)) is shown in red. P_{max} is identified by a red diamond. Bottom panel: The linear relationship between P_{max} and E , with additional E included. The slope of the fit of E vs p_{max} is 639 ± 64 (kV/m)/mbar. Assuming the dipole moment measured from the frequency of small oscillations, this gives $k = \beta_{rot}/P = (4.1 \pm 0.6) \times 10^{-25}$ m³s, which is consistent with $k = \beta_{rot}/P = 3.4 \times 10^{-25}$ m³s predicted in References [123, 124].

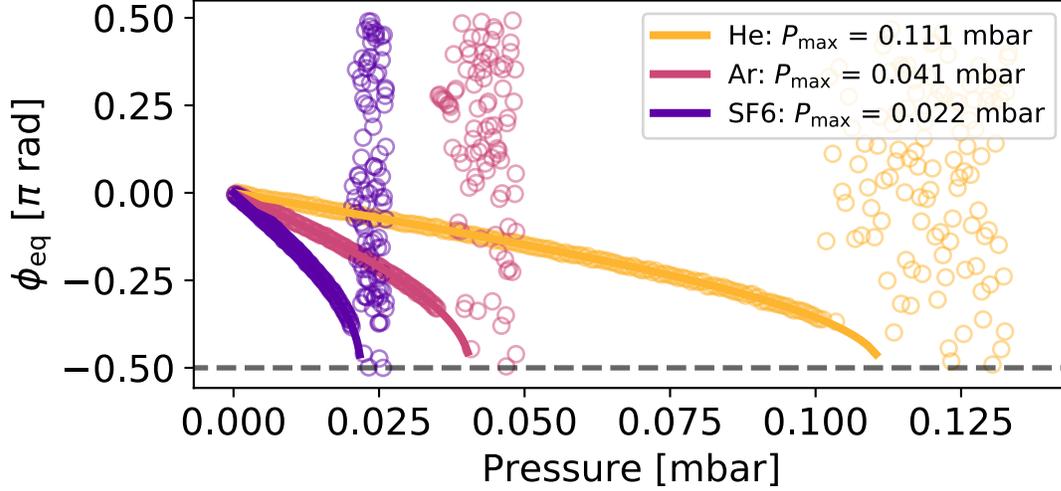


Figure 5.8: Equilibrium phase lag between the orientation of the microsphere dipole moment and a rotating electric field, versus the residual gas pressure determined by two capacitance manometers. Three distinct measurements are shown with He, Ar, and SF₆ residual gas. As the pressure is increased, the phase lags according to Equation (5.18), and P_{\max} is measured from a fit to this expression. As ϕ_{eq} approaches $-\pi/2$, the microsphere rotation becomes unlocked from the driving field (ϕ becomes random) at a pressure slightly below P_{\max} , due to librational motion of the microsphere and torque fluctuations from the increasing pressure of the surrounding gas.

where ω_0 is the angular driving frequency, E is the electric field magnitude, and d is the permanent electric dipole moment of the microsphere. The expression is written in terms of a single parameter P_{\max} , the maximum pressure that has a valid equilibrium solution given particular values of E , d , and ω_0 . Above this pressure, the driving field can no longer provide sufficient torque to maintain the microsphere's rotation. Using Equation (5.10),

$$P_{\max} = \frac{Ed}{\omega_0} \frac{P}{\beta_{\text{rot}}} \equiv \frac{Ed}{\omega_0} \frac{\kappa}{\sqrt{m_0}}, \quad (5.19)$$

where κ is already defined. The pre-factor (Ed/ω_0) is related to the driving torque and can be set by the experimenter, whereas κ is a constant across gas species, as the microsphere radius is unchanged throughout the measurements.

Thus, a gauge cross-check consists of measuring P_{\max} , which is inversely proportional to the damping constant, as a function of m_0 , the particle mass of a residual gas species, while

maintaining a constant driving torque and chamber temperature. P_{\max} can be measured by linearly ramping the pressure, monitored with a capacitance manometer, while continuously measuring ϕ_{eq} until unlocking, at which point the phase lag ϕ becomes random, as shown in Figure 5.7.

This effect is demonstrated for He, Ar, and SF₆ in Figure 5.8. Due to both the librational motion of the microsphere and torque noise from the increasing pressure of residual gas, the microsphere rotation becomes unlocked from the driving field at a pressure slightly below P_{\max} . Hence, P_{\max} is determined by extrapolating the endpoint of the arcsine relationship in Equation (5.18), as seen in Figure 5.8.

The measurement of ϕ_{eq} is used to validate Equation (5.10) and cross-check the microsphere as a spinning-rotor gauge against the capacitance manometer. By comparison, magnetically levitated spinning-rotors often make use of static expansion (wherein a defined amount of gas is allowed to fill a chamber of known volume) in order to provide cross-checks with ion gauges. Practical limitations of the current apparatus, primarily the bandwidth of the driving electronics, place a lower bound on the pressure observable with this method at $\sim 10^{-5}$ mbar, as at lower pressures ϕ_{eq} is too small to measure. Rotation velocities greater than the ~ 100 kHz achievable in the current system may allow for the extension of the method to lower pressure.

Silicon nanorods in counter-propagating beam traps have also been used to measure equilibrium phase lag induced by gas drag, but with the rotation driven optically by modulating the polarization (from linear to circular) of the trapping beams at fixed frequency [104]. Phase lag between the optical drive and the induced rotation of the silicon rotor was demonstrated to accurately predict gas pressures in the range 4 – 10 mbar.

5.8 Gas Composition and Changes in Dipole Moment

In many vacuum chambers, the ultimate base pressure achieved is limited by a single residual gas species, such as water vapor or hydrogen, and Equation (5.19) is directly applicable. In this work, requiring the deliberate introduction of various gas species, a small correction needs to be applied to account for contamination, primarily by water. When there are multiple species, each contributes independently to the total torsional drag. Appealing to the analysis in References [123, 124], the torque noise spectra of multiple gases can be added,

each with partial pressure $P_i = \chi_i P_{\text{tot}}$ and particle mass $m_{0,i}$,

$$S_{N,\text{tot}} = 4k_B T \beta_{\text{rot}} = 4k_B T \left(\frac{P_{\text{tot}}}{\kappa} \sqrt{m_{0,\text{eff}}} \right), \quad (5.20)$$

where the expression has been written in terms of an effective mass, $m_{0,\text{eff}} = (\sum_i \chi_i \sqrt{m_{0,i}})^2$. The mole fractions, χ_i , can be measured with an RGA, appropriately accounting for differences in ionization probability for different gas species. For He, N₂, Ar, and SF₆, the preparation of the manifold results in an apparent purity of $\gtrsim 99.9\%$, limited by systematic uncertainties in RGA ionization probabilities for different gas species. In the cases of Kr and Xe, the apparent purity is $\sim 99\%$, limited by water contamination in the manifold.

The process of characterizing the residual gas leaked into the experimental chamber makes use of an RGA on the experimental chamber itself, which tends to charge the trapped microsphere with an excess of electrons. The microsphere is returned to zero charge, prior to spinning up and performing the drag measurements, by repeated exposure to ultraviolet photons from an Xe flashlamp [24, 25]. The process of charging and discharging appears to change the microsphere electric dipole moment.

In order to measure κ as defined in Equation (5.19), the term (Ed/ω_0) must be measured precisely. The quantities E and ω_0 are controlled and measured precisely, and via the relation between libration frequency and electric field amplitude, it is possible to characterize the microsphere permanent dipole moment d . This is done by analyzing the librational motion about the instantaneous direction of the electric field, together with a calculation of moment of inertia from the measured mass of the microsphere.

Over the course of all measurements presented here, the dipole moment assumed values between approximately 95 and 120 $e\mu\text{m}$ for microsphere No. 1, between 30 and 55 $e\mu\text{m}$ for microsphere No. 2, and between 80 and 110 $e\mu\text{m}$ for microsphere No. 3. Without the RGA filament on, the dipole moment was measured to be constant, within measurement uncertainties, over the course of a day.

5.9 Capacitance Manometer Cross-Check

The microsphere spinning-rotor vacuum gauge is cross-checked via the following procedure: first, the gas manifold is prepared with a particular species and $m_{0,\text{eff}}$ is characterized by leaking a sample of the gas into the experimental chamber where the RGA is present; second,

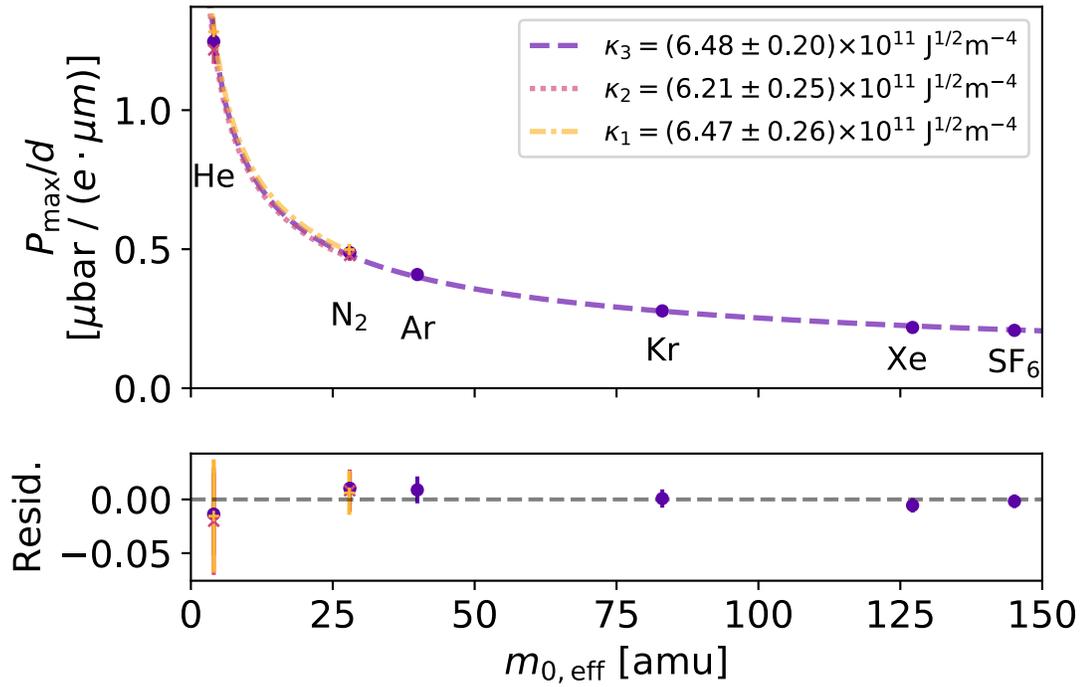


Figure 5.9: The quantity P_{\max}/d is plotted against the effective mass of the gas $m_{0,\text{eff}}$ for different species spanning the 4 to 150 amu range, for three different microspheres. With E and ω_0 known, the constant κ is extracted by fitting Equation (5.19) to the data, with a single free parameter, as shown for each microsphere. A χ^2 -minimization finds $\chi_{1,\text{min}}^2/N_{\text{DOF}} = 1.9/1$, $\chi_{2,\text{min}}^2/N_{\text{DOF}} = 3.5/1$, $\chi_{3,\text{min}}^2/N_{\text{DOF}} = 8.5/5$. Residuals shown below are plotted with the same units as the data.

the chamber is pumped to its base pressure, the microsphere is returned to a neutral state, and its permanent electric dipole moment d is determined by analyzing the libration. Finally, the gas prepared is leaked into the experimental chamber, and the equilibrium phase lag ϕ_{eq} is measured as a function of gas pressure to determine P_{max} . The RGA filament is turned off prior to the second and third steps.

For microsphere No. 1 and microsphere No. 2, this procedure was performed three times for He and N₂, while for microsphere No. 3 the measurement was performed with He, N₂, Ar, Kr, Xe, and SF₆. Successive measurements with each microsphere and gas species are found to agree and are averaged together by considering the quantity (P_{max}/d) , to account for the small differences in dipole moment between measurements. The results of all pressure ramp measurements are shown in Figure 5.9. A χ^2 -minimization with a single parameter is used to fit Equation (5.19) to the data and extract κ for each microsphere. The data are well modeled by Equation (5.19), demonstrating the validity of Equation (5.10) and the analysis in References [123, 124].

The values of κ determined with these analyses are shown in Table 5.2, together with the individual microsphere radii and the expected value of κ , which have both been computed from the known value of the density [19] and the measured value of the microsphere mass, assuming the microsphere is thermal equilibrium with the gas. The consistency between the measured and calculated values suggest that in moderate vacuum, $P \approx 10^{-3} - 10^{-1}$ mbar, the microsphere is indeed in thermal equilibrium with the gas. The uncertainty in the directly measured value of κ is dominated by systematic uncertainties in the measured dipole moment, which, in turn, depends on the uncertainties in the moment of inertia, another derived quantity.

Previous work with magnetically levitated spinning-rotor vacuum gauges usually include a momentum accommodation coefficient relating specular and diffuse reflection of gas particles, and which encompasses rotor geometry and surface composition. Extensive measurements have found this factor to be consistent with unity for steel rotors, with percent-level precision [120–122, 125, 126, 135]. In this work, the accommodation coefficient would appear as a multiplicative constant σ on the right-hand side of Equation (5.10). Comparing the measured and calculated values of κ , the accommodation coefficients σ_i for the i -th microsphere were found to be consistent with unity: $\sigma_1 = 0.98 \pm 0.04$ (stat.) ± 0.08 (sys.), $\sigma_2 = 1.02 \pm 0.05$ (stat.) ± 0.08 (sys.), and $\sigma_3 = 0.96 \pm 0.04$ (stat.) ± 0.07 (sys.). Prior work

Table 5.2: Calculated and measured calibration factors κ for the three microspheres used here, together with their measured masses, from which the values of r and then κ are determined. The microsphere density was assumed to be $\rho_{\text{ms}} = 1.55 \pm 0.05$ (stat.) ± 0.08 (sys.) from Reference [19]. The relatively large uncertainty in ρ_{ms} limits the precision with which κ can be calculated, whereas the precision of the measurement of κ is limited by uncertainties in the measured dipole moment and P_{max} .

	m_{ms} [pg]	κ [$10^{11} \text{ J}^{1/2} \text{ m}^{-4}$]	
		Theory	Experiment
1	84.3 ± 0.2 (stat.) ± 1.5 (sys.)	6.3 ± 0.3 (stat.) ± 0.5 (sys.)	6.47 ± 0.06 (stat.) ± 0.25 (sys.)
2	84.2 ± 0.4 (stat.) ± 1.4 (sys.)	6.3 ± 0.3 (stat.) ± 0.5 (sys.)	6.21 ± 0.06 (stat.) ± 0.24 (sys.)
3	85.0 ± 0.6 (stat.) ± 1.5 (sys.)	6.2 ± 0.3 (stat.) ± 0.4 (sys.)	6.48 ± 0.09 (stat.) ± 0.17 (sys.)

with a single levitated silica nanoparticle, with diameter $d \approx 70$ nm, has found an accommodation coefficient of $\sigma = 0.65 \pm 0.08$ [134] based on an analysis of heating rates. The discrepancy could be the result of the vastly different scale of the two types of rotors, poorly understood material properties, as well as a tendency for nanoscale particles to be highly non-spherical (see scanning electron microscopy images in References [17, 52, 55, 59, 136, 137]).

The consistency between the measured and computed values of κ is also an indirect validation of the work presented in Reference [19] which makes no assumptions about microsphere temperature or thermal equilibrium and demonstrates independence across vacuum pressures $P \approx 10^{-6} - 10^0$ mbar. If the density computed there was incorrect, quantities derived from this density, such as κ , would be inconsistent with independent measurements of those quantities.

5.10 Non-ionizing Gas Analyzer

If the rotational dynamics of a microsphere with a known value of κ are analyzed while the pressure is measured with a calibrated, species-independent vacuum gauge such as a capacitance manometer, the combination can be used as a non-ionizing gas analyzer, operable directly in moderate vacuum. In particular, this system would excel as a binary gas analyzer, comparing the concentrations of two gases, such as one might encounter in nanofabrication with dopants in a carrier gas, or in chemistry with a particular stoichiometric ratio of reagent gases.

From Equation (5.13), the mass of effective residual gas particle can be calculated,

$$m_{0,\text{eff}} = \left(\frac{\kappa\beta_{\text{rot}}}{P} \right)^2 = \left(\sum_i \chi_i \sqrt{m_{0,i}} \right)^2, \quad (5.21)$$

where β_{rot} can be determined from a spindown measurement, $\beta_{\text{rot}} = I/\tau$, or from an equilibrium phase lag measurement, $\beta_{\text{rot}} = Ed \sin(-\phi_{\text{eq}})/\omega_0$, where ϕ_{eq} is strictly negative. The latter quantity can be measured continuously. Such a system has an immediate advantage over ionizing RGAs, as it can operate directly in the pressure regime $P = 10^{-5} - 10^0$ mbar, whereas an ionizing RGA requires $P < 10^{-4}$ mbar and is often connected to a vacuum chamber of interest by a leak valve, which can itself change the relative concentrations of gases. Indeed, the total absence of ionization offers its own advantage.

5.11 Conclusion

We have demonstrated electrically driven rotation of optically levitated particles. Although it would be technically challenging to reach the ~ 1 GHz rotation frequencies achieved with optical rotation [103], electrically driven rotation offers precise control over the direction and magnitude of the torque, making the quantitative measurements presented here possible. This is in contrast to optically driven rotation, where the coupling to the optical torque depends on the geometry and birefringence of the trapped particle, as well as the exact optical properties of the laser beam used to apply the torque. The ability to measure and control system parameters has implications for the field of optical levitation and manipulation.

The libration of an optically levitated dipole trapped in an electric field is a degree of freedom that can be exploited for cooling and precision measurement. Libration of a dipole trapped in a rotating electric field can be cooled by simple phase modulation of the electric field driving the rotation. microspheres made from materials with larger dipole moments could also be used to couple to and measure oscillating electric fields near the rotation frequency. After the dipole coupling to an electric field is measured from the libration frequency, torques measured from the phase lag of the rotation can be calibrated into physical units.

Background forces arising from electric fields coupling to a microsphere dipole moment have been the limiting factor in several measurements with optically levitated microspheres [24, 25]. The techniques presented here provide new tools to control and measure

the orientation of a microsphere's dipole moment, which could be used to mitigate these background forces by averaging over a rotation of the dipole. This would not be possible with optical rotation because the dipole is not guaranteed to be orthogonal to the axis of rotation.

We have further demonstrated the use of an optically levitated silica microsphere as a microscopic spinning-rotor vacuum gauge. The gauge operates on the principle that the torsional drag is proportional to the gas pressure in the vicinity of the microsphere. A spinning microsphere driven by a rotating electric field is released from the driving field, and the subsequent decay of angular momentum induced by torsional gas drag is observed. The residual gas pressure is then inferred from the decay time and the calculated moment of inertia.

Once a microsphere has been calibrated by measuring its mass, pressures within the range $10^{-6} - 10^{-3}$ mbar can be determined within 10 s with a precision of $\lesssim 7\%$, under the given assumptions and limited by the uncertainty in the calibration κ , which, in turn, depends on uncertainty in the assumed microsphere density. The minimum measurable pressure could easily be reduced by using a faster initial rotation velocity, which is limited only by the bandwidth of the driving electronics in the current apparatus.

A second method measures the equilibrium phase lag of the electric dipole moment relative to a driving field, which is induced by the gas drag. This method is used to cross-check the microsphere spinning-rotor gauge against a capacitance manometer. The same bandwidth limitations in the driving electronics limit the minimum measurable pressure of this method to $\gtrsim 10^{-5}$ mbar.

A geometric calibration factor κ related to the microsphere radius is independently determined by measuring the torsional drag as a function of gas species, since the drag depends on the momentum imparted by gas particle collisions and, thus, the molecular mass. The torsional drag β_{rot} has been shown to scale inversely to the thermal velocity and, thus, $\beta_{\text{rot}} \propto m_0^{1/2}$, with m_0 being the molecular mass, validating the rotational dynamics presented here and in References [123, 124], and opening the possibility to measure the effective molecular mass of a mixture of gases without ionization and directly in moderate vacuum.

Chapter 6

Beyond the Standard Model: An Introduction

As an experimentalist, the vast majority of one's time is occupied with innumerable hardware tasks, often followed by periods of data analysis that eventually inform the publication of manuscripts and/or further hardware tasks. It is thus often the case that experimentalists have only effective understandings of the theoretical framework under which their investigations take place.

It would thus be a disservice to both reader and author to drag on at length about the details of the “new physics” that might be explored with the use of optically levitated microspheres, but properly contextualizing the broader impact of one's research is critical. As such, a brief analysis is presented for a select number of theories that may yield measurable, but as-of-yet undetected, new interactions. In particular, these theories would result in short-distance modifications to gravity, or other “fifth forces”, which may be accessible with optically levitated microspheres and other precision force sensors.

6.1 The Problem With Gravity

Among fundamental interactions, gravity has the distinction of simultaneously being the most apparent and yet the least understood. From the theoretical point of view, the universal law of gravitation [138] and General Relativity [139] have been quite successful in describing interactions at macroscopic scales. Under this paradigm, the force of gravity is understood as a classical consequence of curved spacetime, which itself is induced by the presence of

matter and energy, and nowhere requires a quantization or an associated field. This is in stark contrast to the other fundamental interactions, such as electromagnetism, where our current understanding is built on quantum field theory and the Standard Model of particle physics. It should also be noted that the gravitational constant G is not related to any other fundamental constants.

General Relativity, as compared to the Standard Model, suffers from limited and rudimentary empirical knowledge at submillimeter length scales [140, 141], primarily due to the strength of the gravitational interaction relative to residual electromagnetic interactions in experimental apparatuses. At the same time, robust connections between gravitation and quantum mechanics are still obscure, as they have remained since Einstein’s initial attempts to develop a grand unified theory of physics. Yet, much of theoretical physics in the past century has been driven by the assumption that gravity and General Relativity remain unmodified all the way down to the Planck scale of $l_P \approx \sqrt{\frac{G\hbar}{c^3}} \sim 10^{-35}$ m, where G is the gravitational constant, \hbar is the reduced Planck constant, and c is the speed of light.¹ This assumption is necessarily related to the hierarchy problem, given that Planck scale is many orders of magnitude is so vastly different than the electroweak scale.

From an experimentalist’s perspective, this contention is not *a priori* problematic, but presents an interesting avenue for research, regardless of model-dependent motivations. Simply put, if it is assumed that gravity and General Relativity work at both the shortest and the longest scales, it stands to reason that they should be tested at both the shortest and longest scales. For non-relativistic experimental apparatuses (the majority of those on Earth) this generically motivates characterizations of the classical inverse square law of Newtonian gravity, $F_{\text{grav}} \sim 1/r^2$. Any measured deviations from this expectation would surely generate a slew of theoretical models.

Furthermore, the reigning model that describes our universe at the largest scales, Λ CDM Cosmology (see Appendix F), implies that if the equations of General Relativity are assumed and our observations of the universe around us can be trusted, then there are two types of constituents in the universe that we’ve never observed or detected directly, dark matter and dark energy. We only presume their existence given that their role in Λ CDM cosmology

¹It is uncertain whether the Planck scale (length: $l_P = \sqrt{\hbar G/c^3}$, mass: $M_P = \sqrt{\hbar c/G}$, time: $t_P = \sqrt{\hbar G/c^5}$, energy: $E_P = M_P c^2 \approx 10^{19}$ GeV) is actually a fundamental scale of the universe, like the electroweak symmetry breaking scale measured as ~ 250 GeV [142–145], but it does represent the scale at which gravitational interactions should become comparable in strength to electroweak and strong interactions.

is consistent with the observed effects on normal matter and spacetime [146, 147]. Dark energy is particularly mysterious in that it tends to induce expansion of spacetime, whereas normal matter and dark matter both induce contraction, and it appears to take the form of a “cosmological constant” term in Einstein’s field equations. Assuming Λ CDM, one can calculate the expected energy density associated with the cosmological constant, often denoted Λ , which is constrained by observations. With the energy density in hand, it is possible to infer an associated length scale (again, see Appendix F), which is given by,

$$l_{\Lambda} \sim 100 \mu\text{m}. \quad (6.1)$$

It is presumed that above this length scale, any underlying quantum field that gives rise to Λ would be integrated out, and the effective description using a constant vacuum energy density should be sufficient. Below the length scale l_{Λ} , however, dynamics of a quantum field of dark energy may become apparent [140]. Given that Λ is fundamentally tied to a description of the universe making use of General Relativity (and thus gravity), it is natural to consider possible new interactions as modifying gravity.

Based solely on the length-scale, optically levitated microspheres may be useful in this regime, as lasers have wavelengths in the range $0.1 \mu\text{m}$ to $10 \mu\text{m}$, and thus diffraction-limited spot sizes of the same order. Particles that are trappable by such lasers range in size from approximately $0.05 \mu\text{m}$ to $20 \mu\text{m}$ [63]. Given that the optical fields can be so confined and the test masses (i.e. levitated dielectric microspheres) are sufficiently small, only an appropriately sized source mass needs devising. This will be discussed in detail in Chapters 7 and 8, once we understand the possible observables and a few underlying models.

6.2 Parameterizing New Forces

Generally speaking, it’s assumed that the field or fields responsible for dark energy will couple directly to matter (much like gravity), either via the electron, the up/down/strange quarks, or the gluon field itself. This is one of the underlying motivations for probing short-distance forces, wherein two objects composed of normal matter are separated by a distance $d \leq l_{\Lambda}$, and any resulting force between them is observed and characterized. The size and scaling of this force might then depend on the details of the specific theory hoping to explain the interaction. Short-distance deviations from Newtonian gravity may also arise from modifications to General Relativity and/or the structure of spacetime, as well as general

particle exchange that may have no bearing on gravity itself.

Table-top experiments to probe short-distance forces, such as the one described in this work, are usually far into the non-relativistic regime, allowing us to work with classical potentials and simple mechanics. In this regime, it is customary to modify the inverse square law of Newtonian gravity by introducing an additional Yukawa term with a length scale λ . The resulting potential between two point masses can then be written as:

$$V(r) = -G \frac{m_1 m_2}{r} (1 + \alpha e^{-r/\lambda}), \quad (6.2)$$

with G the Newtonian constant of gravitation, m_1 and m_2 the gravitating masses, r their distance, and α the strength of the new interaction relative to gravity. This expression for a modified gravitational potential is simply a particular parameterization of possible new interactions. If an experiment were to detect a new interaction with some value of α and λ , further investigations will be required to verify the specific form of the new interaction, which may, for example, be a power law of r , rather than an exponential term. This form is chosen because a Yukawa potential is naïvely expected if there is massive particle exchange that mediates the new interaction, and α can be either positive or negative, depending on the theory that gives rise to the modification.

One might ask why new interactions are parameterized as modifications to the gravitational potential. A simple argument is that the other fundamental interactions aside from gravity are well described by quantum field theories, and have been experimentally verified to extremely high precision (see References [148–150] for tests of Coulomb’s law). If it is further recognized that laboratory experiments not intentionally involving radioactive sources or nuclear processes will be subject only to gravity, electromagnetism, and possible new interactions, it should be noted that electromagnetism is many orders of magnitude larger than gravity for reasonable physical parameters.² If the new interaction is expected to be small (because if it wasn’t, it would have already been seen), then modifying the classical potential of the weakest interaction makes more sense.

An interesting consequence of this parameterization is that experiments will necessarily

²This is the usual argument presented in undergraduate physics: consider two spheres with equal mass and opposite charge, separated by some distance r . For relevancy’s sake, let $m \sim 100$ pg like our silica microspheres (see Chapter 4) and let $|q| = 1e$. Then we can compare the expected gravitational and electrostatic forces between the two as $F_{\text{grav}}/F_e \sim Gm_1m_2/(kq_1q_2) \sim 10^{-9}$. Clearly, gravitational forces are much smaller than electromagnetic forces for reasonable parameters one might encounter in this and other works.

probe a limited range of the parameter λ . Assume that in any given experiment, r is fixed to a certain value, or a limited range of values. If $\lambda \ll r$, the interaction is heavily suppressed, while if $\lambda \gg r$, the exponential term approaches unity and the interaction amounts to a re-scaling of the gravitational constant $G \rightarrow G(1 + \alpha)$. Given that other parameterizations are possible, experiments with $\lambda \sim r$ are advantageous, as they can more easily distinguish between different underlying models, and are especially sensitive to interactions and parameterizations that scale with large powers of the separation r .

6.3 A Few Examples

One of the seemingly most intuitive explanations of the cosmological constant is the existence of a new quantum field that has not yet been directly detected [140]. New quantum fields might also be the constituents of dark matter [151–153], or may explain the strong-CP problem [154]. Beyond the supposition of new particles and their associated quanta, some Beyond the Standard Model theories attempt to unify gravity and the other fundamental forces, therefore solving the hierarchy problem, by introducing new dimensions to spacetime, where our 4-dimensional space is embedded within the higher-dimensional space [141, 155–158]. It may even be possible that the graviton itself (assuming a quantized gravitational field exists), has some physical size/extent, much different than the point-like nature of Standard Model particles [159].

In nearly all of these scenarios, there arise short-distance modifications to Newtonian gravity, which provide a testable framework for laboratory experiments. The following analyses closely follow References [141, 160] which describe this material in more detail and have extensive bibliographies. There are likely additional models that can yield modifications to Newtonian gravity and this should serve only as a selective sampling.

6.3.1 Bosonic Particle Exchange

Consider a scalar field ϕ which is assumed to couple to matter [161]. Ultra-low-mass scalar and vector bosons may be good dark matter candidates [151, 152], providing general motivation to consider boson-mediated interactions. If ϕ is to be incorporated into the Standard Model, one can write down terms that might appear in the Lagrangian as,

$$\frac{m_e}{\kappa} \phi \bar{f} f, \quad \text{or} \quad \frac{1}{\kappa} \phi G_{\mu\nu}^a G^{a,\mu\nu}, \quad (6.3)$$

where f is a fermionic field (like the electron, up/down/strange quarks), $G_{\mu\nu}^a$ is a gluon field with a indexing gluon color charges, κ is a coupling constant with dimensions of mass, and the chiral-symmetry-breaking operators include the fermion mass, as one might expect from the Standard Model [141]. If this scalar has a mass m_ϕ then the associated strength, α , and length-scale, λ , of the Yukawa modification from Equation (6.2) are given by,

$$\alpha = \frac{M_P^2}{4\pi\kappa^2} \quad \text{and} \quad \lambda = \frac{\hbar}{m_\phi c}, \quad (6.4)$$

where $M_P = \sqrt{\hbar c/G}$ is the Planck mass, and λ is simply the Compton wavelength. As a specific case, consider a coupling to the gluon field, and recall that in the renormalization of the Standard Model with an ultra-violet cutoff scale Λ_{UV} , a scalar field with an interaction term as in Equation (6.3) would expect radiative corrections of order $\delta m_\phi \approx \Lambda_{UV}/4\pi\kappa$. Imposing the condition of naturalness, one expects $\delta m_\phi \lesssim m_\phi$. If it is further assumed this field is related to the dark energy length-scale, $\lambda \sim l_\Lambda$ so that $m_\phi \sim 2 \text{ meV}/c^2$, and take $\kappa \sim m_N$ with m_N the mass of a nucleon (given that the mass of nucleons are largely a result of the gluon field strength and not the masses of the component quarks) this naturalness condition implies,

$$\Lambda_{UV} \lesssim 10 \text{ TeV}, \quad (6.5)$$

above the electroweak scale of $\sim 250 \text{ GeV}$ [142–145] where one might reasonably expect new physics. Particle collider experiments, particularly those at the LHC [162, 163], are accessing these energy scales and provide alternative and complementary techniques for any detection of new physics.

The new physics at this energy scale may be the result of supersymmetry, wherein all the Standard Model particles known have symmetric partners at different energy scales. This sector of extra particles couples to the Standard Model only via gravity and other couplings of similar strength. The supersymmetric partner of the graviton may be related to the cosmological constant, and have an energy scale of $\sim 1 \text{ meV}$ [164]. Other scalar particles of the supersymmetric sector may exhibit gravitational strength interactions with Standard Model particles, motivating investigations of the inverse square law.

Observations of the Cosmic Microwave Background (CMB) [147] and distant supernovas [165] suggest a cosmological model that doesn't allow for the implied relic abundance of a simple scalar field within the standard, but a dynamical scalar field that evolves with

the universe may avoid this limitation [141, 166, 167]. Scalar fields that are manifest products of other physics (such as the QCD axion), or scalar interactions from a supersymmetric sector, are not subject to the same cosmological constraints [154, 160].

6.3.2 Chameleons

Some scalar fields have been postulated to exhibit screening mechanisms, allowing them to avoid detection in typical laboratory experiments or astronomical observations [168]. A specific example is the chameleon mechanism [169, 170], in which the effective mass of the chameleon particle (corresponding to the inverse length scale of the interaction) depends on the local matter density, i.e. there is back-action from the metric. At cosmological distances where the matter density is low, the chameleon field would exhibit the requisite vacuum energy to explain the accelerating expansion of the universe [171], while at shorter distances, the dynamics of the field and its coupling to matter can exert gravitational strength forces that may modify the inverse square law of Newtonian gravity.

Following the initial proposal of the chameleon mechanism [169, 170], consider a field ϕ and construct an effective potential for the field, written in natural units with $\hbar = c = 1$, in the presence of matter density ρ ,

$$\begin{aligned} V_{\text{eff}}(\phi) &= \Lambda^4 e^{(\Lambda/\phi)^n} + \rho e^{\beta\phi/M_P} \\ &\approx \Lambda^4 \left(1 + \frac{\Lambda^n}{\phi^n} \right) + \rho \frac{\beta\phi}{M_P}, \end{aligned} \tag{6.6}$$

where β is the coupling to matter (assumed non-relativistic and averaged over Standard Model particle types), n is an arbitrary power, and where this potential has been explicitly constructed assuming that chameleons are responsible for dark energy and thus includes Λ , the cosmological constant energy scale, as the strength of the self interactions. In reality, the self-interaction energy is a parameter of the theory, but this notation is used here to be consistent with Chapter 7 and Reference [25]. The first exponential has been expanded and only one arbitrary nonlinear term ($1/\phi^n$) has been retained follow the logic in References [166, 168, 172]. The second exponential has been expanded since the exponent is scaled by $1/M_P$ and thus assumed small for any natural value of β .

A wave equation that governs this field can be derived since $\nabla^2\phi = \partial V_{\text{eff}}/\partial\phi$, and thus with a given matter density, ρ , the shape and strength of the chameleon field is calculable. With the specific profile of a chameleon field for a given experimental apparatus, one can

then compute expected the expected force on a test mass within that field. Following the analyses in References [172, 173], the force on a homogeneous test mass of density ρ_0 can be calculated as,

$$\vec{F}_\phi = \xi \frac{\beta \rho_0}{M_P} \int_V (\nabla \phi) dV, \quad (6.7)$$

where V is the volume of the test mass and $\xi \leq 1$ is a screening factor, given that the test mass itself affects the strength of the chameleon field [173]. As noted in References [160, 173–175], a chameleon-mediated interaction would look like a Yukawa-type modification to Newtonian gravity over certain ranges of parameters and apparatus configurations. If any new interactions were detected, special care would have to be taken to distinguish between possible underlying theories

Previous experiments searching for violations of Newtonian gravity at or below the dark energy length scale often employed macroscopic test masses or a conductive shield between the probe and test masses to minimize electromagnetic backgrounds, both of which necessarily limit the scale and possible effects of the chameleon field. Experiments performed in vacuum, with no intervening matter between source and test masses, may allow detection of a field that would otherwise be screened, where the chameleon field is just one particular example.

6.3.3 String Theory and a Multitude of Bosons

By their very construction, string theories often result in a variety of scalar fields that ultimately couple to matter. Some string theories suggest that the strength of the gravitational interaction $G \sim 1/M_P^2$ is not actually a constant, but instead a dynamical scalar field. In this context the scalar field is referred to as a dilaton [176–178], and it is often postulated to evolve with the expanding universe in order to reproduce the currently observed cosmology. As before, modifications to the inverse square law can be expected at distances shorter than the Compton wavelength of the scalar field.

More generally, super-symmetric string theories usually invoke the existence of multiple, massless scalar fields called “moduli” [141, 179] that determine the coupling strength and other parameters of the theory. However, in order for the theory to reproduce Newtonian gravity in the appropriate limit, the moduli need to acquire mass. This is usually accomplished by assuming some supersymmetry-breaking mechanism, similar to the manner in

Table 6.1: Example strengths and length scales for a set of moduli (i.e. scalar fields) from a specific supersymmetric string theory, and with natural assumptions made about symmetry-breaking energy scales. Reproduced from Reference [179].

Scalar field	α	λ [m]
ϕ_u	2×10^{-2}	2×10^{-2}
ϕ_d	8×10^{-2}	1×10^{-2}
ϕ_s	1	5×10^{-4}
ϕ_c	2×10^{-1}	7×10^{-5}
ϕ_b	2×10^{-1}	2×10^{-5}
ϕ_t	2×10^{-1}	5×10^{-7}
ϕ_g	3	8×10^{-4}
σ_Q^2	5×10^{-3}	4×10^{-3}
σ_Q^5	5×10^{-6}	1×10^{-4}
σ_Q^7	5×10^{-5}	1×10^{-4}

which other Standard Model particles obtain their masses from the Higgs field. Depending on the energy scale where the supersymmetry-breaking scale takes place, some of these moduli can have (meV/ c^2)-level masses, and thus macroscopic Compton wavelengths.

The moduli manifest as scalar fields, which generate Yukawa-type modifications of the inverse square law of the form of Equation (6.4). With one particular set of assumptions concerning natural energy scales and other parameters, the authors of Reference [179] derive the expected Yukawa modifications resulting from 10 distinct moduli of a particular theory, focusing on couplings to quarks and gluon fields, which are reproduced in Table 6.1. Different theories with different sets of assumptions may yield different expectations for the masses of the moduli as well as their expected couplings to Standard Model particles. This should simply serve as an illustrative example of that fact that many scalar fields are manifest in supersymmetric string theories, so that if string theories represent a “correct” description of our universe, then one can expect modifications to the inverse square law at some short-distance length scale.

6.3.4 Extra Dimensions and String Theory

String theories can also involve assuming that the universe itself is fundamentally higher dimensional, and that our observable spacetime is just a 4-dimensional subspace of the larger universe which may be anywhere from 5 to 11 dimensions. The extra dimensions are usually space-like, based on their signature within a General Relativity spacetime metric, but there are some formulations that involve extra time-like dimensions [180]. Many of these theories are attractive in that they unify all the fundamental forces, including gravity, while providing natural explanations of the difference in scale between the strength of gravity of the strengths of the other fundamental forces [155–158, 181, 182].

In string theory, the known Standard Model particles are primarily limited to interactions within a 4-dimensional subspace³ that constitutes the observable universe, although this may not hold at all energy scales. Gravity, however, is allowed to freely propagate in the extra dimensions since they’re fundamentally incorporated into the spacetime metric. Some string theories suggest that the extra dimensions are compactified [155, 156], meaning that they have finite extent and in some cases are periodic, such that motion along one dimension eventually returns to its starting position. As one might expect, gravity behaves fundamentally differently at distances shorter than the compactification scale, where the new dimensions play a critical role.

Following a specific formulation [155, 156], let’s assume these extra dimensions can be factorized within the metric, are finite and have been compactified on a torus of radius R_* , and that there is a fundamental mass scale associated with this theory given by M_* , where energy and length scales can be computed accordingly. Then, at distances shorter than R_* , the classical Newtonian potential and the gravitational constant are given by,

$$V(r) \approx -G_{(4+n)} \frac{m_1 m_2}{r^{1+n}} \quad \text{with} \quad G_{(4+n)} = \frac{4\pi}{S_{(2+n)}} \left(\frac{\hbar}{M_* c} \right)^{2+n} \frac{c^3}{\hbar}, \quad (6.8)$$

where $S_{(2+n)}$ is the (dimensionless) area of a unit $(2+n)$ -sphere, $S_{(2+n)} = 2\pi^{(n+1)/2}/\Gamma[(n+1)/2]$. If we take $n = 0$, then $G_{(4+n)} \rightarrow G = \hbar c/M_P^2$ upon recognizing the fundamental scale $M_* \rightarrow M_P$. At distances $r \gg R_*$, the gravitational flux lines cannot penetrate the extra

³These subspaces are often called “branes” in reference to the concept of a 2-dimensional membrane in 3-dimensional space, but generalized to higher dimensions. The exact type of brane we live on is subject to the construction of theory. One example is that the Standard Model fields are localized in the 4-dimensional throat of a 6-dimensional vortex [155]

dimensions, and the potential takes the form

$$V(r) \approx -G_{(4+n)} \frac{m_1 m_2}{R_*^n} \frac{1}{r}, \quad (6.9)$$

with $G_{(4+n)}$ as before. In order for this theory to produce the expected gravitational potential at long distances, it must be that $G = G_{(4+n)} R_*^{-2}$. If it is assumed that $M_* \sim m_{\text{EW}}$ where $m_{\text{EW}} \sim 1 \text{ TeV}/c^2$ is the electroweak scale, so that this assumption solves the hierarchy problem, then the scale of the compact dimension is given by,

$$R_* \sim 10^{\frac{30}{n}-17} \text{ cm} \times \left(\frac{1 \text{ TeV}/c^2}{m_{\text{EW}}} \right)^{1+\frac{2}{n}}. \quad (6.10)$$

Compact dimensions under this construction will manifest at short range via the exchange of Kaluza-Klein gravitons [141, 181, 182], which are massive spin-2 particles with an effective range limited by the size of the extra dimensions. For n new dimensions all compactified on a flat torus with same characteristic dimension R_* , one could expect Yukawa modifications from Equation (6.2) of the form,

$$\alpha = \frac{8n}{3} \quad \text{and} \quad \lambda = R_*, \quad (6.11)$$

where the value of α and λ depend significantly on the underlying model, and different compactification geometries will necessarily produce different couplings.

Similar theories of higher-dimensional gravity allow the size of the extra dimensions to be dynamical quantities, parameterized by spin-0 fields. The linear combination of these size-parameterizing fields is referred to as Brans-Dicke scalar [183] often dubbed the “radion”. A massive radion in a $(4+n)$ -dimensional theory of gravity would mediate an additional gravitational force with a Yukawa modification given by,

$$\alpha = \frac{n}{n+2} \quad \text{and} \quad \lambda = R_*, \quad (6.12)$$

where R_* is still the effective size of those extra dimensions. It is also possible that for “small” extra dimensions $R_* \lesssim 1 \text{ nm}$, the radion can still have a Compton wavelength of the order of $100 \mu\text{m}$, conveniently near the dark-energy length scale.

Following the initial developments of these theories in References [155, 156], it was further postulated that the extra dimensions may exhibit fundamentally more coupling to the usual

dimensions within the higher-dimensional spacetime metric. These are sometimes called “warped” dimensions, and have metrics of the form [157, 158, 184],

$$ds^2 = f(\xi_i)\eta_{\mu\nu}dx^\mu dx^\nu + g_{ij}(\xi_i)\xi_i\xi_j, \quad (6.13)$$

where ξ_i represent the i extra dimensions, f and g are functions of those dimensions with $f(\xi_i)$ called the warping factor, and $\eta_{\mu\nu}dx^\mu dx^\nu$ would yield the usual invariant in the absence of the extra, warped dimensions. The extra dimensions ξ_i can be compact [157], or infinite in extent [158, 184], and the specifics of the warping can greatly alter the masses, Compton wavelengths, and couplings of the Kaluza-Klein graviton modes and the radion. Generally speaking, it’s expected that interactions between the radion of the Kaluza-Klein modes and normal matter would result in modifications to the inverse square law at separations shorter than the Compton wavelength.

6.3.5 Massive Gravitons

In General Relativity, gravitational waves or any perturbations of the spacetime metric are theorized to travel at the speed of light [185]. If a quantum theory of gravity is to be constructed, this would naturally lead to the assumption that the force carrying particle, the graviton, should be massless, so that the gravitational interaction is mediated at the speed of light. However, breaking this assumption and allowing the graviton to have a mass has interesting consequences, including possible modifications to the inverse square law at short-distance. Massive gravity is discussed in great detail in the Reviews [186, 187], and has received significantly increased attention in the last decade.

Some of the first attempts to construct a theory of massive gravitons had issues with the production of “ghosts”, or particle states that are considered unphysical, for example, states with a negative norm or violations of Lorentz-invariance [188]. Ghost-free solutions were developed in recent years [189], which have in turn lead to theories of bigravity, which include two gravitons: one massive and one massless [153, 190]. This lifts a number of phenomenological constraints, for example that a single massive graviton should have an exceedingly small mass so as to reliably reproduce the equations of General Relativity in the regime where they have been tested with great precision.

As with many of the scalar fields and bosonic particles discussed in previous sections, the massive graviton would modify the strength of gravity at distances around the Compton

wavelength. Assume $(1/M_G)$ denotes the energy scale of the coupling between matter fields and the massive graviton, where the inverse of the usual Planck mass $(1/M_P)$ denotes the energy scale of the coupling between matter fields and the massless graviton, and apply a weak field approximation, then Yukawa-type modifications to the inverse square law would be expected to have the form [153],

$$\alpha = \frac{4M_P^2}{3M_G} \quad \text{and} \quad \lambda = \frac{\hbar}{mc}, \quad (6.14)$$

with m the mass of the massive graviton, and λ simply the Compton wavelength. The authors of Reference [153] go further and even suggest that such massive gravitons could be dark matter candidates given certain assumptions related to their relic abundance, which in turn puts constraints on acceptable masses m and coupling strengths M_G^{-1} .

6.3.6 Fat Gravitons

Throughout the Standard Model and General Relativity is the implicit assumption of locality, which, in part, conceptualizes particle as point-like. This isn't the usual statement of the principle of locality, but it is nevertheless a useful viewpoint. If this assumption is relaxed and the graviton is allowed to have a macroscopic extent l_g and soft couplings to heavy Standard Model particles [159], then this has interesting implications. Physics occurring at length-scales much smaller than the size of the graviton would, in a sense, be invisible to the gravity, and the size of the gravity can explain the cosmological constant.

Importantly for our research pursuits, at distances $d < l_g$, the force of Newtonian gravity between two masses composed of normal matter falls off sharply. The modified gravitational force between two point masses m_1 and m_2 would take the form [159, 160],

$$F_{G,\text{fat}} = \frac{Gm_1m_2}{r^2} \left(1 - e^{-3kr/l_g}\right), \quad (6.15)$$

where k is a factor of order unity. Detecting a decrease in the expected gravitational force between two objects as they get closer would be a clear signature of new physics.

6.4 Previous Searches for New Interactions

Astronomical observations combined with terrestrial and laboratory based experiments have placed limits on Yukawa modifications of the inverse square law of Newtonian gravity over

roughly 20 orders of magnitude of the length scale λ , with the most sensitive probe coming from lunar laser ranging measurements at distances of $\sim 10^9$ m [191, 192].⁴ At this length scale, Yukawa-type modifications of the inverse square law have been excluded for $\alpha \gtrsim 5 \times 10^{-11}$ based on the anomalous precession of the moon’s orbit, while longer length scales have been tested by examining the orbits of planets within our solar system and intermediate length scales between the Earth-moon distance and laboratory scales have been tested by satellite-based experiments [141]. At much shorter length scales in the submillimeter regime, the sensitivity of apparatuses to gravitational-strength interactions falls off sharply, and the existing limits are correspondingly much weaker. At $\lambda \sim 10$ μm , Yukawa-type modifications have only been excluded for $\alpha \gtrsim 5 \times 10^3$, and it gets much worse at $\lambda \sim 1$ μm where $\alpha \gtrsim 10^7$. For $\lambda \lesssim 10$ nm, constraints are essentially non-existent and are derived from a variety of scattering and collider experiments. The landscape of current measurements is shown in Figure 6.1.

Traditionally, short-distance gravitational interactions have been experimentally investigated using sophisticated torsion balances [200] which establish some of the most stringent bounds on deviation from the inverse square law at both the meter scale [201–203] and the submillimeter scale [141, 193, 195, 197–199, 204]. Alternative techniques have been developed using nanotechnology to mount test masses at the ends of microcantilevers [194, 196, 205]. In essentially all realizations of short-distance force sensing, the most significant background results from residual electromagnetic interactions between test masses, source masses, and external fields. Adequate environmental shielding to prevent stray fields from interfering with the experiment is necessary, as well as dedicated electromagnetic shielding between the gravitating masses (or a sophisticated understanding of the background forces present) in order to avoid accidentally measuring the electrical properties of the surfaces rather than any of the interactions we’re actually interested in.

Generally, all measurements within this field are limited by systematic effects, such as the reliability and reproducibility (or lack thereof) in the positioning and alignment of the macroscopic objects involved, especially given the small separations required for competitive measurements. Another concern of significant note is the occasional assumption that oscillators with high quality-factor, $Q \gtrsim 10^3$, don’t fluctuate with temperature when operated at $T \sim 300$ K without explicit temperature control. Even small fluctuations away from an

⁴This experiment represents an indirect validation of the moon landings in the late 1960s and early 1970s, given that lunar laser ranging is only possible if the Apollo astronauts placed retroreflectors on the surface of the moon, as opposed to a Hollywood studio in Burbank (as the conspiracy might imply).

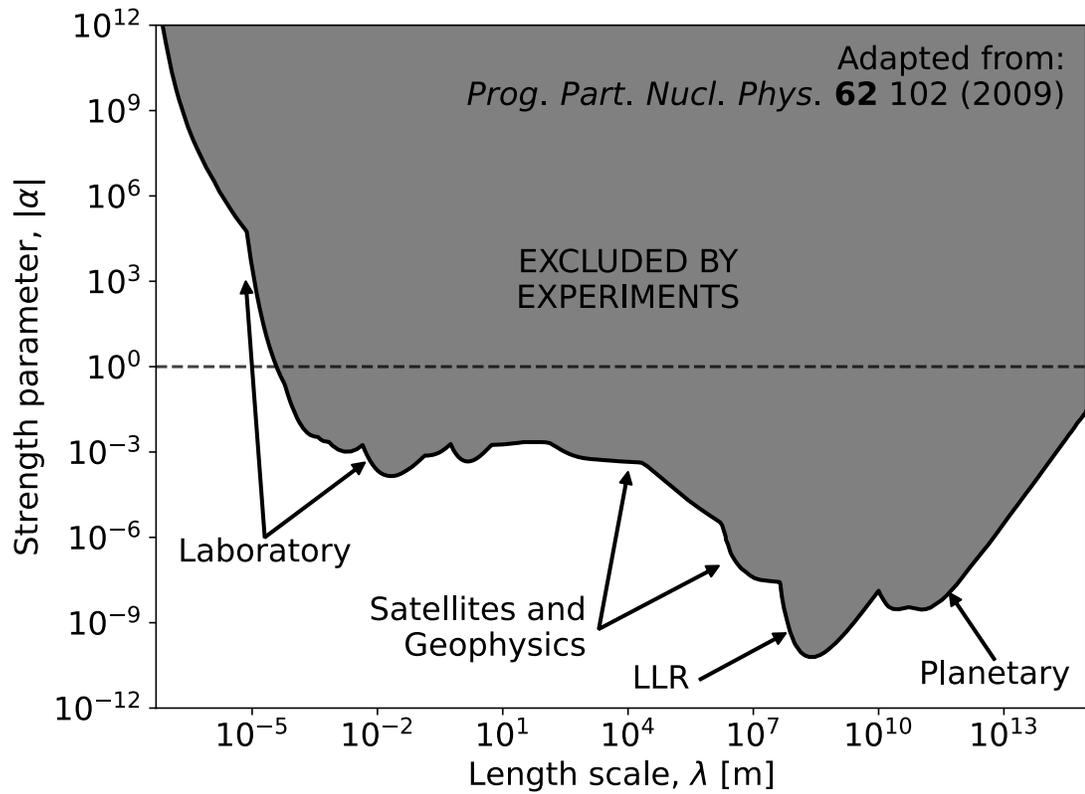


Figure 6.1: Parameter space of Yukawa-type, $\alpha - \lambda$, modifications to gravity, including current experimental constraints. Gravitational strength interactions, $\alpha \approx 1$, are indicated with a dashed line. Adapted from References [141, 160], with significant additions in the submillimeter regime (the region of interest to this work) from more recent publications [193–199].

assumed resonance can significantly limit sensitivity, which in turn results in over-zealous limit estimation if not properly accounted for. Hence, experimental progress calls for new techniques with different attributes and systematics that may eventually contribute to robust discoveries. Short-distance force sensing with optically levitated microspheres offers just such a technique.

6.4.1 Torsion Balances

One of the first tests of the inverse square law was performed by Cavendish in 1797, where he used a large torsion balance in a program he dubbed “weighing the world”. A 2 m wooden rod was suspended by a long wire (the torsional spring), with two spherical lead test masses (roughly 0.75 kg) attached to either end. Two much larger lead spheres (roughly 150 kg) were positioned such that their gravitational interaction with the small spheres produced a net torque on the wire-suspended rod. The torque is transduced to an angular deflection by the torsion balance, which can be then be observed. By comparing the strength of this interaction to the weight of the small spheres (i.e. the strength of the gravitational interaction between the Earth and the test masses), he was able to use Newton’s Law of Universal Gravitation to calculate the density of the Earth [200]. Reinterpretations of his result as a direct calculation of G yield values within a few percent of the currently accepted value of G (from CODATA [206] which averages over distinct measurements, even in the presence of tension).

Since the days of Cavendish, the technology of torsion balances has improved significantly. Tests of the inverse square law at scales $\lambda \sim 1$ m have been made using “large” torsion balances with a variety of geometries [201–203], although we will focus on newer measurements using more compact torsion balances [194, 197–199, 204] that access the shorter length scales, $\lambda \lesssim 1$ mm, of interest to our scientific program. Indeed, if the torsion balance is made sufficiently compact, test and source mass separations of a few hundred nanometers are also possible, allowing these devices to probe even shorter length scales [195].

In one of these newer iterations [198], the torsion balance consists of a metallic disk serving as the test mass ($r \sim 2.5$ cm), which is suspended from a wire serving as the torsion spring, with a regular, 120-fold symmetric azimuthal pattern of voids on the bottom/downward-facing side of the disk. A second, identical disk serves as the source mass, such that when placed below the test mass, with the voids now on the top/upward-facing side, rotation of the source mass relative to the test mass produces an oscillating torque on the test mass

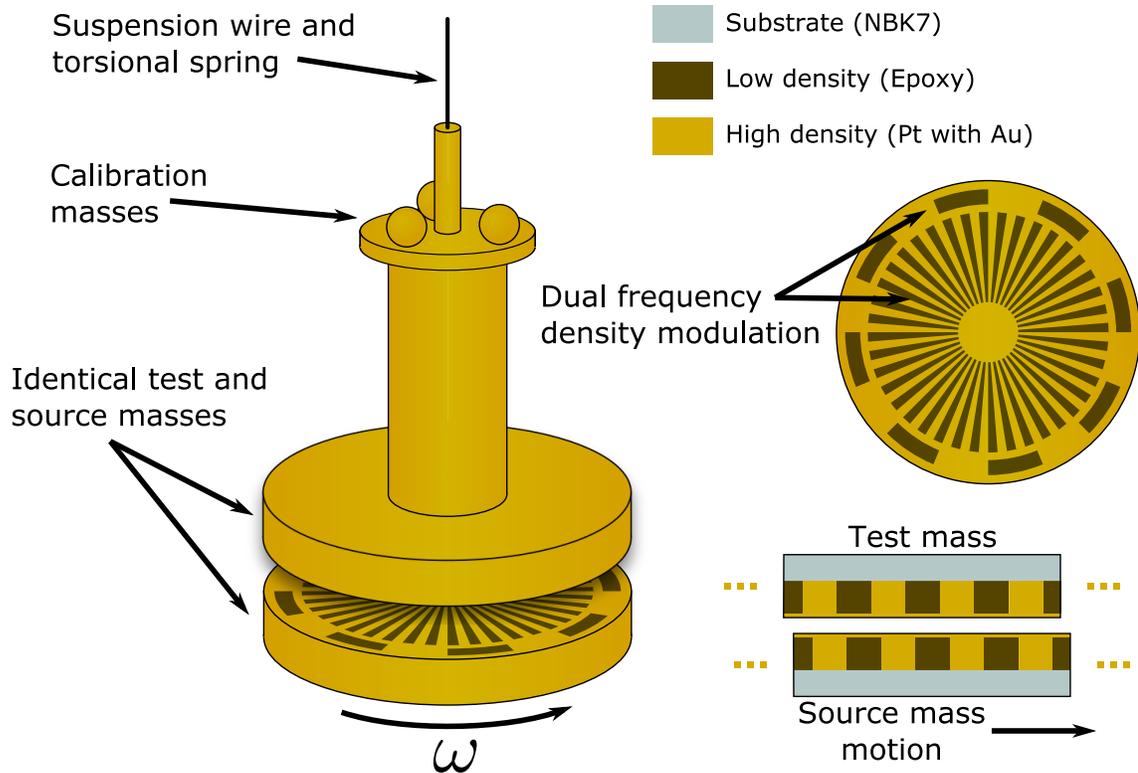


Figure 6.2: Schematic depiction of the torsion balance described by Reference [198]. The test and source masses are patterned platinum foils adhered to a glass substrate, with the epoxy serving as the low-density component of the modulation. The components are all gold-coated to reduce electrostatic interactions. A stationary electrostatic shield (not depicted) sits between the source and test masses to shield anomalous backgrounds. As the source mass rotates at angular frequency ω , an oscillating gravitational torque, plus any new interactions, causes the test mass assembly to torsionally oscillate at $n\omega$ and $m\omega$ where n and m are the number of full periods in the two distinct density modulations present. The density modulations depicted have 8-fold and 40-fold symmetry, while the actual device has 18-fold and 120-fold symmetry.

as the voids rotate past one another. A schematic depiction of this measurement apparatus is shown in Figure 6.2. The test and source masses are brought as close as $\sim 50 \mu\text{m}$, quite remarkable considering their macroscopic size of $\sim 5 \text{ cm}$. This is currently the most sensitive apparatus to gravitational-strength interactions between $10 \mu\text{m}$ and $100 \mu\text{m}$.

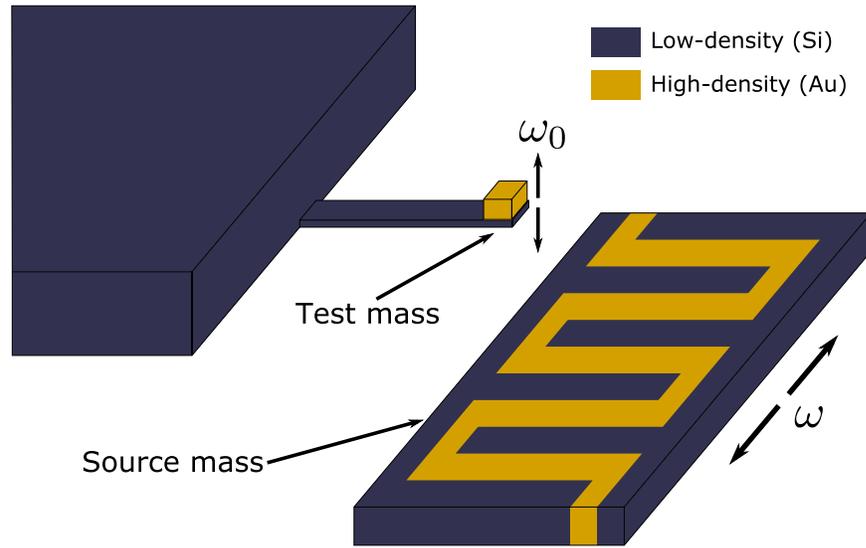
Another scheme involves a more traditional-looking torsion balance [199], in which test masses are attached to the ends of a rod, which is in turn suspended by a wire serving as the torsional spring. A source mass is constructed by placing rectangular tungsten masses in a regular azimuthal pattern with 8-fold symmetry on a glass disk, whose diameter matches the length of the rod. With the source mass disk rotating in a vertical plane (horizontally oriented rotation axis) and positioned such that the rod lies along the face of the disk, the individual tungsten source masses periodically pass the test masses and induce an oscillating torque on the torsion balance, which can be deterministically enhanced by displacing the center of the disk’s rotation from the center of the test mass rod. A clever positioning of displaced compensation masses serves to nearly cancel the Newtonian component of the torque, allowing for a “near-null” measurement. Unlike Reference [198], where the torsion balance was allowed to free-oscillate about a natural equilibrium position, this second scheme requires the use of active feedback to maintain the pendulum about its equilibrium due to anomalous torques from contact potentials.

In both experiments mentioned [198, 199], a thin, conductive membrane is suspended between the source and test masses in order to suppress electromagnetic backgrounds, although this is not a requirement as demonstrated by the success of other works [195]. The sensitivity of the two apparatuses is such that they can easily resolve Newtonian gravity ($\alpha \sim 1$), providing a robust model to which they can compare their data. Deviations from this model are then constrained with profile-likelihood techniques, yielding the reported exclusions shown in Figure 6.1.

6.4.2 Microcantilevers

The experiments making use of torsion balances and discussed in the previous section involve generating oscillating torques, where the frequency of the oscillation is intentionally tuned *away* from the resonance of the torsion balance itself.⁵ Operating on-resonance is generically difficult, especially in the case of oscillators with high quality-factor, as small fluctuations

⁵For small angles, the motion of a torsion pendulum is that of a damped harmonic oscillator, and thus Appendix B is applicable.



Drive n periods
"on resonance" : $\omega_0 = n\omega$

Figure 6.3: Schematic depiction of the microcantilever apparatus described by Reference [194]. The test mass consists of a solid gold block mounted to the end of a silicon cantilever, with some loaded resonant frequency ω_0 . A nanofabricated chip with a regular $25\ \mu\text{m}$ density modulation serves as the source mass, and is driven to oscillate n periods of the modulation at a frequency ω such that $\omega_0 = n\omega$. A stationary electrostatic shield (not depicted) sits between the source and test masses to shield anomalous backgrounds. The high-density gold portion is also continuous, allow a current to be driven through it for electromagnetic calibrations.

of environmental conditions such as temperature can alter the resonant frequency and push the measurement off-resonance, significantly limiting sensitivity. However, if environmental conditions can be robustly controlled while the resonant frequency is simultaneously monitored, operating with drive forces on-resonance can offer huge improvements in sensitivity, since the system response is amplified by the quality factor $Q \sim \omega_0/\gamma$, relative to the low-frequency response, where ω_0 is the resonant frequency of the system and γ is the damping (see Appendix B).

This is generally the approach of a class of experiments using microcantilevers [194, 205], wherein a homogeneous test mass is mounted to a cantilever, then subjected to possible new interactions from a source mass, such that the new interactions excite motion at the test-mass-loaded-cantilever's resonant frequency, ω_0 . Essentially, a source mass is designed with some spatial periodicity in mass density, such as alternating regions of gold and silicon [194], or alternating regions of tungsten and indium antimonide [205]. The source mass is then made to rotate or oscillate at some frequency ω such that n full periods of the density modulation are exposed to the test mass in one period of the rotation or oscillation, exciting the test mass at $n\omega$. By actively monitoring the resonant frequency of the loaded cantilever, which can fluctuate with environmental conditions, ω can be continuously tuned such that $n\omega = \omega_0$. A schematic depiction of this type of measurement is shown in Figure 6.3.

Micro- and nanofabrication techniques allow for a wide range of cantilever dimensions, and thus resonant frequencies, although the most common type is a rectangular prism akin to a diving board, where one end is allowed to free oscillate, while the other is rigidly connected to a handle. Resonant frequencies can range anywhere from 100 Hz to 1×10^5 Hz, but this number is reduced significantly when a test mass is attached to the free end, and requires robust and continuous characterization for a proper experiment.

6.5 Extending the Experimental Reach

At submillimeter length scales, the most sensitive measurements of gravitational forces come from a number of torsion balance experiments. At a few specific values of λ in this regime, microcantilever experiments approach the same sensitivity as the best torsion balances, but do not surpass them. Thus, almost the entirety of the $\alpha - \lambda$ parameter space is dominated by a single type of measurement framework. For robust characterizations of gravity, as well as

any eventual discoveries of new physics, multiple approaches to the same task provide necessary cross-validation. Optically levitated microspheres are one such alternative approach, subject to entirely different set of systematic effects.

Chapter 7

Search for Screened Interactions: The Chameleon Field

The content of this chapter has been adapted from Reference [25], for which the author was a primary contributor. Apparatus and generic experimental method descriptions have been removed, appealing to Chapter 2 of this thesis instead, while some other sections have been expanded for completeness.

7.1 Introduction

It might be possible to detect the presence of a scalar field constituting dark energy by searching for new interactions between objects separated by distances below the dark energy length scale [140, 168, 204, 207]. In many cases, the resulting forces can be substantially larger than Newtonian gravity at short distances [168, 208]. The most sensitive previous searches for violations of Newtonian gravity at or below the dark energy length scale employed macroscopic test masses or a conductive shield between the probe and test masses to minimize electromagnetic backgrounds [160, 194, 195, 204, 209].

Although these experiments place stringent constraints on deviations from Newtonian gravity, it is possible to construct theories of dark energy involving new forces that could have avoided detection due to the geometry and scale of previous experiments [168, 173, 175, 208]. For these screened interactions, recent searches using microscopic test masses such as atoms [172, 210] or neutrons [211–213] often provide the strongest constraints.

Several screening mechanisms have been proposed to evade existing experimental constraints on scalar interactions in the laboratory and solar system [168]. A specific example is the chameleon mechanism [169, 170], in which the effective mass of the chameleon particle (corresponding to the inverse length scale of the interaction) depends on the local matter density. At cosmological distances where the matter density is low, the chameleon field would constitute the vacuum energy density necessary to explain the accelerating expansion of the universe [171] (see Chapter 6). However, most laboratory experiments are carried out in regions of high matter density, where the forces arising from the chameleon interaction are suppressed.

This work presents a search for screened interactions below the dark energy length scale using optically levitated μm -size dielectric spheres as test masses. Levitated microspheres in high vacuum [21, 32, 53, 54, 56, 57, 62, 80, 214] can be used to detect forces $\ll 10^{-18}$ N [24, 56, 60, 215, 216], and in many cases their small size avoids screening effects.

7.2 Experimental Setup

The test masses used in this work consist of amorphous silica microspheres with radius $r = 2.35 \mu\text{m}$ and mass $m = 0.13 \text{ ng}$ [76]¹ levitated in a single-beam, upward-propagating 1064 nm laser trap [7, 24]. The radiation pressure from the laser counters Earth's gravity and acts as an optical spring pulling the microsphere to the center of the Gaussian beam [6]. The resonant frequencies of the trap are ~ 150 Hz for the 2 degrees-of-freedom orthogonal to the Earth's gravity and ~ 100 Hz for the degree-of-freedom parallel to Earth's gravity. The position of the microsphere is measured by focusing secondary 650 nm Gaussian laser beams on the microsphere and imaging the pattern of scattered light onto a position-sensitive photodiode (PSPD). For small displacements from the center of the trap, the PSPD produces a differential current that is a linear function of the position of the microsphere.

When the microspheres are loaded into the optical trap, they typically have an electric charge of $\sim 100e$ [24]. The charge is measured by monitoring the response to an oscillating electric field. Microspheres are discharged with UV radiation from a Xenon flash-lamp. As shown in [24], clear charge quantization can be observed at the end of the discharging cycle, providing a force calibration for the system.

¹At the time of initial publication [25], the precision mass measurement (see Chapter 4) had not been performed, so the results presented here were limited by systematic uncertainty from the assumed microsphere radius and density.

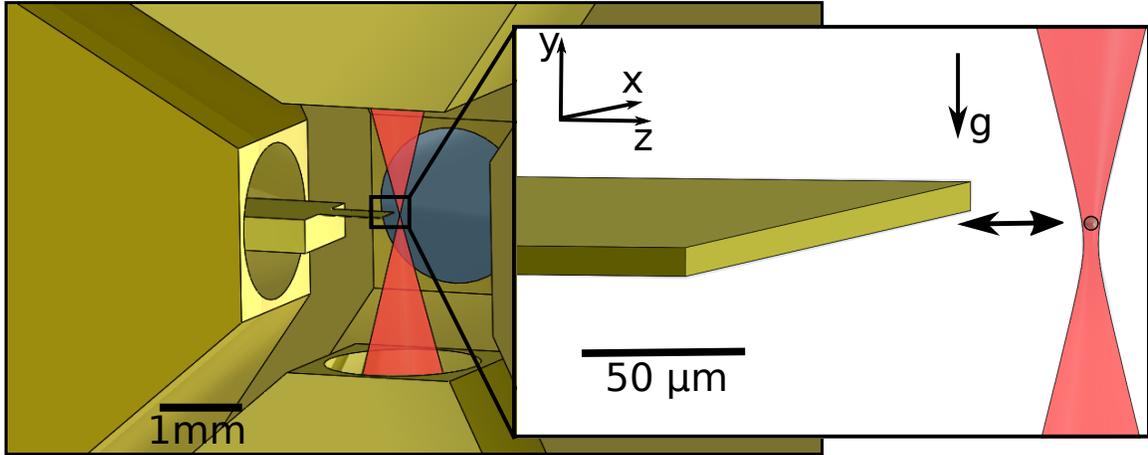


Figure 7.1: (left) Schematic of the optical trap and shielding electrodes. The electrode in the foreground is removed to show the inside of the trap. (right) Zoom in on the region near the trap. A $4.7\ \mu\text{m}$ diameter microsphere is suspended at the focus of an upward propagating laser beam. The $10\ \mu\text{m}$ thick Au-coated Si cantilever is positioned at $20\ \mu\text{m}$ to $200\ \mu\text{m}$ separations from the microsphere and oscillated in the z direction using a nanopositioning stage. (This figure is nearly identical to Figure 3.1, aside from a redefinition of the coordinate system. Both figures are included in their respective chapters to maintain clarity and consistency.)

The microspheres are levitated inside of a vacuum chamber to reduce the force noise coming from collisions with residual gas. Due to reduced gas damping, the trap becomes unstable below $0.05\ \text{mbar}$. To stabilize the trap, active feedback is applied by measuring the microsphere's position and modulating the position of the trap. Measurements are performed at pressures below $10^{-6}\ \text{mbar}$ where the noise for force measurement is limited to $2 \times 10^{-17}\ \text{N Hz}^{-1/2}$ by imaging noise. The optical setup and calibration methods are improved versions of those discussed in [24].

A schematic view of the apparatus is shown in Figure 7.1 where a coordinate system is defined. The microsphere coupling is probed with a silicon cantilever with dimensions $500\ \mu\text{m} \times 2000\ \mu\text{m} \times 10\ \mu\text{m}$ and a $500\ \mu\text{m}$ thick handle, fabricated from a silicon on insulator (SOI) wafer using optical photolithography and plasma etching. The $10\ \mu\text{m}$ dimension is oriented so that the cantilever clears the Gaussian beam waist of the laser and the $500\ \mu\text{m}$ dimension is approximately centered on the trap in the x direction. A $200\ \text{nm}$ gold shielding layer was evaporated onto the cantilever to minimize its electrostatic interactions with the microsphere. The cantilever is mounted on a 3-axis nanopositioning stage used to control

its spacing from the microsphere with a precision of 3 nm and a travel of 80 μm . The trap and cantilever are electrically shielded inside a cube consisting of six gold-plated electrodes separated by 4 mm, whose potentials are controlled by external digital-to-analog converters (DACs). The nanopositioning stage is mounted on a piezoelectric motor driven stage with 12 mm travel in the z direction to provide coarse positioning.

7.3 Dipole Response Measurement

To measure electrostatic interactions between the cantilever and the microsphere, each shielding electrode was set to a nominal potential of 0 V while the cantilever was biased to a non-zero potential. The z position of the nanopositioning stage was driven with an 18.3 Hz sine wave over its full 80 μm travel. The microspheres were aligned with the center of the cantilever in the y direction by determining the position at which the maximum electrostatic response was seen as the cantilever was swept in the z direction at fixed bias. The microsphere and stage positions were recorded in 50 s long integrations. Data were acquired for coarse stage positions with closest approach of 20, 60, 100, and 150 μm . This procedure was repeated for each of three microspheres considered in this work.

To eliminate low frequency drifts, the microsphere positions were first mean subtracted. The data were then averaged in 10 μm cantilever position bins and calibrated to force units using the single-charge-step calibration discussed previously. Data at neighboring coarse-stage positions were matched in the 30 μm to 40 μm overlap region. The measured electrostatic force versus spacing between the cantilever and the microsphere is shown in Figure 7.2.

Although electrically neutral microspheres are used, they still contain $\sim 10^{14}$ charges and interact primarily as electric dipoles. The force on a microsphere with dipole moment \vec{p} is given by $\vec{F} = (\vec{p} \cdot \nabla) \vec{E}$ [4], where $\vec{p} = \vec{p}_0 + \alpha \vec{E}$ consists of a permanent dipole, \vec{p}_0 , and an induced dipole, $\alpha \vec{E}$, for polarizability α . The latter term incorporates any dipole induced by an electric field, including the linear dielectric response as well as any non-zero surface charge mobility. Figure 7.2 shows a fit to the model $\vec{F} \cdot \hat{z} \equiv F_b(z) = (p_x \partial_x + p_y \partial_y + p_z \partial_z) E_z \approx p_{0z} \partial_z E_z + \alpha E_z \partial_z E_z$.

A finite-element method (FEM) was used to solve for \vec{E} within the geometry of the trapping region. Dipole moments and polarizabilities were extracted by fitting the microsphere responses at non-zero cantilever bias to $F_b(z)$. The results of this fit for each microsphere are

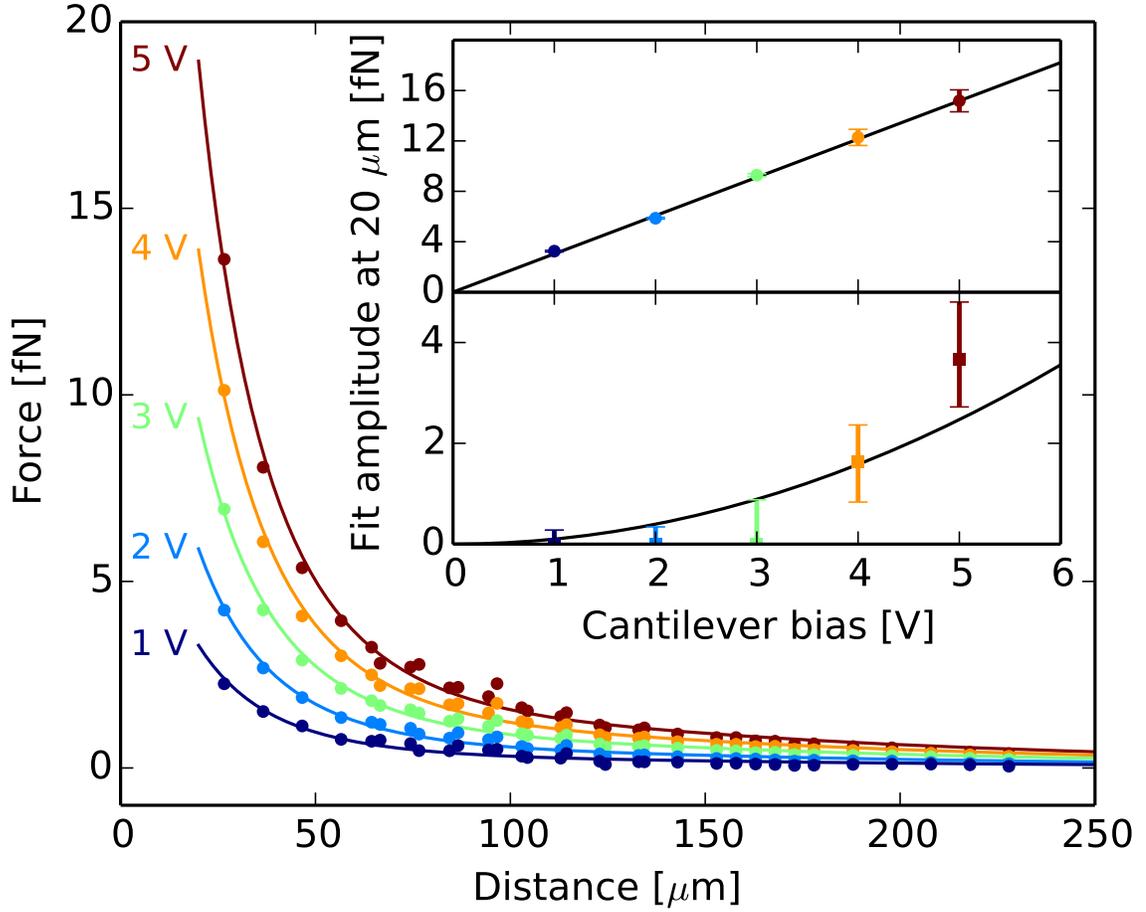


Figure 7.2: Measured response of microsphere #1 versus distance from the cantilever face as the cantilever is swept in z with a constant bias of 1, 2, 3, 4, and 5 V. The data points are shown as dots and the best fit model as solid lines. (inset) Amplitude of the fit component $\propto \partial_z E_z$ (top) and the fit component $\propto E_z \partial_z E_z$ (bottom). Fits to the expected linear and quadratic dependence on the voltage are also shown (solid lines).

Table 7.1: Dipole moments and polarizabilities measured for each microsphere.

Microsphere	p_{0z} [$e \mu\text{m}$]	α/α_0
#1	151 ± 6	0.21 ± 0.13
#2	89 ± 10	0.00 ± 0.33
#3	192 ± 30	0.25 ± 0.14

shown in Table 7.1. The dipole moments are measured in units of $e \mu\text{m}$ and the polarizabilities are reported relative to $\alpha_0 = 3\epsilon_0 \left(\frac{\epsilon_r - 1}{\epsilon_r + 2}\right) \left(\frac{4}{3}\pi r^3\right)$ assuming $\epsilon_r \sim 3$ and $r = 2.35 \mu\text{m}$. The reported values of polarizability, which are smaller than α_0 , could arise from systematics in the determination of a small induced dipole on top of a much larger permanent dipole, an unexpectedly low value of ϵ_r , or a smaller than expected volume.

7.4 Probing for Chameleons

Following the measurement of the electrostatic interaction at a given coarse stage position, the cantilever was set to a nominal potential of 0 V, and twenty additional 50 s long integrations were acquired to search for new screened interactions. This procedure was then repeated to obtain three 1000 s measurements at each coarse stage position in order to quantify time dependent variation in the measured response over a period of several hours. The standard deviation of the repeated measurements at each position bin was included as an additional systematic error.

The resulting data can be used to set constraints on new screened interactions at distances of the order of the dark energy length scale, with a sensitivity that is limited by the presence of the residual electrostatic backgrounds. As a concrete example, we consider the presence of a non-relativistic, steady-state chameleon field, ϕ , that mediates a force between the microsphere and cantilever. Following [172, 173, 208], we assume an inverse power law form of the effective potential,

$$V_{\text{eff}}(\phi) = \Lambda^4 \left[1 + \left(\frac{\Lambda}{\phi} \right)^n \right] + \frac{\beta \rho}{M_{Pl}} \phi, \quad (7.1)$$

where Λ is the scale of the chameleon self interaction, often chosen at the dark energy scale $\Lambda \sim 2.4 \text{ meV}$. The coupling to matter of density ρ is determined by the scale $M = M_{Pl}/\beta$ where M_{Pl} is the reduced Planck mass and β is unitless. Although other power laws are possible, $n = 1$ was chosen as a characteristic example for this search.

Similar to the electric field calculation described above, an FEM was employed to solve the non-linear equation of motion $\nabla^2 \phi = \partial_\phi V_{\text{eff}}$ in the geometry described previously. The residual gas pressure of $\sim 10^{-6} \text{ mbar}$ was included, but has negligible effect on ϕ for values of the matter coupling considered in this work. Boundary conditions were set to the equilibrium value of the field within the cantilever and electrodes, following the detailed treatment of

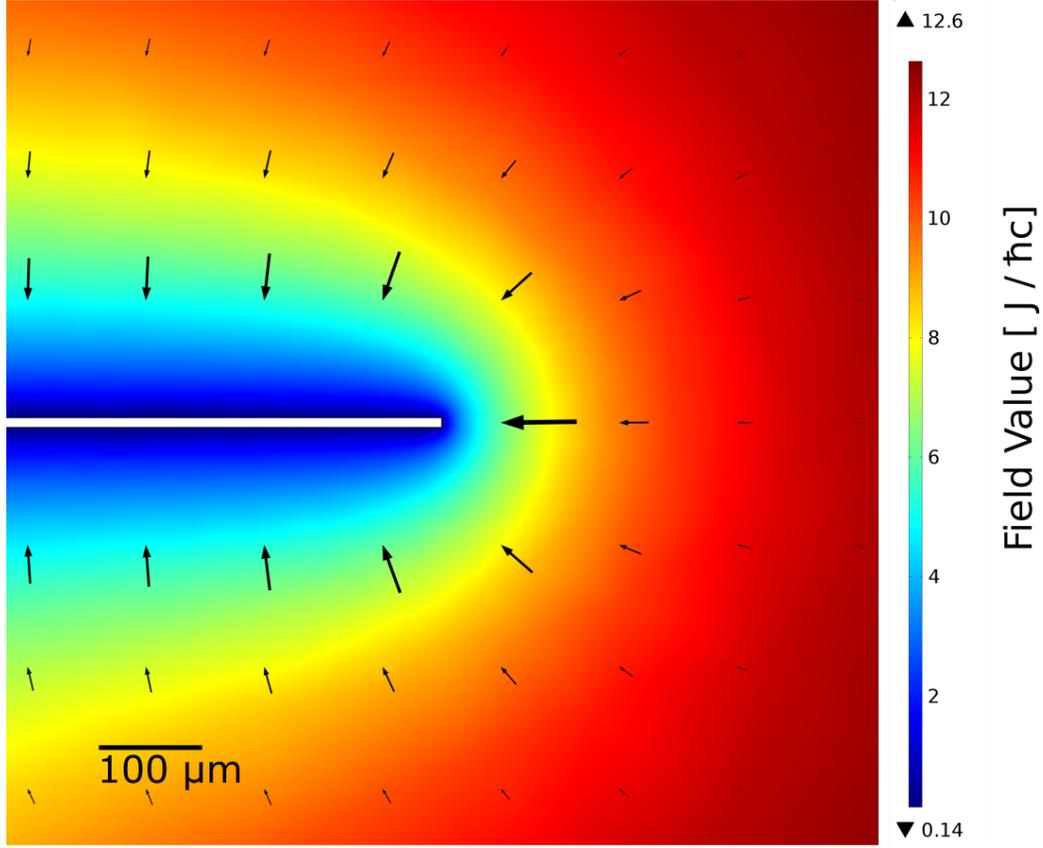


Figure 7.3: An example FEA calculation of the chameleon field. A cross-section in the yz -plane is shown, so that the cantilever appears edge-on and extends into and out-of the page. The residual vacuum pressure of $\sim 10^{-6}$ mbar does not affect the maximum value of the field here, which is instead limited by the finite size of the vacuum chamber, as the boundary condition enforces a low value at the vacuum-matter interface

matter-vacuum interfaces in [172]. An example field calculation is shown in Figure 7.3, where the field has been computed assuming $\Lambda = 2.4$ meV and $\beta \approx 10^5$.

The resulting chameleon force on a microsphere in the z direction was calculated as

$$F_c(z, \beta, \lambda) = \lambda(\beta\rho/M_{Pl}) \int_V (\partial_z \phi) dV, \quad (7.2)$$

where ρ and V are the density and volume of the microsphere, and λ is a screening factor [172, 173]. In the region of parameter space where $\rho r^2 < 3M_{Pl}\phi/\beta$, the microsphere is unscreened and $\lambda = 1$. However, when β becomes sufficiently large, the force on the microsphere is suppressed by $\lambda < 1$ [172, 173].

The measured force versus position for each of the three microspheres is shown in Figure 7.4. A small residual force $\lesssim 10^{-16}$ N can be seen for each microsphere. This response is consistent with electrostatic forces resulting from the permanent electric dipole moment of the microspheres coupling to the electric field produced by potential differences between the cantilever and shielding electrodes of $\lesssim 30$ mV. Contact potentials of this scale are expected to arise between connections to the electrodes in the vacuum chamber and external electronics.

To search for possible new interactions, the data for each microsphere were fit to a model $F(z) = A_c F_c(z, \beta, \Lambda) + A_b F_b(z) + A_0$, where $F_b(z)$ is the shape of the empirical background measured for each microsphere when the cantilever was biased to 5 V, A_b is the unknown electrostatic background amplitude due to residual contact potentials on the electrodes, and A_0 accounts for the arbitrary offset subtracted from the data at each coarse stage position.

A_c , the normalization of the chameleon force, was constrained in the fit by the following systematics. The microsphere mass was not directly measured, but the radius of the spheres was determined by the manufacturer to be (2.50 ± 0.24) μm , leading to a 35% uncertainty on the chameleon force.² Fits of the calibration data to the electric field simulations indicate that the microsphere was centered in y relative to the cantilever within 4 μm , leading to an uncertainty on the amplitude of the chameleon force of 1.8%. The z position of the coarse stage was determined from microscope images of the cantilever to $\lesssim 10$ μm , at each coarse stage setting. Using the positions and uncertainties determined from the calibration images, the coarse stage positions were further refined by allowing z -position offsets to float at each coarse stage position in the electrostatic fit. The best fit positions were used in the final chameleon fit, and their uncertainties contribute an additional systematic error of 6%. All errors were added in quadrature for a total systematic error of 36% on A_c , dominated by the uncertainty in the microsphere masses.

At each value of Λ , the profile of the negative log likelihood (NLL) was calculated by minimizing the NLL for the fit at each value of β over the nuisance parameters A_c (including its Gaussian constraint), A_b and A_0 . The 95% confidence interval for β was determined from the combined profile from all three microspheres following Wilks' theorem [218, 219]. This was done assuming that 2NLL follows a χ^2 distribution with one degree-of-freedom (DOF). The χ^2 statistic at the best fit point and for the background only model indicates that both

²Again, this uncertainty is large, as the precision mass measurement (see Chapter 4) had not been performed.

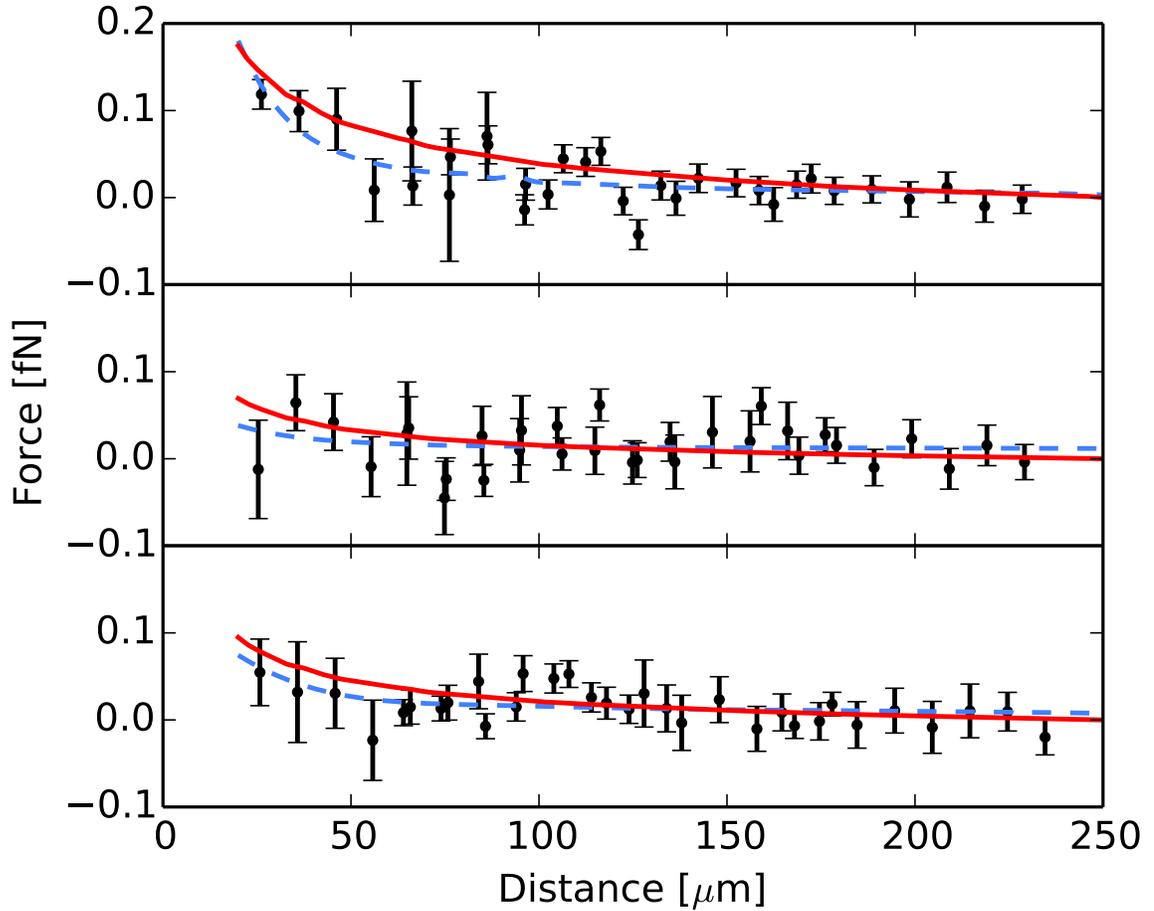


Figure 7.4: Measured response for microspheres #1 (top), #2 (middle), and #3 (bottom) versus distance from the cantilever face as the cantilever is swept in z with a nominal bias of 0 V. The best fit electrostatic background-only model (dashed) and the amplitude of a chameleon force at the 95% CL upper limit for $\Lambda = 10$ meV (solid) are also shown.

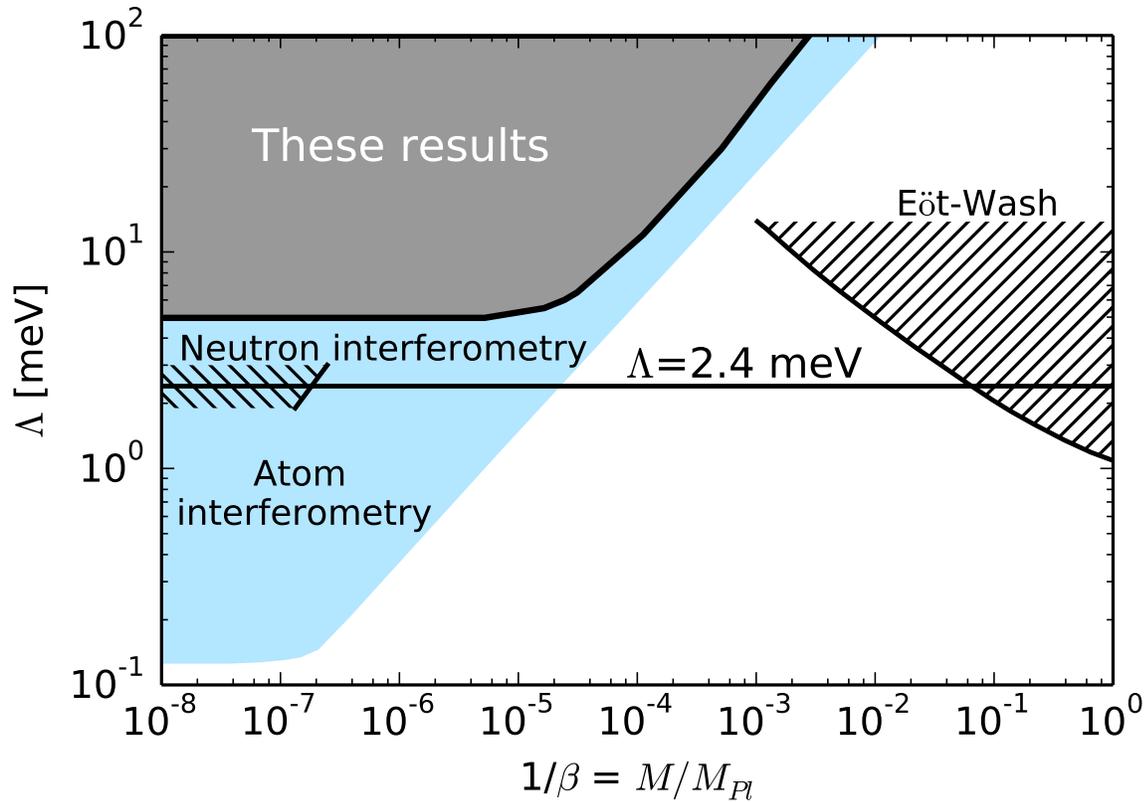


Figure 7.5: Limits on Λ versus $1/\beta = M/M_{Pl}$ for the chameleon model discussed in the text. The 95% CL exclusion limits from this search are denoted by the dark (gray) region. Recent constraints from atom interferometry are shown by the light (blue) region [172, 210]. The horizontal line indicates $\Lambda = 2.4$ meV. Limits from neutron interferometry [211–213] and from the Eöt-Wash torsion balance experiment [175, 204, 217] are denoted by the hatched regions. These limits are shown only in the restricted regions of parameter space considered in Refs [211] and [175].

provide a good fit to the data. At the best fit point, $\chi^2 = 97.8$ for 87 DOF, while for the background only model $\chi^2 = 98.9$ for 88 DOF. For all Λ , the data are consistent with the background-only model at the 95% confidence level (CL). The background-only fits are shown in Figure 7.4, together with the amplitude of a chameleon force at the 95% CL upper limit.

The resulting limits on $1/\beta = M/M_{Pl}$ are shown in Figure 7.5 and compared to existing limits on chameleon interactions. Due to the self-screening of the microsphere at large values of β , these results are not able to constrain forces arising from chameleons for $\Lambda = 2.4$ meV given current backgrounds. However, at values of $\Lambda > 4.9$ meV, the self-screening is reduced, and these data are able to constrain chameleon interactions. These bounds are within a factor of 3 of the best existing constraints from atom interferometry using an entirely independent technique.

7.5 Conclusion

The analysis presented here constrains screened interactions that would produce a force between the cantilever and the microsphere greater than 0.1 fN at separations greater than 20 μm . This search is limited by backgrounds from fixed dipole moments in the microspheres coupling to electric fields caused by contact potentials. One method for reducing such backgrounds is to spin the microspheres by applying an optical [52] or electrostatic torque [220]. It might be possible to anneal the microspheres in situ [132] to increase the rate at which separated charges within the microspheres recombine. Finally, commercially available microspheres made from different materials might have smaller permanent dipole moments. Some combination of these techniques may be used in the future to enhance the sensitivity reached here.

These results provide the first search for interactions below the dark energy length scale using isolated mesoscopic objects separated by mesoscopic distances without an intervening electrostatic shield. This experimental technique is complementary to previous searches and could be sensitive to interactions that have evaded detection to date. The determination of the electric field near the cantilever and measurement of the interaction of electrically neutral silica microspheres with these fields provides important constraints on the expected backgrounds for future searches using similar methods. Future work will feature a search optimized for unscreened Yukawa interactions.

Chapter 8

Search for non-Newtonian Interactions that Couple to Mass

The content of this chapter has been adapted from Reference [29], for which the author was a primary contributor. Apparatus and generic experimental method descriptions have been removed, appealing to Reference [37] for apparatus details, and occasionally Chapter 2 of this thesis for general statements. Text includes supplementary material prepared for journal submission that is not in the original arXiv posting (at the time of writing).

Recall from Chapter 6 that it is customary to modify the inverse square law of Newtonian gravity by introducing an additional Yukawa potential with a length scale λ . The resulting potential between two point masses can be written as:

$$V(r) = -G \frac{m_1 m_2}{r} (1 + \alpha e^{-r/\lambda}), \quad (8.1)$$

with G the Newtonian constant of gravitation, m_1 and m_2 the gravitating masses, r their distance, and α the relative magnitude of the new interaction. A few possible examples of theories that would result in deviations from Newtonian gravity at short-distance were detailed in Chapter 6, although investigating gravity and gravitational-strength forces at short-distance has intrinsic value regardless.

In the present work, we describe the first investigation of the inverse square law in the $1 < \lambda < 100 \mu\text{m}$ range using an optical tweezer in vacuum, where radiation pressure is used to counter the Earth's gravity and to provide the restoring force against which the interaction is compared. As first discussed in [215], the motion of an optically levitated

silica microsphere [1, 7] is studied to infer its coupling with an attractor system in which regions of different mass density (silicon and gold) are alternated on a microscopic scale. To our knowledge, this is the search using the smallest objects to both source and sense a new interaction or modified gravity. So far, experiments probing the micrometer regime have been mainly conducted with greater separations between the source and the test mass, and/or using substantially larger test and source masses. In this study, the separation between microsphere and attractor system, the scale of the test mass, and the attractor system density modulation, are all matched to the length scale λ of the interaction. This results in measurements with broader applicability, including to non-Newtonian potentials that cannot be described by the form in Equation (8.1).

8.1 The Experiment

The microsphere, acting as a force sensor, is isolated from the environment so that its center of mass motion can be reduced to very low effective temperatures [36] in an otherwise room temperature setup. The charge state of the microsphere can be controlled with exceptional accuracy (see Chapter 2) to provide an empirical force calibration and, during inverse square law test measurements, ensure overall neutrality. Directly measuring the force vector on the microsphere (see Chapter 3) provides more dimensions to understand backgrounds and provides sensitivity to the sign of α , in contrast to experiments only sensitive to a deviation from $|\alpha| = 0$ [194–196]. Finally, many methods developed in quantum optics can be applied to this technique in the future, with the potential for substantial advances in an all important problem of experimental physics.

The overall apparatus layout, microsphere trapping, force calibration, charge neutralization, metrology, and the force sensitivity achieved, are described in detail in Reference [37].¹ Briefly, the central part of the system, shown in Figure 8.1, is a $(7.56 \pm 0.19) \mu\text{m}$ diameter silica microsphere [88] trapped in an upward-propagating, single-beam optical tweezer, formed by 1064 nm light focused down to a waist size of $3.2 \mu\text{m}$ by an off-axis parabolic mirror with a focal length of 5 cm. The mass and density of the microsphere are estimated to be $m = 414 \pm 15 \text{ pg}$ and $\rho = 1.83 \pm 0.15 \text{ g/cm}^3$ from a combination of measurements *in situ* for another microsphere from the same batch, following the method in Reference [19],

¹The revised apparatus is similar to that described in Chapter 2 of this work, with major improvements to the auxiliary optics and a modified optical trap using focusing mirrors instead of lens.

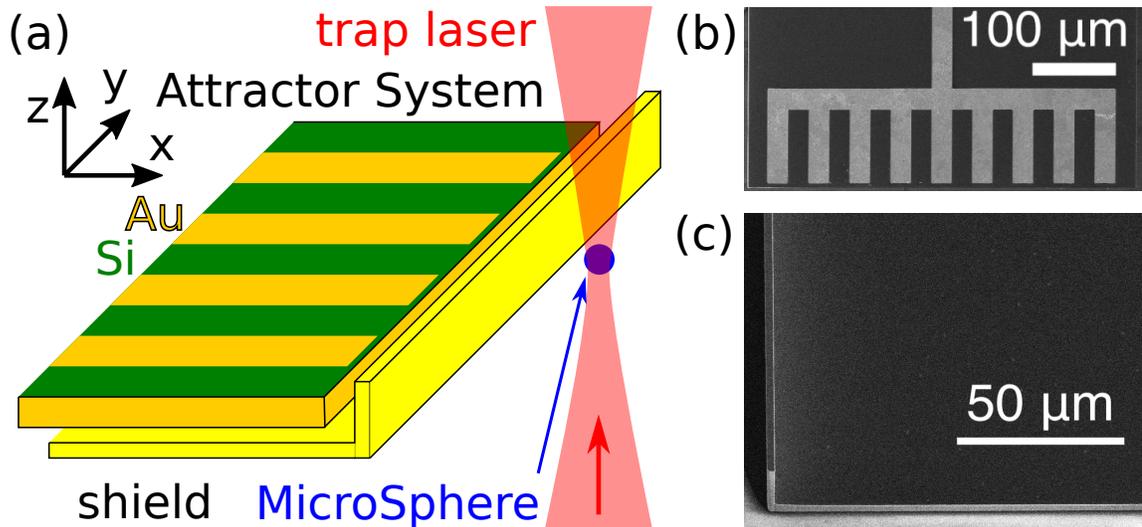


Figure 8.1: (a) Central portion of the experimental setup: a microsphere is trapped in an optical tweezer. A stationary shield centered about the trapped microsphere, with the closest surface within a few microns of the microsphere, and the attractor system is behind it. (b) Scanning electron microscope (SEM) image of the attractor system. The dark (bright) regions correspond to silicon (gold). (c) SEM image of the shield, viewed at a 40° angle to highlight the three-dimensional structure, with the vertical wall to the left.

and manufacturer's specifications [88].

The x and y positions of the microsphere are measured by interfering the recollimated forward-scattered light with a reference wavefront and projecting the result onto a quadrant photodiode (QPD). The z position of the microsphere is measured by interfering the light retroreflected by the microsphere with another reference wavefront, whereby motion along z produces a change in the path length and thus in the phase of the retroreflected light. Both interference measurements make use of heterodyne detection, in which the reference wavefronts are frequency-shifted by -125 kHz relative to the trapping beam. The photocurrent signals are then amplified, digitized, and digitally demodulated. The resulting measurements of the x , y , and z degrees of freedom are used both for real-time feedback control and offline analysis.

The trapping region is surrounded by six identical electrodes resulting in a cubic cavity in which the microsphere is shielded from external electric fields. The electrodes have holes for optical and mechanical access from six directions, and they can be individually biased to control translational and rotational degrees of freedom of the microsphere. This feature

is used to calibrate the force sensitivity of the system by adding a well-defined charge to the microsphere and driving its motion with AC fields applied to the three pairs of opposite electrodes [24–27, 37]. These manipulations are generally done with the attractor system and shield in their retracted position, so that the applied electric field at the microsphere location is well understood and approximately uniform.

Prior to the inverse square law measurements, a neutralized microsphere is driven to rotate at 6 kHz, by coupling a rotating electric field to the permanent electric dipole moment in the microsphere [28, 38] (see Chapter 5). This results in a lower and more consistent force noise. At the 4×10^{-7} mbar vacuum employed here, the microsphere’s angular velocity decays exponentially with a time constant >8 hours [28] in the absence of a driving field and while inverse square law measurements are performed. The natural oscillation frequency of the trapped microsphere is ~ 380 Hz for both x and y , while feedback in the z direction results in a similar trapping frequency (cf. the optical spring constant without the feedback in the z direction corresponds to ~ 30 Hz). Slow drifts in the z position, which may be attributed to changes in the optical path, are corrected at ~ 10 s intervals by an auxiliary measurement performed using a camera-based microscope installed at a side-view port.

The attractor system (Figure 8.1b) is a cantilever device, nanofabricated in silicon and measuring $500 \mu\text{m} \times 475 \mu\text{m} \times 9 \mu\text{m}$ in the x, y, z directions, and supported by a thick silicon handle [69]. The front portion of the attractor system, closest to the trapped microsphere, is patterned with nine rectangular trenches filled with gold, regularly spaced along the y axis with a pitch of $50 \mu\text{m}$, measuring 25 (100) μm in the y (x) direction to create the required density modulation.

Although the attractor system is coated with 150 nm of gold over a 50 nm titanium adhesion layer, a separate shield is employed to further reduce both scattered light and electrostatic backgrounds. The shield (Figure 8.1c) is also nanofabricated in silicon, to obtain an L-shaped cross-section in the $x - z$ plane. The horizontal plane of this device is $350 \mu\text{m} \times 1000 \mu\text{m} \times \mu\text{m}$ in the x, y, z directions, and the vertical wall nearest to the trap is $22 \mu\text{m}$ tall (z) and $\sim 2 \mu\text{m}$ thick (x). The shield, also sputter-coated with 150 nm gold over 50 nm titanium, is maintained stationary during a measurement, while the attractor system scans along the y direction with reciprocating motion. This arrangement is designed to reduce the background from electric field gradients, originating from both a contact potential and patch potentials of the surface of the attractor system [27, 78], as it scans in front of the microsphere. Additionally, the shield reduces backgrounds due to

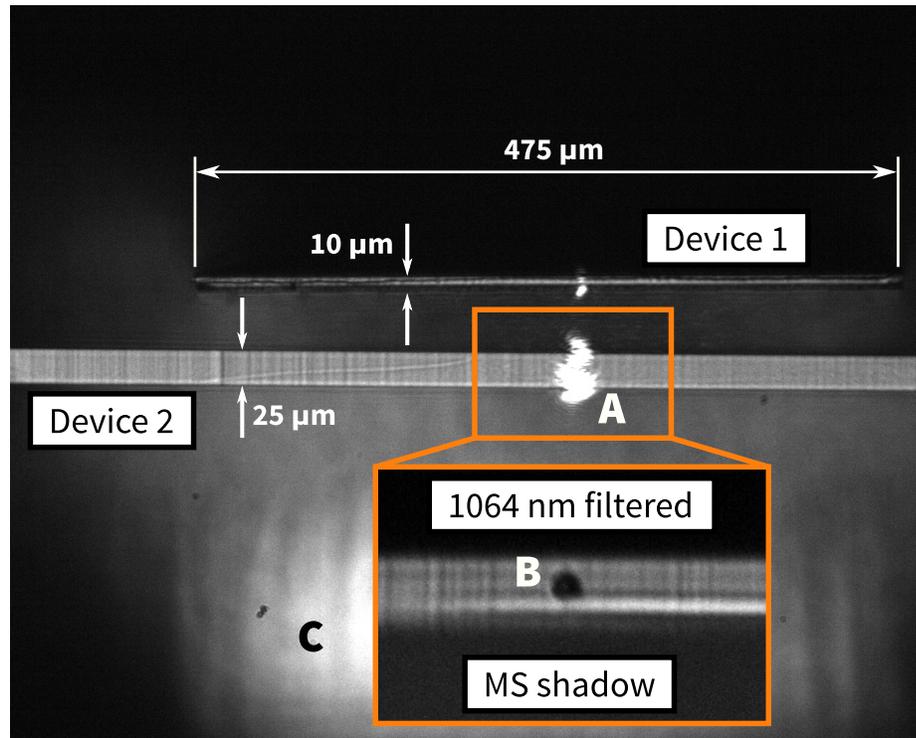


Figure 8.2: Composite image of two nanofabricated devices in the vicinity of the trap, together with a trapped microsphere, captured by the side-view microscope. Device 1 is the attractor system, mounted on the main nanopositioning stage, and Device 2 is the electrostatic shield, mounted on the auxiliary nanopositioning stage. The L-shaped cross-section of the shield, which usually houses the attractor system, is not visible here. In this image, the attractor system is translated vertically from its nominal position and is out of focus, the latter reducing the apparent vertical extent. The main frame of the image shows the scattering of the trapping laser by the microsphere (A), which leads to saturation of several pixels of the camera. The inset is taken with a notch filter that has an optical depth of 6 for 1064 nm light (Thorlabs NF1064-44), in order to demonstrate the shadow of the microsphere (B) blocking the 870 nm illumination of the shield. The bright and blurry region (C) behind the device is caused by a reflection of the illumination from the handle of the attractor system.

modulations of the halo of the trapping beam or other stray light, which mimic minute shifts in the centroid of light on the QPD. An example image from the auxiliary side microscope showing all devices in close proximity to a trapped microsphere is shown in Figure 8.2.²

²This side microscope, detailed in Reference [37], is significantly more sophisticated than that described in Chapter 2, and has been designed for this newer iteration of the apparatus specifically to perform quantitatively accurate alignment of all devices in close proximity to a trapped microsphere.

With all devices in position as in Figures 8.1 and the apparatus calibrated as described, the attractor system undergoes harmonic reciprocating motion with a frequency of 3 Hz and a peak-to-peak amplitude of 202 μm along the y direction, corresponding to ~ 4 full periods of the density modulation. During a 10-s-long measurement, the motion of the microsphere, the position of the attractor system in three dimensions, as well as various power-monitoring photodiodes, and feedback monitors, are synchronously digitized at 5 kHz and stored in a single binary file with timestamps. Environmental variables such as temperature and atmospheric pressure are sampled at a lower rate stored separately. A total integration of 10^5 s is obtained by repeating such 10 s measurements 10^4 times.

8.2 Analysis

8.2.1 The Data Sample

For the 7.56 μm silica microspheres used here, a force sensitivity of $\leq 1 \times 10^{-16}$ N/ $\sqrt{\text{Hz}}$ in the 1 Hz to ~ 50 Hz frequency range is achieved [37]. For neutral microspheres, this performance is also observed when both attractor system and shield are in close proximity, as shown by a typical force amplitude spectral density (attractor systemD) displayed in Figure 8.3, with the closest shield surface at 11 μm from the center of the microsphere. The observed baseline noise is of a statistical nature, and can be integrated for multiple days without encountering an irreducible floor. The 10^5 s data set used here was collected with one microsphere. The distance between the center of the microsphere and the front surface of the attractor system in the x direction is 13.9 μm , and the offset between the center of the microsphere and the center of the attractor system is 4.9 (-15.7) μm in the y (z) direction. The uncertainties and drifts of these parameters over the entire run are about ± 1 μm or less and are specifically shown in Table 8.1. Although the expected sensitivity for this exposure at the noise limit corresponds to $\alpha \approx 1 \times 10^7$ for $\lambda = 10$ μm , the actual sensitivity is limited by backgrounds, which manifest when the attractor system scans. This is illustrated by Figure 8.3, as there are specific frequencies at which a response well above the noise results from the scanning of the attractor system.

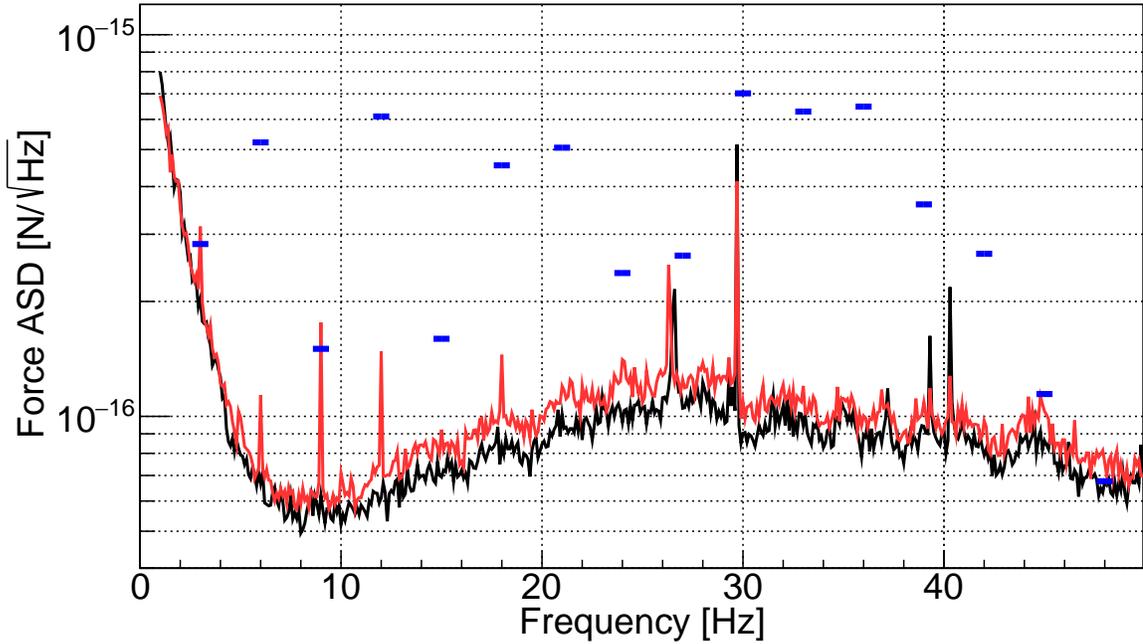


Figure 8.3: Amplitude spectral density (attractor systemD) of the z component of the force on a $7.56 \mu\text{m}$ diameter microsphere. The black (red) curve shows actual data with the attractor system stationary (scanning along y at 3 Hz with $202 \mu\text{m}$ peak-to-peak amplitude). The blue bars show a comparison to an expected microsphere response produced by the potential described by Equation 8.1 with $\alpha = 10^{10}$ and $\lambda = 10 \mu\text{m}$. The data displayed here is the average of 100 distinct 10-s integrations.

8.2.2 Backgrounds

Backgrounds can originate from several sources. Interactions between electric field gradients induced by the attractor system and the electric dipole moment of the microsphere, estimated to be $10^2 - 10^3 e \cdot \mu\text{m}$ [25, 28, 38] with e the fundamental charge, are expected in all directions, with different levels of attenuation from the shield. In the xy -plane, backgrounds may also arise from small variations in the halo or stray light, driven by the scanning motion of the attractor system. In the z direction, this background is expected to be substantially smaller as the shield blocks the attractor system in the image plane of the retroreflected photodiode, although couplings between z and $x-y$ at the 20% level exist. The $x-y$ components of the background observed at individual frequencies are as large as $1.5 \times 10^{-15} \text{ N}$, which is equivalent to $\alpha \gtrsim 10^{11}$ for $\lambda = 10 \mu\text{m}$.

While the three dimensions can eventually be used to provide a more sensitive measurement, the asymmetry in the current background levels makes the measurement along z substantially superior for the present analysis. By modeling the system with a finite element method, it was found that a contact potential difference of ~ 50 mV between the attractor system and the shield can account for backgrounds in z at the observed order of magnitude. Backgrounds from patch potentials on the attractor system are found to be subdominant because of strong attenuation from the shield.

8.2.3 Signal Model

In order to conduct a search for non-Newtonian forces that couple to mass, a signal model is built from mesh calculations of the force between the attractor system and microsphere as a function of their relative displacement, for various length scales λ . This is done by first computing the analytical expression for Yukawa-modified gravitational force between a perfect homogeneous sphere and a point mass. The attractor system is then discretized and treated as a regular grid of point masses, with the total force between the microsphere and attractor system being the sum of the constituent analytical results. An example result of this calculation for the $d = 7.5 \mu\text{m}$ microspheres is shown in Figure 8.4, at four different attractor-face to microsphere-center separations, and for $\lambda = 10 \mu\text{m}$ and $\alpha = 1$. The expected force along all three axes is shown, including the “DC” component from the interaction between the bulk of the attractor system and the microsphere, as well as the expected modulation. The signal scales proportionally to α , which is the parameter of interest in the statistical inference procedure.

This model is subsequently sampled by the measured position of the attractor system during each 10 s run to generate the expected force on the microsphere as a function of time. The microsphere response is expected to have different amplitudes at several integer multiples of the fundamental frequency f_0 of the attractor system motion, as shown Figure 8.3.

As some background sources, such as vibration, are expected to affect mainly the fundamental frequency, we exclude 3 Hz and use only harmonics which contain an expected signal stronger than that of 3 Hz. Also excluded are the 6 Hz, 2^{nd} harmonic, because of a potential background arising from non-linearities in the system, and the 30 Hz, 10^{th} harmonic, because of an unidentified large spectral feature at 29.7 Hz (also present with the attractor system stationary). Therefore, the search is performed using the harmonics at 12, 18, 21,

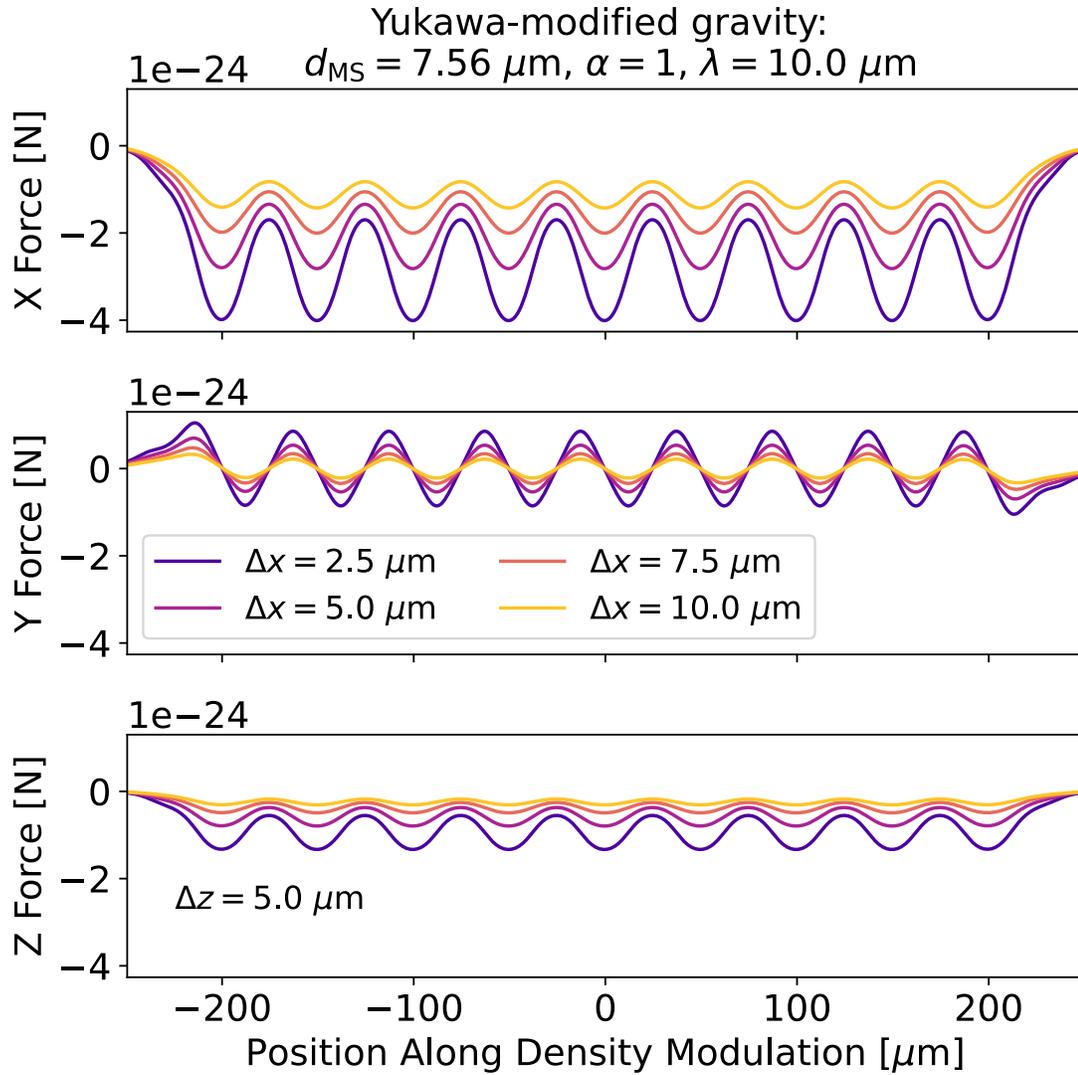


Figure 8.4: An example result of the mesh calculation for the Yukawa-modified gravitational force with $\alpha = 1$ and $\lambda = 10 \mu\text{m}$, in all three orthogonal axes, on a $d = 7.5 \mu\text{m}$ and sourced from the attractor system described. The expected force is shown for a few distinct horizontal separations, and a single fixed vertical separation, included to demonstrate exciting a response in the axial direction.

33, 36, and 39 Hz. In addition to the amplitude information, the phase of the expected signal relative to the attractor system motion is incorporated for all those harmonics.

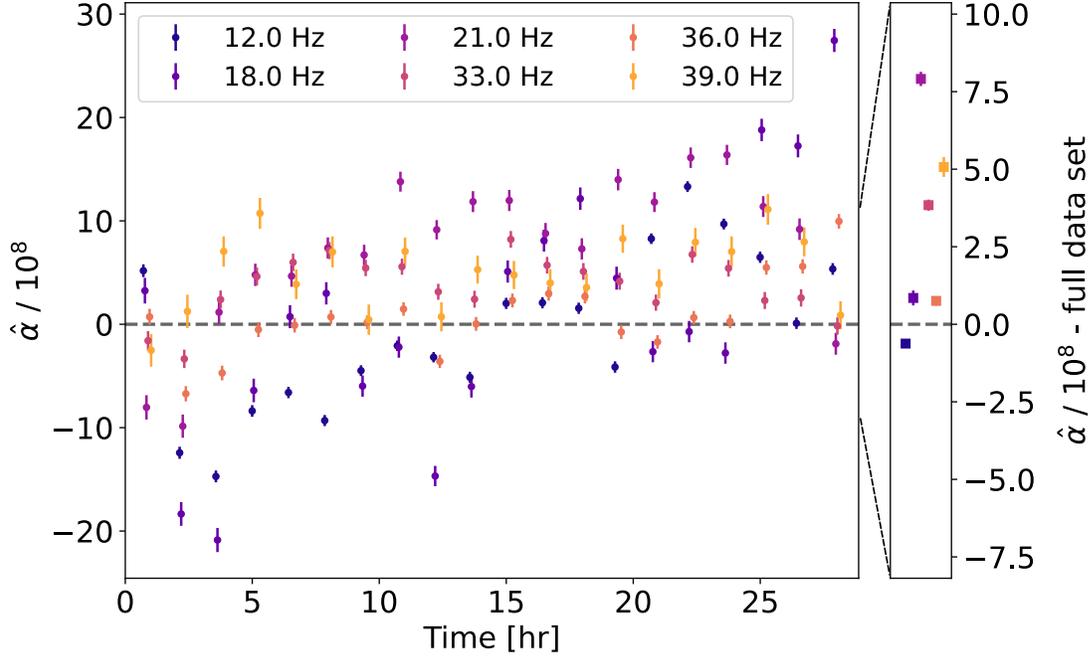


Figure 8.5: The single harmonic maximum likelihood estimator (MLE) $\hat{\alpha}_i$ for $\lambda = 10 \mu\text{m}$ as a function of time for the six harmonics used in the analysis. Each harmonic f_i is evaluated separately taking into account its own phase response and noise level. Here, each estimation of $\hat{\alpha}_i$ comes from 5000 seconds of data. The error bars represent 95% confidence intervals about the MLEs. The panel to the right shows the MLE for each harmonic, integrating over the entire data set (note the expanded vertical scale).

8.2.4 Statistical Procedure

For each harmonic f_i , we define the following likelihood function,

$$\mathcal{L}_i(\alpha, \lambda) = \prod_j \left(\frac{1}{\sqrt{2\pi\sigma_{ij}}} \right)^2 \exp \left\{ -\frac{[\text{Re}(F_{ij} - \tau_i(\alpha, \lambda, \vec{x}_j))]^2}{2\sigma_{ij}^2} - \frac{[\text{Im}(F_{ij} - \tau_i(\alpha, \lambda, \vec{x}_j))]^2}{2\sigma_{ij}^2} \right\}, \quad (8.2)$$

where F_{ij} is the value of the single-sided Fourier transform of the z -force (normalized to units of $\text{N}/\sqrt{\text{Hz}}$) in the frequency bin corresponding to f_i , $\tau_i(\alpha, \lambda, \vec{x}_j)$ is the value of the Fourier transform of the expected signal force in the same frequency bin for a given α and λ and attractor system displacement \vec{x}_j , σ_{ij} is the standard deviation of the Gaussian white noise in the frequency bin for f_i , estimated from 10 neighboring sidebands, j indexes the

10^4 , 10-second-long, data-files, and $\text{Re}()$ and $\text{Im}()$ are the real and imaginary components of the complex-valued Fourier transforms, respectively.

Specifically, σ_{ij} is calculated as follows for a single harmonic, f_i , and continuous integration, j , from the observed variance of neighboring sidebands f_k :

$$\sigma_{ij}^2 = \frac{1}{2N_{\text{sb}}} \sum_{k=1}^{N_{\text{sb}}} [\text{Re}(F_{kj})^2 + \text{Im}(F_{kj})^2], \quad (8.3)$$

where $N_{\text{sb}} = 10$ is the number of sidebands, F_{kj} is the value of the Fourier transform of the z -force in the frequency bin corresponding to the sideband f_k , and the factor of $(1/2)$ yields the expected uncertainty for either the real or imaginary component independently.

Each $\mathcal{L}_i(\alpha, \lambda)$ can be used individually to provide the maximum likelihood estimator, $\hat{\alpha}_i$, for each harmonic, as shown in Figure 8.5. It is confirmed that the measured signals are background-like and not due to a novel interaction by observing that the amplitudes extracted for each selected harmonic do not exhibit the expected ratio from the signal as shown in Figure 8.3. In addition, the expected time-invariant behavior is not found in the data.

Due to the different levels of background in different harmonics, each is treated independently in the statistical procedure and combined in an approach following [221]. This utilizes the fact that a gravity-like force should be present in all harmonics, increasing the sensitivity when backgrounds are correlated differently than the expected signal.

Since the described experiment is sensitive to the direction of the force, upper limits can be set separately on positive and negative values for α . Harmonics with $\hat{\alpha}_i > 0$ are used to constrain an upper limit on $\alpha > 0$, while those with $\hat{\alpha}_i < 0$ constrain $\alpha < 0$, following the procedure in [221]. A test statistic for harmonics f_i with $\hat{\alpha}_i > 0$ is defined as

$$q_{\alpha,i} = \begin{cases} -2 \log \left(\frac{\mathcal{L}_i(\alpha, \lambda)}{\mathcal{L}_i(\hat{\alpha}_i, \lambda)} \right) & \alpha \geq \hat{\alpha}_i \\ 0 & \alpha < \hat{\alpha}_i \end{cases}, \quad (8.4)$$

where a nearly identical function is defined for harmonics with $\hat{\alpha}_i < 0$, but with the conditions flipped appropriately for the change in sign. The final test statistic used to establish upper limits on alpha is simply the sum over all harmonics, $q_\alpha = \sum_i q_{\alpha,i}$, and is profiled independently for $\alpha > 0$ and $\alpha < 0$. For this work, the entire procedure was completed

with three completely independent analysis frameworks, in order to provide a level of cross-validation.

The method introduced above was thoroughly investigated by injecting artificial software signals on top of actual experimental noise. Data sets with a total length of 10^4 seconds were used, in which the relative positions of microsphere and attractor system are nearly the same as in the primary measurement, but with no scanning motion and hence with no signal or background. This was done repeatedly for a range of both parameters, and an upper limit was estimated for each unique data set. This process validates the analysis, quantifying the deviation from Wilks' theorem [221] (and the expected χ^2 distribution), and finding the critical values corresponding to the 95% CL upper limit. In a separate process, constant and time-varying backgrounds were added together with a simulated signal, testing various scenarios and demonstrating that the procedure is robust against under-coverage.

8.3 Results

For values ranging from $\lambda = 1 \mu\text{m}$ to $\lambda = 100 \mu\text{m}$ the results are shown in Figure 8.6. The proximity of the upper limit on $|\alpha|$ for both directions implies that the background is of the same order of magnitude in the most sensitive harmonics. This provides a degree of robustness against possible cancellations with backgrounds and signal in opposite directions, as the expected signals (in terms of $\hat{\alpha}$) should be consistent between harmonics. The limit is constant for $\lambda \gtrsim 10 \mu\text{m}$, and degrades exponentially as the length scale becomes shorter than the separation between the attractor system and the microsphere.

The main systematic uncertainties are summarized in Table 8.1. The dominant effect is the uncertainty in the distance between the attractor system and the microsphere in the x direction. Further significant contributions come from uncertainties in the phase response of the microsphere as measured in the calibration procedure, uncertainty about the attractor system thickness, as well as drift of the amplitude response. microsphere properties, distances in x and y , and alignment stability and accuracy of the attractor system movement have been found negligible.

The main limitation of the investigation presented here is the existence of backgrounds originating from electrostatic interactions, stray light modulated by the attractor system motion, and vibrations of components inside the vacuum chamber. As mentioned above, the interaction between the microsphere and an electric field gradient arising from a contact

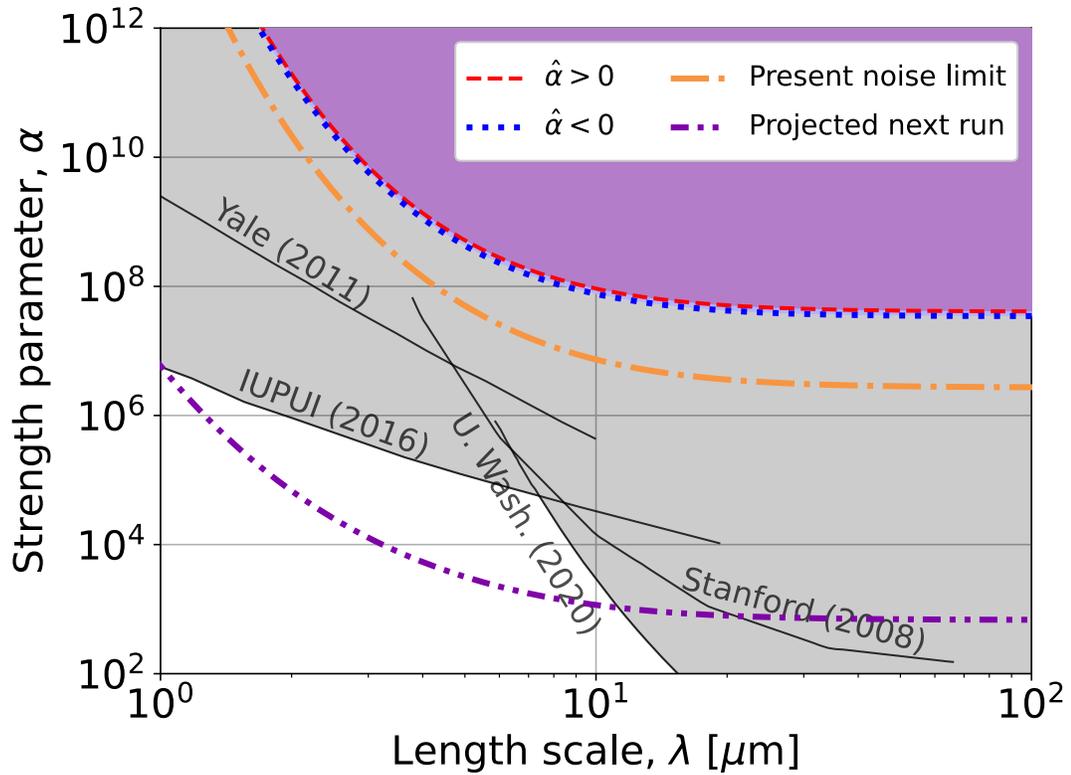


Figure 8.6: Limit curve in the $\alpha - \lambda$ parameter space. The region above and to the right of the red and blue lines indicates the parameter space excluded by this experiment for positive and negative α , respectively, with a 95% confidence level. The gray region shows the parameter space covered by previous searches [194–196, 198]. The background-free sensitivity for this run, using the current attractor system-microsphere separation, noise conditions, and integration time is shown by dash-dotted orange line. In addition, the projected sensitivity for the next run, given the improvements outlined in the text, is shown by the dash-double-dotted purple line. This assumes a noise floor of 1×10^{-18} N/ $\sqrt{\text{Hz}}$, attractor system-microsphere separation of 7.5 (-5) μm in the x (z) direction, and an integration time of 30 days. The reach could be extended further by using larger microspheres [34] or smaller separations [37].

potential can be calculated to provide an adequate model for the electrostatic backgrounds. This model can be constrained and validated by a three-dimensional scan in which the attractor system is placed in different regions around the microsphere, as in Reference [27]. The interaction will then be minimized by applying a bias between shield and attractor

system to null the contact potential. The stray light background is being investigated with a combination of measurements and ray tracing analysis, to inform the design of light baffles, in parallel with the development of a new multi-pixel sensor to replace the QPD, that will provide discrimination between actual shifts of the microsphere and changes in the halo. Finally, critical optical components inside the vacuum chamber are being stiffened to minimize vibrations. Those efforts, along with the multi-harmonic analysis technique presented above, are expected to push the experiment into the noise dominated regime for the next run. The improvement in sensitivity from those changes, without altering other parameters, can be seen in Figure 8.6.

Beyond background suppression, an improvement in sensitivity in terms of noise reduction is targeted with the next iteration of the experiment. It is important to emphasize that the force sensitivity of the system is limited by pointing fluctuation of the trap beam [37] and not by shot noise or residual gas damping as already demonstrated in References [34, 60]. Therefore, an enclosure of the input external to the vacuum chamber, possibly replacing air with helium to lower the refractive index and hence the effects from its fluctuation, along with the stiffening of mechanical components, are expected to lead to substantial reduction of the noise floor, down to $\sim 1 \times 10^{-18}$ N/ $\sqrt{\text{Hz}}$ as demonstrated in Reference [34]. In addition, a significant gain in sensitivity will be achieved by changing the position and proximity of the attractor system, which, in the current run was limited by misalignment of the attractor system and the electrostatic shield. The projected sensitivity assuming the lower noise floor, 7.5 (-5) μm separation in the x (z) direction, and an integration time of 30 days is shown in Figure 8.6.

Table 8.1: List of systematic uncertainties.

Effect: ϵ	$\Delta\epsilon$	$\Delta\alpha/\alpha$
Drift of amplitude response	10%	10%
Attractor thickness	1 μm	11%
Phase response	~ 0.1 rad	12%
Distances in Y	< 0.2 μm	$< 3\%$
Distances in Z	< 0.9 μm	$< 6\%$
Distances in X	1.5 μm	30%
microsphere weight	15 pg	3.5%

8.4 Conclusion

We have described the results of the first experiment searching for non-Newtonian forces which couple to mass using optically levitated test masses. The effects observed in the data are not consistent with a new interaction, and the result is interpreted in terms of upper limits on the Yukawa parameter α . These are $\alpha > 9 \times 10^7$ and $\alpha < -8 \times 10^7$ with 95% confidence level at $\lambda = 10 \mu\text{m}$. The length scales involved in the experiment, in terms of dimensions of the test masses and feature size of the source of the interaction, and the separation between the two, are, for the first time, all similar to the characteristic length scale being probed. Therefore, this method provides a more robust test that applies also for interactions that cannot be parameterized with a Yukawa potential. Substantial improvements in sensitivity are expected for the next round of measurements.

Conclusions

Optical levitation and trapping is a technique that leverages the interaction between a focused laser beam and a dielectric particle. By virtue of both the reflection and deflection of incident photons, particles with an index of refraction larger than that of the surrounding medium are attracted to the focal point of the laser. The radial confinement is easily understood in the paraxial limit, where motion of the trapped particle away from the beam center induces deflections of the trapping beam, which necessarily generate optical restoring forces on the trapped particle due to momentum conservation. Axial confinement can then be achieved with the gradient forces associated to a highly focused laser beam (numerical aperture $\text{NA} \gtrsim 0.5$), or by balancing the axial radiation pressure force against another force such as gravity, either of which results in fully three-dimensional trapping about some equilibrium position.

As was alluded, displacement of an optically trapped particle about its equilibrium position scatters the trapping beam. The scattering is necessarily associated to an optical restoring force, equal in magnitude and opposite in direction to whatever force initially induced the displacement. This leads naturally into force-sensing applications, where measurements of the scattered light can be used to infer any forces applied to the trapped particle. The length scales associated to typical optical traps, order of magnitude $\sim 0.1 \mu\text{m}$ to $10 \mu\text{m}$, present certain advantages for studies of new physics.

In particular, we consider tests of the universal law of gravitation and searches for fifth forces. This is motivated by our limited understanding of gravity at submillimeter length scales, as well as some outstanding problems in cosmology whose answers may be found at these distance scales. Assuming General Relativity is an accurate description of the dynamics of gravity, and we also trust our observations of the cosmic microwave background and baryon acoustic oscillations, then there is significant implied energy content in the universe that has yet to be understood or observed experimentally. This has come to be called this

dark energy. It is possible to derive a length-scale associated to this missing energy density, which is approximately $100\ \mu\text{m}$. At distances below this length scale, new interactions mediated by new particles and fields or modifications to the dynamics of gravitation itself can be expected, suggesting that precise characterizations of gravity in this regime have intrinsic value.

In this thesis work, I have presented the development and use of a vertically-oriented optical tweezer to trap silica microspheres, stabilized axially by gravity and operated in vacuum. The trap was generated by a fiber-optic laser source with wavelength $1064\ \text{nm}$, focused to a $\sim 3\ \mu\text{m}$ spot with a numerical aperture of $\text{NA} = 0.12$. All of the measurements and techniques described made use of one of two types of silica microspheres, sourced from two different manufacturers and with radii of $2.35\ \mu\text{m}$ or $3.76\ \mu\text{m}$. The focal point of the laser, the nominal position of the optical trap, was enclosed by six identical shielding electrodes in order to mitigate the influence of stray electric fields. The electrodes could be individually biased in order to generate known forces and torques on a trapped microsphere, which was particularly useful for calibrating the system.

Calibration of a microsphere's translational response to external forces was accomplished by applying a known oscillating force to a charged microsphere, while simultaneously measuring the amplitude of the response. The charge state of the microsphere was then modulated with flashes of UV photons, which ejected individual electrons photoelectrically, changing the amplitude of the microsphere response in unit steps with size proportional to the elementary charge. This provided an empirical calibration of the system based only on measured quantities and a single fundamental constant, avoiding unnecessary assumptions and maintaining a degree of robustness.

Once calibrated, an optically trapped microsphere was used for force sensing, simply by observing the translational motion in response to a chosen excitation. The capability of this force sensing was demonstrated with the development of a technique referred to as force-field microscopy. A nanofabricated mechanical device of microscopic dimensions was introduced in proximity to a charged microsphere. The device was positioned over a three-dimensional array of points in the vicinity of the trap, and was biased with a voltage and used as an excitation electrode. At each position within a $10 \times 10 \times 10$ grid of relative positions between the electrode and microsphere, the electrode was biased with an AC voltage and measured the response of the microsphere along three orthogonal degrees of freedom. The result is a force-field: i.e. a three-dimensional array of points, with a three-dimensional measured

force response at each point within the array. The measurement was compared to a finite element analysis of the expected force around the biased electrode, yielding good agreement between observation and expectation.

Force-field microscopy was then applied to characterize a possible source of electromagnetic background forces in precision short-distance force sensing. Optically trapped microspheres have been found to have permanent electric dipole moments, which couple to electric fields and electric field gradients. In the searches for new physics, gravity-like interactions were sourced by a mechanical device in close proximity to a trapped microsphere. Metallically coated silicon and gold devices were used as sources, whose sputtered titanium/gold surfaces are known to have patch potentials, or microscopic domains within the coating that have measurably different contact voltages. The random nature of these patch potentials induce large and randomly distributed electric fields and gradients near to surfaces on which they lie. Using force-field microscopy, the response of a charged microsphere to the patch potentials on the end of an aforementioned source mass device were characterized. The scaling of the root-mean-square force was found to be consistent with a simple toy model of patch potentials, and the product of the characteristic patch size and expected variance of the patch potentials was computed.

When testing forces that couple to mass, it is clear that accurate and precise knowledge of the test and source masses is critical to making quantitative measurements and inferences. As such, a technique to measure the mass of optically trapped microspheres *in situ* was developed, and together with radius characterization of individual spheres using electron microscopy, the density of the levitated microspheres was inferred. The mass was first measured by applying an axial electrostatic force to a charged microsphere whilst actively modulating the optical power of the trapping beam to maintain a constant suspension force balancing Earth's gravity acting on the microsphere. By extrapolating the measured relation between applied electric strength and measured optical power to the point of zero optical power, the algebraic relation $qE = mg$ was used to infer the mass m , given that q was known by virtue of the single-electron-precision control of the microsphere charge state, E was the extrapolated value of the electric field strength whose scale is known by construction, and g was known. Mass measurements were performed on several microspheres in multiple vacuum pressure regimes and with both signs of charge states, all yielding consistent results. The technique was demonstrated for ~ 80 pg and ~ 420 pg silica microspheres.

A select few microspheres whose masses were measured were then individually collected

from the optical trap with a nanofabricated silicon device. These microspheres were transferred to a scanning electron microscope where their apparent sizes were compared to calibrated diffraction gratings in order to provide an accurate measurement of their radius. With the mass and radius characterizations of specific and individual microspheres, their densities were calculated directly, assuming homogeneity. It was found that for certain types of silica microspheres, distinguished by their manufacturer source, their measured density (1.55 g/cm^3) was significantly less than that of both fused silica and the value provided by the manufacturer itself ($\sim 1.8 \text{ g/cm}^3$). The density of a second type of microsphere, also with a larger radius, was found to be consistent within measurement uncertainty with the value provided by the manufacturer (1.86 g/cm^3). These results and inherent discrepancies suggested that precise characterization, beyond assuming datasheets are correct, of the levitated test mass properties is absolutely critical to any application dependent on those properties.

While the permanent electric dipole moment of these silica microspheres is a possible source of background forces, it also provides a handle by which to control the rotational degrees of freedom of a trapped microsphere, given that a dipole moment tends to align with the direction of an externally applied electric field. By driving four of the six shielding electrodes, a coplanar set, with four properly phased sinusoidal voltages, an electric field with a constant magnitude and a rotating orientation was applied to trapped microspheres. The resulting torque induced a trapped microsphere to rotate in the plane of the driving electrodes. This rotation was measured using the residual birefringence of the microspheres: as they rotate, they couple some of the linearly polarized trapping beam light into the orthogonal polarization, at a rate of twice their rotation frequency. The cross-polarized light was separated downstream of the optical trap and collected on a photodiode to recover a measurement of the microsphere's angular position. Various dynamics of the rotational motion were then exhibited.

A rapidly rotating microsphere was released from the electric field and allowed to decelerate in the presence of drag torque from residual gas at high vacuum, and the resulting exponential decay of angular velocity was observed. With the electric field left continuously driving, moderate pressures of nitrogen gas were introduced to the experimental chamber such that the equilibrium orientation of the microsphere's dipole moment lagged behind the orientation of the electric field due a drag torque comparable in strength to the driving electrical torque. Again with a continuously driving electric field, phase modulations of the

cross-polarized rotational signal were observed and further identified as librational motion. In the frame co-rotating with the applied electric field, the orientation of the electric dipole moment of a trapped microsphere oscillated harmonically about the direction of the electric field, yielding phase modulations of the rotation signal. The scaling of the fundamental libration frequency as a function of electric field strength was then characterized to infer the magnitude of the electric dipole moment of a spinning microsphere. Finally, gyroscopic precession was induced by rapidly changing the orientation of the electric field by 90° .

The spindown of a rotating microsphere in the absence of electric field and the phase lag induced by surrounding gas were applied to operate the apparatus as spinning-rotor vacuum gauge. The kinetic theory of gases implies a direct proportionality between the torsional drag coefficient and the residual gas pressure, so that measurements of the torsional drag are pressure measurements. At the system's base pressure, limited by residual water content, spindowns from a 100 kHz angular velocity were found to have exponential decay times of ~ 2000 s, implying pressures of a few 10^6 mbar. The accuracy of the gauge was cross-checked against a capacitance manometer by inducing equilibrium drag with non-negligible pressures of six different gas species. Geometric coefficients related to the surface properties of the microspheres were found to be consistent with theoretical expectations, and across multiple microspheres from the same lot. Given that the torsional drag depends specifically on the gas species, this type of spinning-rotor gauge could be employed for species identification if used in tandem with a species-independent gauge like a capacitance manometer, with applications such as process gas control for nanofabrication.

With the myriad of metrological tools discussed, the apparatus was applied to investigate fundamental physics, specifically searching for new interactions that may manifest at short-distance. An intriguing example of possible new physics are screened interactions, where the strength of the underlying quantum field, and thus the mass and effective range of any force-mediating particle excitation of this field, depends on the nearby matter content. Non-negligible matter content, such as is found nearly everywhere on Earth, tends to suppress the field strength and significantly limit the range and strength of the interaction. Out in space, the field is no longer screened and can mediate long-range interactions that may help explain dark energy, the universe's apparent missing energy content.

A particular screened interaction called the chameleon field was investigated by using a trapped microsphere as a test mass, where the evacuated region in the vicinity of the optical trap allows for a non-negligible chameleon field strength. It turns out that test masses

are actually sensitive to the gradient of the chameleon field. These were induced with a high aspect ratio mechanical device, akin to a silicon cantilever, in vacuum, so that there was a large gradient near the matter-vacuum interface, given that the mechanical device significantly screens the chameleon field and the vacuum environment does not. Furthermore, the gradient was modulated by driving the mechanical device toward and away from a trapped microsphere, allowing for a precise characterization of any anomalous forces at the chosen driving frequency. A maximum likelihood parameter inference found our data to be consistent with no new interactions, and instead, dominated by an electrostatic background due to a contact potential on the mechanical device interacting with the electric dipole moment within a trapped microsphere.

While probing specific theories and models clearly has value to the physics community, more general characterizations of gravity-like forces that manifest at short-distance are applicable to broad classes of theories. Under this paradigm and assuming tabletop experiments operate in the classical (non-relativistic) regime, it is customary to write down a modified Newtonian gravitational potential of the form $V(r) = (Gm_1m_2/r)(1 + \alpha e^{-r/\lambda})$, which includes the usual $(1/r)$ term, plus a Yukawa-type modification with a strength α relative to gravity and effective range λ . Proposed extensions to our current understanding of physics, such as supersymmetry, string theory, modified General Relativity, massive gravitons, and more, each imply possible values of α and λ that one could measure with an appropriate apparatus. This $\alpha - \lambda$ parameter space has been probed over a huge range of scales, making use of observations of the orbits of planets within our solar system, as well as microscopic resonators and torsion pendulums. It should be noted that this is only a particular parameterization, and other functional forms have been proposed.

Optically levitated microspheres were applied to this investigation as part of the conclusion of this thesis work. With a now well-characterized microsphere acting as a force sensing test mass, gravitational interactions including possible Yukawa modifications were probed. The interactions were sourced with a custom nanofabricated device consisting of alternating regions of gold and silicon with a sputtered gold coating to ensure an isopotential surface and to minimize possible electrostatic backgrounds. The source mass was made to oscillate laterally in front of the test mass at a fixed separation such that the microsphere was exposed to multiple regions of the alternating mass density. This induces an oscillating gravitational response of the test mass at integer multiples of the fundamental driving frequency. To further suppress electrostatic backgrounds, a custom nanofabricated shield device, consisting

entirely of silicon with sputtered gold coating, was held stationary between the test and source masses for the duration of the measurement.

This search for new interactions with optically levitated microspheres was ultimately background limited, due to both contributions from anomalous scattered light contaminating the measurement, as well as real background forces from contact potentials interacting with the trapped microsphere's electric dipole moment. A simultaneously cautious and robust statistical inference procedure was proposed and implemented, making use of the test mass's independent response at different frequencies and the highly-correlated harmonic structure of the expected signal. This procedure set upper limits of the order $\alpha \lesssim 10^8$ for $\lambda \sim 10 \mu\text{m}$. Notably, **this result represents the first use of optically levitated microspheres in a search for new interactions that couple to mass at submillimeter range**, where searches in this range are highly motivated by the dark energy length scale of $100 \mu\text{m}$. This technique is entirely complementary to existing methodologies, and is subject to an entirely different set of systematic effects, so that in the wake of potential discoveries, reliable cross-checks are available and attuned.

Summary for Non-Scientists

The “Why”

From a physicist’s perspective, everything that happens in the universe is governed by some physical law, either one we intimately understand, or one we are working toward understanding. This is true for the largest and the smallest scales, ranging from galaxies to microscopic system. Some physical laws have observable consequences at all of these scales, most notably the universal law of gravitation. Gravity, however, presents an interesting and unique problem: the fundamental equations that tell you how gravity should behave, i.e. Einstein’s General Relativity, are vastly different than any of our other equations that describe such things as how atoms are held together, or how photons propagate in space. Our understanding of gravity can be tested with astronomy and cosmology, by looking out into space and seeing what all the stars and galaxies are doing. When we do this, we observe motion and distributions of matter that can be explained reasonably well by General Relativity, with the caveat that the reigning description includes two mysterious contributions to the mass and energy density of the universe: dark matter and dark energy. There is no current consensus on exactly what dark matter and dark energy are, although we can infer their presence based on these cosmological observations. Yet, their nature has significant consequences for our understanding of the origin and the ultimate fate of the universe.

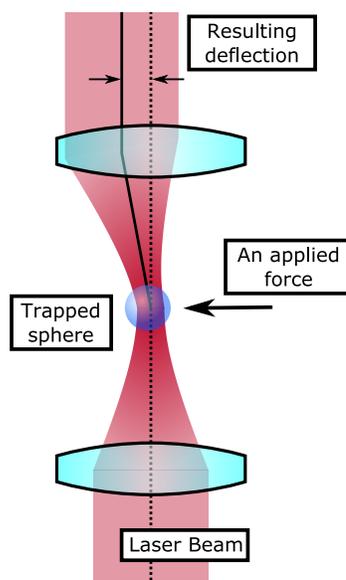
New theories or modifications to General Relativity that attempt to explain these mysterious contributions often have specific observable consequences, for example an increase in the expected strength of gravity when two objects are close together, or far away from one another. However, there is a complication that gravity is generally very weak compared to other fundamental forces like electromagnetism, such that precise investigations of gravity have historically required large masses like planets and stars, since the expected strength is

proportional to the masses of the interacting objects. Gravity is not well understood when the separations are small, of microscopic proportions less than about 1 millimeter. In this short-distance regime, there are number of testable theories explaining dark matter or dark energy that would result in anomalously large or anomalously small gravitational forces. If forces at this scale can be measured, it is possible to learn something about the nature of gravity and General Relativity. The work described in this thesis seeks to explore this regime of gravity.

The “How”

The interaction of light and matter is critical to a variety of different physics experiments and techniques. Particularly useful is laser light, which can be focused down to extremely small and bright spots, much like sunlight can be focused with a magnifying glass. Small particles that bend light, such as water droplets or microscopic glass spheres, can then be confined within the focused spot of a laser beam. This is the basis of optical tweezers, the technique around which this thesis work has been built, wherein a microscopic glass sphere is held to a specific location in three-dimensional space, and suspended against Earth’s gravity with the radiation pressure applied by the laser beam.

Once confined to an optical tweezer, subsequent motion of a trapped sphere scatters the laser light used to form the tweezer. For example, if the sphere moves slightly to one side of the center of the tweezer, it bends the focused laser beam in the direction in the direction of motion. If the scattered laser beam is monitored with light-sensitive detectors (similar in principle to digital cameras), then the motion of the sphere is also monitored. This leads us naturally into applications of force-sensing, in which external forces pushing on an optically trapped sphere cause it to move, which in turn deflects the laser beam and can thus be measured. It possible to derive a direct relation between the applied force and an observed deflection, both with computer-driven simulations as well as empirical calibrations. In this fashion, an optically trapped sphere can be used as a force sensor, specifically useful for short-distance force sensing given the microscopic size of the glass spheres that can be trapped with laser light. This framework is detailed schematically in the figure below.



A sphere is shown confined by an optical tweezer. An applied force pushes the sphere to one side, resulting in a deflection of the laser beam used to form the optical tweezer. The deflection can then be used to monitor the motion of the sphere, and eventually measure the applied force.

- Force-field Microscopy

A number of applied physics techniques were developed in this thesis work making use of an optically trapped sphere acting as a force sensor. The first of these we call force-field microscopy. Consider for a moment the concept of a van de Graaff generator, the classic physics demonstration wherein static electricity is generated on a large metallic sphere. If you get close to the van de Graaff generator, especially if you touch it, the electrical forces produced cause your hair to stand up. Clearly, if you stand far away, your hair doesn't stand up because the electrical force is smaller the farther away you are. This same principle holds for a variety of different forces: the strength (and direction) of a force depends on your position relative to the source of that force. If you can measure this position-dependence, it is possible to better understand the origin and fundamental nature of the force being sensed.

This is the basic principle of force-field microscopy. We demonstrate this technique by placing a source of electrical forces in proximity to an optically trapped sphere. The electrical forces push and pull on the sphere depending on the position of the source relative to the sphere, e.g. the force is smaller for greater separations between the two. By measuring the deflections of the laser beam that result from these pushing and pulling forces, it is possible

to reconstruct the strength and direction associated to the force. Piecing together results from the many different positions tested, we measure the electric field around the source of electrical forces. It is also possible to simulate the expected electric field, and our measured data very closely matches the simulation result. An example of this measurement can be seen in Chapter 3, specifically Figure 3.5.

- Mass and Radius Measurements

Generally speaking, proper inference of any forces sensed by a trapped sphere requires detailed knowledge of the properties of the sphere, in particular its mass and radius. For example, the more massive the sphere is, the stronger a gravitational force would be. As such, we have developed a technique for the precision mass and radius measurement of individual optically trapped spheres. This is accomplished with an electrical balance of sorts, wherein the inherent optical suspension provided by the optical tweezer itself is supplemented with an electrical suspension. By simultaneously varying the strength of the (known) electrical and (unknown) optical suspension forces, whilst measuring the strength of the electrical suspension, it is possible to infer the mass of the sphere by equating these suspension forces to the gravitational force pulling down on the trapped microsphere. We found that the typical sphere used in our trap weighs approximately 100 picograms, or about one ten-billionth of the mass of a dollar bill. The measurement scheme and two example measurements are shown in Figures 4.1 and 4.2, respectively.

Following a measurement of their mass, a few specific spheres were mechanically extracted from within the optical trap, and transferred to an electron microscope, which is a common characterization tool used for microscopic objects. The electron microscope used allows for exceptionally precise imaging, much like a camera, which we apply to determine the size of the spheres that were collected. Typical spheres measure about 5 microns in diameter, about size of a human red blood cell. Importantly, with both mass and radius measurements of individual spheres, we can make robust calculations particle density, which have implications for the internal structure and physical properties of the spheres. Various images captured with the electron microscope are shown in Figure 4.4.

- Spinning Spheres

In addition to their overall side-to-side and up-and-down motion, trapped spheres can also rotate in any direction within the optical trap. We have developed a method to control

this rotation by using electrical torques. Essentially, a strong electric field is applied to a trapped sphere, where the electric field points in a specific direction, and by virtue of the physical properties of the spheres used, they tend to orient themselves with the direction of the electric field, much like a loaded die always has the same number facing up when thrown. If the direction of the field is then changed, the orientation of the sphere follows. In this way, the sphere can be made to rotate by applying a rotating electric field, similar to how a basketball can be spun on the end of a fingertip. We have demonstrated electrically driven rotation up to speeds of approximately 100000 rotations/second. Minuscule wobbling of the sphere as it rotates is actually related to the physical properties of the sphere, and is used to extract information about how electric charge is distributed within the sphere.

The rotational motion of a trapped sphere has an immediate technical application as a vacuum gauge in order to measure gas pressures. If the sphere is spun-up to a fast rotation velocity and then released from the electric field, gas in the immediate vicinity of the sphere induces a drag force much like air resistance. This drag force causes the sphere's rotation to decelerate, where the rate of deceleration can be related directly the gas pressure around the sphere. An example of this deceleration is shown in Figure 5.4. For force-sensing applications, air molecules colliding with a trapped sphere cause a lot of random motion of the sphere, which contaminates any potential measurement. To get around this limitation, we remove nearly all of the gas from a chamber that houses our optical tweezer, although there is a small amount of residual gas leftover which sets a limit on how sensitive the apparatus can be. With a spinning sphere serving as a vacuum gauge, we find that, once evacuated, our chamber reaches a base pressure of approximately 1 nanobar, about one billionth of the atmospheric pressure at sea level.

The Results

So far, there has been significant attention to the technical capabilities of the apparatus described by this thesis work. In addition to developing these methods, we have applied them to investigations of fundamental physics. In particular, we have performed short-distance force sensing for gravitational forces, in order to further our knowledge of gravity and the underlying physics when two objects are very close together. As was hopefully motivated before, it is precisely this regime where we might find evidence of new physics or refinements of our current physical laws. Some of these theories are discussed in Chapter 6.

At their core, these investigations involve a very simple setup. First, a sphere is trapped and stabilized within the optical tweezer, serving as the “test mass”, i.e. the portion of the apparatus that responds to applied forces. Once the sphere is in place, a mechanical device is carefully introduced very near to the sphere. This device is a relatively fancy contraption, but can be conceptualized quite simply as a stick. The stick has mass and thus interacts gravitationally with a trapped sphere, serving as the “source mass” and exerting a force that may be either attractive (like conventional gravity) or even possibly repulsive, depending on any new physics that might manifest at these short distances. The experiments themselves then amount to a careful positioning of the source mass relative to the test mass, while simultaneously measuring the motion of the test mass to look for any response that might be indicative of new forces.

Two different searches have been performed thus far, with differing emphasis. The first was as simple as described: a sphere is trapped and a source mass is introduced in proximity. The source mass is made to oscillate toward and away from the trapped sphere, and the resulting response of the trapped sphere is measured. In a second search, the overall framework and setup is similar, with the addition of a stationary shielding device between the test and source masses in order to protect the test mass from possible background forces that can both limit sensitivity, as well as mimic possible signals. A schematic depiction of this latter measurement setup is shown in Figure 8.1. After careful analysis of data collected during both searches, we found no compelling evidence for new physics, although it’s distinctly possible the effect was smaller than our measurement capability, a constant concern for physics experiments.

Importantly, and despite the null result, **these measurements mark the first published instance of tests of gravity with optically levitated test masses.** This statement should not be taken lightly, as nearly all of the other tests of gravity in the short-distance regime make use of one particular type of apparatus (a torsion pendulum), and thus they might all be subject to similar uncertainties and errors. Complementary and competitive techniques are absolutely essential to making robust discoveries in physics, and this thesis presents just such a complementary method.

Appendix A

The T -Matrix Method

Following the construction of Helmholtz equation (Equation (1.2)), the spectral theorem is applied again in order to write both incident and scattered optical fields as superpositions of basis functions that are themselves solutions to the Helmholtz equation [13],

$$E_{\text{inc}} = \sum_n^{\infty} a_n \psi_n^{(\text{inc})} \quad \text{and} \quad E_{\text{scat}} = \sum_m^{\infty} p_m \psi_m^{(\text{scat})}, \quad (\text{A.1})$$

where a_n and p_m are coefficients of the basis functions. The summations are not usually taken to infinity, rather, they are truncated at some value set by the convergence behavior of the coefficients for a particular scenario. Assuming the dielectric microsphere has a linear response to the incident optical fields, the scattered field is related to the incident field by a simple matrix equation,

$$p_m = \sum_n^{\infty} T_{mn} a_n. \quad (\text{A.2})$$

The T -matrix from Equation (A.2) is geometric in nature, depending only on the size and shape of the dielectric particle (here, assumed to be a microsphere), the index of refraction, and the wavelength of light. It is usually calculated with the coordinate origin at the center of the microsphere, so that any translations or rotations of the particle relative to the beam can be considered by transforming the basis functions themselves, as opposed to the T -matrix, and thus the expansions of the incident and scattered electric fields. Details of this calculation will be discussed after a choice of basis functions.

An excellent set of basis functions known to solve the Helmholtz equation are the vector spherical harmonics [11, 13]. Following the treatment in Reference [13], choose a coordinate system centered on the dielectric microsphere, and consider two expansions of the incoming and outgoing electric fields

$$\begin{aligned}\vec{E}_{\text{in}} &= \sum_{n=1}^{\infty} \sum_{m=-n}^{m=n} a_{nm} \vec{M}_{nm}^{(2)}(k\vec{r}) + b_{nm} \vec{N}_{nm}^{(2)}(k\vec{r}) \\ \vec{E}_{\text{out}} &= \sum_{n=1}^{\infty} \sum_{m=-n}^{m=n} p_{nm} \vec{M}_{nm}^{(1)}(k\vec{r}) + q_{nm} \vec{N}_{nm}^{(1)}(k\vec{r}),\end{aligned}\tag{A.3}$$

where the functions \vec{M} and \vec{N} are given by,

$$\begin{aligned}\vec{M}_{nm}^{(1,2)}(k\vec{r}) &= N_n h_n^{(1,2)}(kr) \vec{C}_{nm}(\theta, \phi) \\ \vec{N}_{nm}^{(1,2)}(k\vec{r}) &= \frac{h_n^{(1,2)}(kr)}{kr N_n} \vec{P}_{nm}(\theta, \phi) + N_n \left(h_{n-1}^{(1,2)}(kr) - \frac{nh_n^{(1,2)}}{kr} \right) \vec{B}_{nm}(\theta, \phi),\end{aligned}\tag{A.4}$$

where $h_n^{(1,2)}(kr)$ are spherical Hankel functions of the first and second kinds, $N_n = [n(n+1)]^{-1/2}$ are normalization constants, and the vector spherical harmonics are given by,

$$\begin{aligned}\vec{C}_{nm}(\theta, \phi) &= \nabla \times (\vec{r} Y_n^m(\theta, \phi)) \\ \vec{P}_{nm}(\theta, \phi) &= \hat{r} Y_n^m(\theta, \phi) \\ \vec{B}_{nm}(\theta, \phi) &= \vec{r} \nabla Y_n^m(\theta, \phi),\end{aligned}\tag{A.5}$$

where $Y_n^m(\theta, \phi)$ are the usual scalar spherical harmonics with complex exponentials and associated Legendre polynomials [4], θ is the azimuthal angle measured relative to the $+z$ axis, and ϕ is the polar angle measured in the x - y plane relative to the $+x$ axis. There is a minor convention subtlety in the construction of Equation (A.3): one can choose to expand in terms of purely incoming and outgoing fields (as written); or one can specify an ‘‘incident’’ field with both incoming and outgoing terms, together with a scattered field that is purely outgoing. The latter approach requires using regularized basis functions, since those in Equation (A.4) have singularities at the origin, but there are well-defined and easily calculable transformations between the two solution conventions [14, 15].

To actually calculate the T -matrix itself, a third expansion, \vec{E}_{int} , is required to describe the field internal to the dielectric microsphere [14], an intermediate step of sorts, where

it is assumed that the microsphere is homogeneous and isotropic given use of these basis functions. This expansion has the form of Equation (A.3), but uses the regularized basis functions that are singularity free, with their own set of coefficients,

$$\begin{aligned}\overline{RgM}_{nm}^{(1,2)}(k\vec{r}) &= \frac{1}{2} \left[\overline{M}_{nm}^{(1)}(k\vec{r}) + \overline{M}_{nm}^{(2)}(k\vec{r}) \right] \\ \overline{RgN}_{nm}^{(1,2)}(k\vec{r}) &= \frac{1}{2} \left[\overline{N}_{nm}^{(1)}(k\vec{r}) + \overline{N}_{nm}^{(2)}(k\vec{r}) \right].\end{aligned}\tag{A.6}$$

With these three multipole expansions, a calculation proceeds as follows: first, the matrix relating the incoming field, \vec{E}_{in} , to the internal field, \vec{E}_{int} , is determined; then the matrix relating the internal field to the outgoing field, \vec{E}_{out} , is determined; finally, the product of the two matrices yields the desired T -matrix. Both of the aforementioned transformation matrices are calculated by enforcing the following boundary conditions at the surface of the dielectric microsphere, when $r = d/2$,

$$\begin{aligned}\hat{n} \times \left(\vec{E}_{\text{inc}}(r) + \vec{E}_{\text{scat}}(r) \right) &= \hat{n} \times \vec{E}_{\text{int}} \\ \hat{n} \times \left(\vec{H}_{\text{inc}}(r) + \vec{H}_{\text{scat}}(r) \right) &= \hat{n} \times \vec{H}_{\text{int}},\end{aligned}\tag{A.7}$$

where the magnetizing fields $\vec{H} = \vec{B}/\mu$, are given by similar expansions: $\vec{H}_{\text{inc}} = (1/k_{\text{medium}})\vec{E}_{\text{inc}}$, $\vec{H}_{\text{out}} = (1/k_{\text{medium}})\vec{E}_{\text{out}}$, and $\vec{H}_{\text{int}} = (1/k_{\text{dielectric}})\vec{E}_{\text{int}}$, with $k = 2\pi n/\lambda$ as before.

Now we can return to the original task of understanding the scattering of a focused laser beam off of a dielectric microsphere in the vicinity of the beam's focus. Given the microsphere's dimensions and material properties (i.e. index), and an input beam wavelength, the T -matrix is first calculated to some desired precision of multipole expansion,¹ as this is one of the more computationally expensive operations. With the T -matrix in hand, a numerical solution for the scattered electric field is constructed by first matching \vec{E}_{in} from Equation (A.3) to the known electric field of an input beam, which yields the input expansion coefficients a_{nm} and b_{nm} . This can be done in the far-field away from the focus, or at the focal plane itself, where both methods have certain advantages depending on one's application [15].

¹Consider a numerical solution that has been truncated at $n = N_{\text{max}}$. Given the form of Equation (A.3), there are a total of $2N_{\text{max}}(N_{\text{max}} + 2)$ coefficients, which determines the size of the T -matrix that must be calculated.

Then, the scattered field is given directly by a simple matrix operation,

$$p_i = T_{ij}a_j, \tag{A.8}$$

where i and j run from 1 to $2N_{\max}(N_{\max} + 2)$, and $\{p_i\}$ ($\{a_j\}$) here represents the collection of both p_{nm} and q_{nm} (a_{nm} and b_{nm}) with the solution truncated at $n = N_{\max}$.

Appendix B

Harmonic Oscillators

For small oscillations about the equilibrium position, the center-of-mass motion of an optically trapped microsphere can be reasonably approximated by that of a harmonic oscillator, as was shown in Chapter 1. This is also true of the librational degree of freedom of a rotating microsphere. For the center-of-mass motion, the harmonic approximation forms the basis of our physical calibration from measured microsphere response to actual applied forces

It goes without saying that physicists at nearly all levels of training are familiar with the harmonic oscillator, given its broad applicability across essentially every subfield. As such, there are many phenomenal resources for the interested reader, especially introductory textbooks, and reproducing some of the associated derivations is clearly outside of the scope of this work. However, presentations vary in their choice of convention, so we include a sufficient set of equations such that the convention used throughout this work is known.

B.1 The Equation of Motion

The governing equations for harmonic oscillators are derived from Hooke's law [222], which relates the displacement/deformation of an elastic body to the force required to generate that deformation. Consider the classic model of a point mass, m , attached to one end of a massless spring with spring constant k , and examine a displacement, x , of the mass along a single axis. Then the displaced mass is subject to a spring force given by,

$$F_{\text{spring},x} = -k(x - x_0) \hat{x}, \tag{B.1}$$

where x_0 is the equilibrium position of the mass with the spring unstretched and the minus sign indicates a restoring force. Given that $F = m a = m \ddot{x}$, Hooke's law is a second-order linear differential equation.

A general equation of motion can then be constructed by including external driving forces, as well as damping forces like drag. Let $x_0 = 0$, and consider $\Sigma F = F_{\text{spring}} + F_{\text{drive}} + F_{\text{drag}} = m a$. Skipping a few steps and collecting the final result,

$$\ddot{x} + \gamma \dot{x} + \omega_0^2 x = \text{Re}\left[\frac{F_0}{m} e^{i\omega t}\right], \quad (\text{B.2})$$

where the dot indicates a time derivative, γ is a damping factor with units of angular frequency, $\omega_0 = \sqrt{k/m}$ is the natural resonant frequency of the harmonic oscillator, and F_0 and ω are the amplitude and frequency of the external driving force.

The steady-state solution to the above EOM can be calculated using the ansatz method, assuming a solution oscillating at the frequency of the driving force and given by,

$$x(t) = \text{Re}[A e^{i(\omega t + \phi)}]. \quad (\text{B.3})$$

The amplitude A and phase ϕ (relative to a drive phase of 0) can then be determined by substituting this ansatz into the differential equation. This yields,

$$A = \frac{F_0}{m [(\omega_0^2 - \omega^2)^2 + \omega^2 \gamma^2]^{1/2}} \quad \text{and} \quad \phi = \text{atan} \left(\frac{\omega \gamma}{\omega^2 - \omega_0^2} \right), \quad (\text{B.4})$$

which are the amplitude and phase of the harmonic oscillator response to a single-frequency driving force of the form $\text{Re}[(F_0/m)e^{i\omega t}]$. In the case of multiple driving forces, the total solution is a superposition of these single-frequency solutions.

Appendix C

Phasors and Heterodyne Interference

The following calculations have been included in this appendix, as most introductory literature (and even some peer-reviewed manuscripts) seems to avoid discussing potential pitfalls when using phasor notation for interference problems. There are even select instances where the calculations appear to be incorrect, although usually only by a factor of 2.

C.1 A Gaussian Beam

The following analysis will be performed in the paraxial limit, with perfectly aligned TEM₀₀ modes. The electric field of a Gaussian beam at its focal plane is given by,

$$\begin{aligned} \vec{E}(\vec{x}, t) &= E(\vec{x}) \cos[\Phi(t)] = \text{Re} \left(E(\vec{x}) e^{i\Phi(t)} \hat{p} \right) \\ \text{with } E(\vec{x}) &= e^{-\frac{|\vec{x}|^2}{w^2}} \quad \text{and} \quad \Phi(t) = \omega t + \phi, \end{aligned} \tag{C.1}$$

where $E(\vec{x})$ is the spatial mode of electric field with peak value E and beam waist w (the radius at which the intensity has dropped by $1/e^2$); $\Phi(t)$ is the temporal phase, with $\omega = 2\pi c/n\lambda$ the optical frequency, c the speed of light, n the index of refraction, λ the vacuum wavelength, and ϕ an arbitrary constant phase (included for generality); and \hat{p} is the polarization of the beam, assumed orthogonal to the direction of propagation.

Given that traditional optical detectors are sensitive to intensity and not the electric field directly, the square of the electric field $|\vec{E}(\vec{x}, t)|^2$ is generally considered. Furthermore,

the intensity will be integrated over an infinite (or half-infinite) plane to compute the optical power, and subjected to a time average, given that semiconductor-based detectors are not capable of responding at optical frequencies of \sim PHz,

$$\langle \mathcal{P} \rangle = \left\langle \int dA I(\vec{x}) \right\rangle = \int dA \left\langle \frac{1}{\eta} |\vec{E}(\vec{x}, t)|^2 \right\rangle, \quad (\text{C.2})$$

where $\eta = 1/c\epsilon$ is the wave impedance and ϵ is the dielectric constant. Since the integral is over spatial coordinates, it's assumed to commute with the time-averaging operator. Consider the total power in a single Gaussian beam. This calculation will be performed both with the real-valued field *and* in phasor notation, making use of the expression $|A|^2 = A^* A$ with (*) the complex conjugate,

$$\begin{aligned} \langle \mathcal{P} \rangle &= \int dA \left\langle \frac{1}{\eta} |E(\vec{x}) \cos[\Phi(t)]|^2 \right\rangle & \langle \mathcal{P} \rangle &= \int dA \left\langle \frac{1}{\eta} |E(\vec{x}) e^{i\Phi(t)}|^2 \right\rangle \\ &= \frac{\langle \cos^2[\Phi(t)] \rangle}{\eta} \int dA \left(E^2 e^{-2\frac{|\vec{x}|^2}{w^2}} \right) & &= \frac{\langle |e^{i\Phi(t)}|^2 \rangle}{\eta} \int dA \left(E^2 e^{-2\frac{|\vec{x}|^2}{w^2}} \right) \\ &= \frac{\langle \frac{1}{2} \{1 + \cos[2\Phi(t)]\} \rangle}{\eta} \left(\frac{\pi w^2 E^2}{2} \right) & &= \frac{\langle e^{-i\Phi(t)} \cdot e^{i\Phi(t)} \rangle}{\eta} \left(\frac{\pi w^2 E^2}{2} \right) \\ &= \frac{\pi w^2 E^2}{4\eta}, & &= \frac{\pi w^2 E^2}{2\eta}, \end{aligned} \quad (\text{C.3})$$

where $\langle \cos[2\Phi(t)] \rangle = 0$ trivially since $\Phi(t) = \omega t + \phi$ is linear in t and $\omega \sim$ PHz is well above the response time of any practical detector; and $e^{-i\Phi(t)} \cdot e^{i\Phi(t)} = 1$ for all t . This result clearly presents a problem, given that the answers are not the same. The issue arises from assuming that the real-value operator, $\text{Re}(\dots)$ commutes with the squared norm, $|\dots|^2$, which was done implicitly in first line of the right column. This is often the case for interference problems: people make use of the phasor notation, since complex exponentials are powerful and make some mathematical operations trivial, and arrive at quantitatively inaccurate solutions.

For single beams, the phasor notation can be used naïvely to compute the power, simply including an “extra” factor of $(1/2)$ from the time-average of the oscillating terms.¹ Then, we arrive at the correct expression relating the optical power, a directly measurable quantity,

¹Although this is by no means rigorous, one can recall Euler's formula ($e^{i\phi} = \cos \phi + i \sin \phi$) and say that computing the complex-square of an electric field in phasor notation “double counts” the power since there are two sinusoids (cos and sin) in a single phasor ($e^{i\phi}$). A simple factor of $(1/2)$ adjusts for the “double-counting”.

to the beam waist and electric field of a TEM₀₀ mode beam,

$$\langle \mathcal{P} \rangle = \frac{\pi w^2 E^2}{4\eta}, \quad (\text{C.4})$$

which is consistent with well-established formalism [5, 20].

C.2 Two Gaussian Beams

Now let's examine the interference of two ideal TEM₀₀ beams with distinct frequencies, and consider their interference at a mutual focal plane (the surface of some imagined detector) where both wavefronts are flat, but may have differing temporal phase. By the principle of superposition, the total electric field is simply the sum,

$$\begin{aligned} \vec{E}_{\text{tot}}(\vec{x}, t) &= \vec{E}_1(\vec{x}, t) + \vec{E}_2(\vec{x}, t) \\ &= E_1(\vec{x}) \cos[\Phi_1(t)] \hat{p}_1 + E_2(\vec{x}) \cos[\Phi_1(t)] \hat{p}_2 \\ &= \text{Re} \left\{ E_1(\vec{x}) e^{i\Phi_1(t)} \hat{p}_1 + E_2(\vec{x}) e^{i\Phi_2(t)} \hat{p}_2 \right\}. \end{aligned} \quad (\text{C.5})$$

Following the same general procedure as before, let's consider the optical power of this superposition, first using the sinusoidal form of the expression for the total field. The polarizations are assumed to be aligned so that $\hat{p}_1 = \hat{p}_2$, noting that the interfering portion of the final power would be suppressed by $\hat{p}_1 \cdot \hat{p}_2 = \cos\theta$ for some misalignment θ between the polarizations. This yields,

$$\begin{aligned} \langle \mathcal{P} \rangle &= \int dA \left\langle \frac{1}{\eta} \left| \vec{E}_{\text{tot}}(\vec{x}, t) \right|^2 \right\rangle \\ &= \frac{1}{\eta} \int dA \left\langle \left| E_1(\vec{x}) \cos[\Phi_1(t)] \hat{p}_1 + E_2(\vec{x}) \cos[\Phi_1(t)] \hat{p}_2 \right|^2 \right\rangle \\ &= \frac{1}{\eta} \int dA \left\langle E_1(\vec{x})^2 \cos^2[\Phi_1(t)] + E_2(\vec{x})^2 \cos^2[\Phi_2(t)] \right. \\ &\quad \left. + 2E_1(\vec{x})E_2(\vec{x}) \cos[\Phi_1(t)] \cos[\Phi_2(t)] \right\rangle \\ &= \frac{\pi w_1^2 E_1^2}{4\eta} + \frac{\pi w_2^2 E_2^2}{4\eta} + \frac{E_1 E_2}{\eta} \left\langle \cos[\Phi_1(t)] \cos[\Phi_2(t)] \right\rangle \int dA e^{-\frac{|\vec{x}|^2}{w_1^2}} e^{-\frac{|\vec{x}|^2}{w_2^2}} \\ &= \langle \mathcal{P}_1 \rangle + \langle \mathcal{P}_2 \rangle + \frac{2E_1 E_2}{\eta} \left\langle \frac{1}{2} \{ \cos[\Sigma\Phi(t)] + \cos[\Delta\Phi(t)] \} \right\rangle \left(\frac{\pi w_1^2 w_2^2}{w_1^2 + w_2^2} \right), \end{aligned} \quad (\text{C.6})$$

where $\Sigma\Phi(t) \equiv \Phi_1(t) + \Phi_2(t)$ and $\Delta\Phi(t) \equiv \Phi_1(t) - \Phi_2(t)$ are defined for brevity, and the

result from Equation (C.2) is used. The term oscillating at $\Sigma\Phi(t)$ is clearly well above detector bandwidth and thus has a trivial time-average, but the $\Delta\Phi(t)$ will stay around as it is assumed that $(\omega_1 - \omega_2)$ is within the detector bandwidth (by design!). Thus,

$$\begin{aligned}\langle \mathcal{P} \rangle &= \langle \mathcal{P}_1 \rangle + \langle \mathcal{P}_2 \rangle + \frac{\pi w_1^2 w_2^2 E_1 E_2}{\eta(w_1^2 + w_2^2)} \cos [\Delta\Phi(t)] \\ &= \langle \mathcal{P}_1 \rangle + \langle \mathcal{P}_2 \rangle + \frac{4w_1 w_2}{w_1^2 + w_2^2} \sqrt{\langle \mathcal{P}_1 \rangle \langle \mathcal{P}_2 \rangle} \cos [\Delta\omega t + \Delta\phi],\end{aligned}\tag{C.7}$$

where the relation from Equation (C.2) has been used to replace E_i with $\langle \mathcal{P}_i \rangle$, and defined the difference/interference frequency $\Delta\omega \equiv \omega_1 - \omega_2$ and the phase offset $\Delta\phi \equiv \phi_1 - \phi_2$. If $w_1 = w_2$ and $\langle \mathcal{P}_1 \rangle = \langle \mathcal{P}_2 \rangle$, then a familiar result is recovered: $\langle \mathcal{P} \rangle = 2 \langle \mathcal{P}_{1,2} \rangle [1 + \cos (\Delta\omega t + \Delta\phi)]$, i.e. that the total incident power of the interference of two co-aligned Gaussian beams with the same spatial mode, but different frequencies, oscillates between 0 and $4\times$ the individual beam power [5, 20].

Now let's do the same thing with phasors, using the usual formula for the squared-norm of a complex quantity,

$$\begin{aligned}\langle \mathcal{P} \rangle &= \int dA \left\langle \frac{1}{\eta} \left| \vec{E}_{\text{tot}}(\vec{x}, t) \right|^2 \right\rangle \\ &= \frac{1}{\eta} \int dA \left\langle \left| E_1(\vec{x}) e^{i\Phi_1(t)} \hat{p}_1 + E_2(\vec{x}) e^{i\Phi_2(t)} \hat{p}_2 \right|^2 \right\rangle \\ &= \frac{1}{\eta} \int dA \left\langle E_1(\vec{x})^2 + E_2(\vec{x})^2 + E_1(\vec{x}) E_2(\vec{x}) \left(e^{i[\Phi_1(t) - \Phi_2(t)]} + e^{i[\Phi_2(t) - \Phi_1(t)]} \right) \right\rangle \\ &= \frac{\pi w_1^2 E_1^2}{2\eta} + \frac{\pi w_2^2 E_2^2}{2\eta} + \frac{E_1 E_2}{\eta} \langle 2 \cos [\Delta\Phi(t)] \rangle \int dA e^{-\frac{|\vec{x}|^2}{w_1^2}} e^{-\frac{|\vec{x}|^2}{w_2^2}} \\ &= 2 \langle \mathcal{P}_1 \rangle + 2 \langle \mathcal{P}_2 \rangle + \frac{2E_1 E_2}{\eta} \cos [\Delta\Phi(t)] \left(\frac{\pi w_1^2 w_2^2}{w_1^2 + w_2^2} \right) \\ &= 2 \langle \mathcal{P}_1 \rangle + 2 \langle \mathcal{P}_2 \rangle + \frac{8w_1 w_2}{w_1^2 + w_2^2} \sqrt{\langle \mathcal{P}_1 \rangle \langle \mathcal{P}_2 \rangle} \cos [\Delta\omega t + \Delta\phi],\end{aligned}\tag{C.8}$$

where Euler's formula, $2 \cos \phi = e^{i\phi} + e^{-i\phi}$, has been employed. Thus we find the same result: if one uses phasors and proceeds naïvely, ignorant of the underlying complexities, then the total, time-averaged power computed needs an extra factor of $(1/2)$ in order to match the correct result,

$$\langle \mathcal{P} \rangle_{\text{correct}} = \frac{1}{2} \langle \mathcal{P} \rangle_{\text{phasors}}.\tag{C.9}$$

Appendix D

System Control and Data Acquisition

As the text may allude, there is a vast array of hardware and instrumentation for the table-top experiment that this thesis describes. The control of and subsequent data acquisition from this apparatus, the focus of this Appendix, is centralized to a single computer, with a number of peripherals. There are a few “aged” instruments that are immune to remote control (either digital or analog), but for the most part, the single central computer provides the user interface with the apparatus and manages the analog-to-digital conversions, as well as the digital-to-analog conversions.

Generally speaking, the controls and data acquisition are divided into two groups: fast and slow. The fast control and acquisition is accomplished with National Instruments hardware housed within the host PC or an external PXI chassis, making use of LabVIEW as an interface. Various slow controls are achieved with USB, RS-232, and GPIB protocols, essentially using whatever input the instrument in question has available. A small subset of slow controls are implemented with a National Instruments Compact FieldPoint system, although this hardware is outdated and can’t interface with newer versions of LabVIEW.

D.1 Controllers and Their Assignments

The various control systems are tabulated below, with a short verbal description of the device function. Further details will be provided later in relation to the specific tasks performed, such as digitization.

Table D.1: Components housed within the central **PXI-1042** chassis and their functions.

Model	Function	Use
PXI-8331	Controller	Communication between host PC and PXI chassis, data transfer
PXI-7854R	FPGA with integrated ADCs and DACs (both 16-bit)	Primary signal digitization, digital demodulation, active feedback, synchronization
PXI-6259	Reconfigurable analog/digital inputs and outputs (both 16-bit)	Electrode and stage monitors, high-frequency digitization for rotational signal
PXI-6723	Multi-channel analog output (13-bit)	Electrode driver, attractor stage driver
PXI-6229	Reconfigurable analog/digital inputs and outputs (both 16-bit)	Valve control for main vacuum system, pressure and temperature monitors

D.1.1 PXI Hardware

Nearly all of the fast digitization and signal generation is performed with hardware conforming to the PCI eXtensions for Instrumentation (PXI) platform, specifically from National Instruments. A chassis houses the various hardware units, listed in Table D.1, providing power and interconnectivity. A controller module within the chassis communicates with a companion module within the host PC and allows commands to be sent to individual hardware units, as well as facilitates data transfer for offline analysis. Virtual instruments that interface with and make use of these hardware units are created with LabVIEW and executed on the host PC.

The chassis itself provides a hardware synchronization interface, with a 10 MHz analog output that can be used to phase-lock the digital timebase of other instruments, a critical feature for phase-locked digital demodulation. The interface also has an analog input that can be used to synchronize the master timebase within the PXI chassis to any external 10 MHz source. A number of different synchronization configurations have been explored and tested, the most robust being the use of a stable external signal to which the chassis locks, together with other instruments.

Table D.2: PCIe components housed within the host PC and their functions.

Model	Function	Use
PCIe-7841	FPGA with integrated ADCs and DACs (both 16-bit)	Multi-channel DDS for generating phase-continuous rotational drive signals, and feedback for a single rotational degree of freedom
StarTech PEX2PCIE4L	PCIe to PCI/PCIe Expansion Chassis	Provides PCI slot (not in computer) for the controller module for the main PXI chassis

D.1.2 PCIe Hardware

The FPGA module used for the primary signal digitization (coordinates x , y , z , power \mathcal{P}) is at nearly full utilization¹ with the various timekeeping/synchronization tasks, demodulation, downsampling, and finally feedback signal generation. As a result, the whole apparatus requires extra hardware in order to perform similar functions, but for different input signals and/or possible actuators. The most obvious solution would be a secondary FPGA module within the PXI chassis, but for economical reasons, the old FPGA module from the first iteration of the apparatus was repurposed. This old module has a PCIe form factor, and thus can sit directly inside the host PC without an external chassis or a separate controller. Its inputs and outputs are dedicated to signals and controls for the rotational degrees of freedom of a trapped microsphere.

As mentioned briefly, the PXI chassis connects to the host PC via a pair of modules. The module on the PC side of things (as opposed to the one within the chassis itself) is actually a PCI (not PCIe) form factor, and the motherboard within the control computer does not have any PCI slots, as most newer systems have transitioned entirely to PCIe. Thus, in order to connect to and control the PXI hardware, there is a PCIe-to-PCI/PCIe expansion chassis. Despite the existence of this nested hardware, the host PC just “sees” the expansion chassis as normal components, as if they were plugged directly into the motherboard. This subtlety will be ignored beyond the acknowledgment within this section.

¹This is likely a result of poor optimization on the part of the user. To ease the digital signal processing and remove concerns about data loss from various operations that might result in overflow or changes in numerical precision, the digitized 16-bit signals end up being stored in 32-bit, and sometimes 64-bit, datatypes. Thus, any digital buses, shift registers, buffers, etc. need the same bit-width, which can significantly increase the required elements within the FPGA itself.

Table D.3: Components housed within the auxiliary **cFP-1808** chassis and their functions.

Model	Function	Use
cFP-AI-112	ADC module (16-bit)	Dedicated to various power and fault monitors for the fiber-optic laser amplifier
cFP-AO-210	DAC module (12-bit)	Enable/disable and level control for laser amplifier
cFP-DO-400	Sourcing digital output module	Valve control for custom gas manifold

D.1.3 Compact FieldPoint

The Nufern fiber-optic laser amplifier in use on this setup intentionally does not have a standalone controller and is instead controlled and monitored directly by Compact FieldPoint (cFP) hardware and LabVIEW VIs. The amplifier has a 25-pin DSUB connector, with TTL level enable pins and fault monitors (i.e. power fault, high back-reflection fault), as well as an analog control to set the output level and analog power monitors from photodiodes within the amplifier itself. All of the inputs and outputs for the laser amplifier control can be effectively handled with two dedicated ADC and DAC modules, respectively. The cFP hardware is advantageous for this application, as the outputs are meant to run continuously (the laser amplifier has a long stabilization time, so enabling and/or changing output level incurs a time cost), unlike much of the PXI hardware which executes output tasks and digitizations in triggered bursts.

The cFP hardware also serves a secondary purpose: controlling the valves within the custom gas manifold discussed in Section 2.1.1. A digital output cFP module that can source significant current to multiple channels is used to drive electronic relays, which in turn actuate the solenoid valves controlling the compressed air, which then actuate the pneumatically-driven valves controlling gas flow within the manifold and thus to the chamber. There is a 1-1-1 correspondence, so each digital output drives a unique relay, solenoid valve, and finally pneumatic valve, in order to fill the manifold with a single gas species, or open the manifold to its dedicated scroll pump.

Given that these two functions are relatively straightforward and explained in full above, this control hardware and its application will not be discussed further in this chapter.

Table D.4: Instruments controlled with a single GPIB Hub, the NI-GPIB-HS, connected to the host PC via USB.

Model	Function	Use
SG382	RF Signal Generator	Generates 150 MHz drive for trapping beam AOM, with amplitude modulation for axial feedback
HP8643A	RF Signal Generator	Generates 143.375 MHz drive for reference beam AOM
HP3325B	Function Generator	General purpose, multiple applications including chirp drive signal for piezo-launcher based microsphere trapping
WW5064	Arbitrary Waveform Generator	General purpose, multiple applications including electrode driver

D.1.4 GPIB

The IEEE-488/General Purpose Interface Bus (GPIB) is a digital bus and communications protocol that was developed in the early 1960's by Hewlett-Packard (called HP-IB at the time). Despite its age, GPIB is still used quite frequently for test and measurement equipment due to the proliferation of the standard, as well as the unique feature that the connectors are stackable, meaning instrument connections can be daisy-chained, up to a total of 15 instruments given the 4-bit addressing used. Each device connected to the bus can “take control”, sending commands to other connected devices and reading any subsequent responses, whilst locking out commands from other devices.

In this apparatus, the host PC is connected via a single USB to a National Instruments GPIB Hub, which in turn is daisy-chained to a number of signal-generators. Instrument setup commands (waveform type, amplitude, frequency, offsets) are sent to the GPIB hub as ASCII-encoded strings with specific destination addresses, which then get forwarded along to the individual instruments specified by the addresses. It should be noted that most instruments have unique and very specific formats for these command strings. At the time of writing, each device listed in Table D.4 has a publicly available operations manual with details concerning the format and syntax of the accepted GPIB commands.

Table D.5: Instruments connected to the host PC via USB.

Model	Function	Use
PS-5 Eternal Laser Controller	Controller for the primary Orbits Lightwave Yb-fiber seed laser	Sets optical pump power, monitors temperature throughout the laser module itself, provides both hardware (turn-key) and software enable
CLD1015	Laser diode and TEC controller for butterfly packages	Controls the laser illumination for the top-view microscope
EO-0513	Color CMOS camera	Side-view microscope to verify if a microsphere is trapped, and roughly determine height within the trap
Mako U-029B	Monochrome CMOS camera	Top-view microscope to assess and refine alignment of mechanical devices in proximity to a trapped microsphere
WW5064	Arbitrary Waveform Generator	General purpose, multiple applications including electrode driver

D.1.5 USB, Serial, and Ethernet

Any device not included in the previous sections is likely connected to the host PC via USB, or else not amenable to remote control. However, many older instruments and peripheral devices use the serial communications protocol RS-232 and 9-pin DSUB connectors, while newer instruments use USB. Instruments within this apparatus requiring RS-232 are connected to the host PC via USB-to-RS-232 converters, and appear as serial “COM” ports, assuming the correct drivers have been installed. ASCII- or byte-string commands can then be sent via a serial communications protocol to be interpreted by the instruments, and queried values or instrument states can be returned to the host PC.

A single instrument, a residual gas analyzer (MKS, eVision+), makes use of ethernet to communicate with a host computer. Proprietary software running on the host PC can identify and communicate with any instruments connected to the local network.

D.2 LabVIEW Primer

In nearly all of the methods discussed below, the control and data acquisition hardware is interfaced and controlled with LabVIEW running on a single host PC. The specific implementation within LabVIEW is available to the interested reader via a GitHub repository [223], however it should be noted that this is only one particular implementation, and likely a sub-optimal one. Not only can the same tasks be accomplished with different frameworks entirely within LabVIEW, but low-level control could also be implemented by directly interfacing with the PXI chassis and the instruments within, given that PXI is an industry standard, although one invented and championed by National Instruments.

As such, we refrain from including specifics of the LabVIEW and instead focus on the hardware and its configuration for a desired control or data acquisition task at hand.

D.3 Analog Signal Production

Despite the apparent complexity of this apparatus, there are a limited number of analog signals that are applied at any given time. Once a trapped microsphere has been lowered to the focus, feedback has been engaged, and the experimental chamber has been brought to high vacuum, the only driven inputs to the system are the AOM RF drives, the feedback signals, the electrode voltages, and the attractor stage drive signals.

D.3.1 Synthesizers

Two different RF synthesizers are used to drive the fiber-coupled AOMs that control the total amount of light in the trapping and reference beams. Both are operated in a continuous fashion, with level control implemented via built-in amplitude modulation, whose source signals will be discussed in the next section.

The trapping beam synthesizer is a Stanford Research Systems SG382, capable of producing signals up to 2 GHz. The synthesizer output first passes through a 10 dB attenuator before being sent to a broadband amplifier with a gain of 30 dB (MiniCircuits, ZHL-03-5WF+). The amplified signal passes through another 10 dB attenuator before driving the AOM itself. The repeated attenuators have a dual purpose: first and foremost they provide a degree of protection for the RF electronics, from both over-driving any of the components and from any reflections due to mismatched impedances, as reflections must pass through

each attenuator twice, preventing damage to the active components; the attenuators can also be swapped out for different values to tune the dynamic range of the control over optical power, a critical feature if the apparatus is to be used with different sizes of microspheres, which require different optical power to levitated at the focus of the vertically-oriented, gravity-stabilized optical trap used here.

It serves to elaborate slightly on these. Firstly, the input of the RF amplifier can only accept up to +10 dBm signal without suffering damage, while the synthesizer full-scale is +13 dBm. Some amount of attenuation here prevents over-driving the amplifier. Similarly, the AOM can only accept up to +33 dBm of RF power without suffering damage. Thus, the second 10 dB attenuator prevents the amplifier from over-driving the AOM, regardless of input level, and is left permanently affixed. The first attenuator, as well as control of the synthesizer's full-scale, can be modified to tune the maximum optical power, which necessarily sets the dynamic range since the modulation port is driven with 16-bit DAC, yielding a resolution of roughly $1/(2^{16} - 1)$ of the full-scale.

The reference beam synthesizer is a Hewlett Packard HP8643A, capable of producing signals up to 1 GHz. As with the trapping beam synthesizer, the reference beam synthesizer output first passes through a 10 dB attenuator before being sent to a broadband amplifier with a gain of 37 dB (EIN, 400AP). The amplified signal passes through another 6 dB attenuator before driving the trapping beam AOM. This RF drive is usually left at a constant value, so that the reference beam is both constant and stationary. Future improvements of the apparatus may make use of active level control here to stabilize the interference measurements against amplitude fluctuations, since the oscillating interference signals are proportional to the square root of the beam power.

D.3.2 Feedback Signals

There are a total of three feedback signals and their associated actuators for active stabilization of a trapped microsphere's translational degrees of freedom. All three signals are first generated digitally within the primary FPGA (to be discussed below), and subsequently converted to analog voltages by ± 10 V, 16-bit DACs, operating at 5 kHz, well above the natural resonant frequencies of the microsphere center-of-mass motion. However, this voltage range is too large for all of the feedback actuators, thus requiring a degree of level-control.

The amplitude modulation port for the trapping beam synthesizer accepts a ± 1 V signal for RF amplitudes from 0 to full-scale. To adapt to this input, the z /axial feedback signal

from the primary FPGA drives a simple 10:1 voltage divider, passively converting the ± 10 V down to ± 1 V, whose output is connected directly to the synthesizer modulation input. The FPGA output impedance and modulation port input impedance are sufficiently low- and high-impedance that the voltage divider can be constructed with a variety of resistor values, where $10\text{ k}\Omega$ and $90\text{ k}\Omega$ have been chosen to accomplish the 10:1 division.

For the radial degrees of freedom, a piezoelectrically-driven deflector mirror serves as the feedback actuator in order to control the radial position of the trap (see Chapter 2). The deflector requires relatively high voltage, 0 V to 75 V ,² as well as an output with a large slew rate in order to drive the capacitive load presented by the piezoelectric chip. A dedicated piezoelectric amplifier (Thorlabs, MDT693A) is used to drive the two orthogonal deflections of the trapping beam, although this amplifier has a full-scale of 150 V set by an analog input of 0 V to 10 V . Given the positive-definite nature of the required drive voltage, an active circuit is employed to adjust the $\pm 10\text{ V}$ output from the FPGA. This level-control circuit, documented in Appendix E, divides the signal by a factor of four and adds a bias of 2.5 V . Thus the $\pm 10\text{ V}$ is converted to 0 V to 5 V , which in turn yields a deflector drive voltage of 0 V to 75 V following amplification.

D.3.3 Electrode and Stage Drives

The final collection of signals required for control of the apparatus are drive signals for the electrodes and the nano-positioning stage that controls the motion of the attractor device used in the fundamental science measurements. The electrodes have two different drive systems for distinct applications, one of which shares hardware with the stage driver and thus is discussed here. This first system makes use of a dedicated PXI module, the PXI-6723, which has a total of $32\ \pm 10\text{ V}$, 13-bit DACs, also operated at 5 kHz . By virtue of the PXI chassis synchronization, these outputs are triggered simultaneously with digitization tasks, and phase-locked to other signal sources.

Given that the hardware is controlled and operated by LabVIEW, there exists a well-defined formalism to programmatically define desired output signals with a fixed length (DAQmx tasks). These can be simple sinusoids, or arbitrarily defined waveforms with complex spectral content, the latter being specifically useful for testing the frequency-response of a trapped microsphere.

²Piezoelectric chips generally do not like negative voltages, thus passive voltage division is not a viable solution to use the full dynamic range of the FPGA.

A total of 11 of the 32 output channels are used: three for the x , y , and z degrees of freedom of the attractor's nano-positioning stage, and eight for the electrodes (there are 6 truncated, pyramidal electrodes surrounding the trap, together with the attractor and shield which can be independently biased). The electrode signals output from the PXI-6723 are buffered by a unity-gain amplifier circuit which also produces a separate 1:1 monitor signal to be digitized and stored for offline analysis. Thus the range of applicable voltages for the electrodes is simply ± 10 V. Further amplification can be used if different signal levels are required.

The attractor's nano-positioning stage accepts 0 V to 10 V signals, so an active level control circuit (see Appendix E) is used to divide the ± 10 V by a factor of two, and add a 5 V DC offset. These adjusted signals are then sent to the stage driver to control the attractor's position. The position of the attractor's stage is measured by internal strain gauges (for closed-loop control of the hysteric piezos), with output monitor ports for subsequent digitization and offline analysis.

D.3.4 DDS Electrode Drive

While the simultaneously triggered drive signals and digitization tasks are critical for a number of applications, most notably force-sensing, there are some instances where a drive signal must be continuous for long periods of time, usually much longer than an individual drive/digitization task (usually these are about 10s long). For example, if the rotational drive signals used in Chapter 5 were to turn off at regular intervals set by the finite-length drive/digitization tasks, the orientation of a trapped microsphere's dipole moment would easily become uncorrelated with the phase of the spinning field, and many of the measurements detailed in that chapter would be fundamentally impossible. At the same time, these measurements require well-defined phase relationships, $\pm\pi/2$ rad, between neighboring electrode drives in order to produce the necessary spinning fields.

In order to generate phase-continuous and indefinite drive signals for the electrodes, a direct digital synthesis (DDS) drive system is implemented within an FPGA. This makes use of a secondary FPGA module with a PCIe form factor that sits directly inside the host PC. The centerpiece of this system is a single DDS set to one specific frequency, which outputs both sine and cosine components ($\Delta\phi = \pm\pi/2$ rad). These outputs can then be duplicated and inverted in order to generate the four, properly phased, sinusoids necessary to drive a rotating electric field with four of the trapping electrodes. The signals are generated and

scaled digitally, then sent to dedicated, ± 10 V, 16-bit DACs operated at 500 kHz, where the much larger bandwidth is critical for high-frequency spinning applications.

This signal generation structure is also convenient for planned future developments of the apparatus, as a phase modulation of the centerpiece DDS results in a common-mode phase modulation of each of the electrodes, corresponding to a modulation of the rotating electric field's angle relative to unmodulated rotating reference frame.

D.4 FPGA-based Digital Demodulation and Feedback

The central hardware component for the system control and data acquisition is the PXI chassis, which houses a number of different modules with various tasks. All modules are from National Instruments and interface natively with LabVIEW for programmatic control. There are a huge number of subtleties and intricacies to proper LabVIEW programming, and it's quite likely that not everything was implemented in the most optimal manner. Additionally, successive versions of LabVIEW always seem to change some important aspect of the LabVIEW programming, such that new projects may not be able to implement the same code structure as older projects using different versions of LabVIEW. Thus, the following sections will focus on effective descriptions while, for the interested reader, full details can be obtained by perusing the GitHub repository for the control computer [223].

The most critical module within the PXI chassis is the FPGA module that manages the primary signal digitization and demodulation, as well as the calculation and generation of the feedback signals used to actuate various hardware components of the apparatus. Amplified interference signals are sent from the QPD and the retroreflected PD directly to 16-bit ADC inputs embedded within the FPGA module, where the signals are subsequently demodulated to produce digital estimates of the amplitude and phase of each of the digitized interference signals. The details of this demodulation are discussed in the next section, followed by a description of the implementation in LabVIEW. In addition to the interference signals, the FPGA module also digitizes the amplified output from the power-monitoring photodiode, although this doesn't not require the same demodulation as the primary interference signals that encode the motion of the trapped microsphere.

Once the interference signals have been demodulated, estimates of the microsphere displacement are calculated from the signal amplitudes and phases. These estimates are compared to desired setpoints (determined empirically and set externally), to generate an error

signal for feedback, in particular: proportional, integral, and derivative (PID) feedback, which was discussed in Section 2.7.

A critical feature of the implementation of the digital demodulation as described is the phase-locking of the signal generators and the primary ADCs. Both the SG382 and the HP8634A signal generators, as well as the PXI chassis itself, include connections for synchronizing with external timebases, based on a shared 10 MHz timing signal that can be sourced by any of the three instruments. The most reliable phase-locking has been found with the SG382 acting as the primary source for the timebase, which is then shared with the PXI chassis, and further redistributed to the HP8643A and any other instruments.³ Within the PXI chassis, backplane connections allow the distribution and subsequent use of this timing signal by each of the modules, usually as triggers for the ADCs and DACs.

The FPGA module itself is configured and controlled with LabVIEW. Specifically, a LabVIEW FPGA project involves two executables that are intimately connected, one that operates within the FPGA, and one that operates on the host PC. Both are configured to run independently, although the host PC can write digital values to shift registers within the FPGA in order to adjust settings such as feedback strength, setpoints, digital variable scaling, and sampling frequency. Importantly, the host PC and FPGA are setup with Direct Memory Access (DMA) First-In First-Out (FIFO) buffers, so that when triggered, the host PC can extract data stored within the FPGA, such as the estimates of the microsphere position and the demodulated amplitudes and phases.

D.4.1 Timekeeping

Although the FPGA is the source of the primary signals (x , y , z , power \mathcal{P}), other hardware elements within the PXI chassis and otherwise connected to the host PC digitize signals of importance, such as electrode voltage monitors. It is important that these digitized signals can be synchronized in time with one another, for such things as a proper measurement of a trapped microsphere's frequency response. A nearly universal and easy to manage representation of time is the Unix epoch, a measure of the number of seconds that have transpired since January 1, 1970. All modern computers keep track of this epoch, often synchronizing to GPS time to account for slight variations and imperfections in their own

³Although all of the instruments expect a timing signal with a frequency of 10 MHz, they have different input impedances and amplitude requirements/limitations. This particular timing configuration best matched the requirements of our specific signal generators, and appeared to generate the least phase noise, as observed with the digital demodulation.

internal clock. The FPGA, much like a computer, can keep track of this number with an internal counter. Often, this is a 64-bit counter, where the 64-bit integer represents the number of nanoseconds since 1/1/1970.

The first level of timekeeping/synchronization matches the 64-bit counter within the FPGA to the 64-bit counter within the host PC, which is assumed to be synchronized with GPS time by virtue of the operating system and a robust internet connection. Once every second, when not actively sampling and transferring data for offline storage, the FPGA queries the counter within the host PC and loads this value into its own counter from which it continues counting up. In this way, the FPGA and host PC maintain consistent timekeeping so that data associated with particular timestamps can be compared. This is sufficient to match slow environmental monitors to the data digitized by the FPGA module.

For faster, synchronous digitization between multiple modules with ADCs, a more creative solution was required. In an ideal world, backplane triggers from the PXI chassis could easily be distributed among different modules, but the different acquisition structure naively prevents this. Dedicated digitizer modules that have been properly configured to operate based on external triggers can immediately begin digitizing once the trigger has been received. The FPGA, however, is always running, generating values of x , y , and z continuously in order to maintain active feedback. LabVIEW's interaction with DMA FIFOs is structured such that unless those values are being read out constantly, the buffer fills up and ceases to record new data. Thus, instead of constantly writing data into a FIFO, reading it out on the PC, and then discarding unwanted values (an extremely CPU intensive operation), an external trigger tells the FPGA to begin writing values to the FIFO at the same as the PC starts to read the values, thus preventing pile-up.

Here's where it gets interesting: some aspect of this FIFO startup is non-deterministic and asynchronous, likely depending on current CPU load and/or interrupt time. If the same PXI trigger is sent to the FPGA module and the dedicated digitizer modules, the DMA FIFOs on the FPGA start and end late, with a quasi-random delay compared to the other digitizers. This is clearly problematic. The adopted solution instead has one of the digitizer modules generate an extremely specific TTL-encoded digital sequence⁴, where the sequence starts at the same time as the first sample of the properly-triggered digitizer. This synchronization sequence drives a digital input on the FPGA module, so that the FPGA can correlate the measured sequence with the expected sequence, essentially allowing

⁴The sequence is a UTC-8 encoded byte string that says "At least it's not barium tagging".

a determination of the first sample's position in time. If the FPGA is triggered prior to the other modules and set to digitize a little longer than would normally be necessary, then it records all of the relevant data plus some extra samples before and after the nominal digitization period, as set by the dedicated digitizer modules.

Validating the synchronization and cropping the excess FPGA data happens offline, just prior to analysis. The raw data files stored include the excess samples.

D.4.2 Digital Demodulation

Generally speaking, heterodyne interference measurements involve mixing two higher frequencies down to a moderate difference frequency, where the amplitude and phase of the signal oscillating at the difference frequency encode the desired information to be measured. The first mixing occurs in the interference between the reference wavefronts and the transmitted/retroreflected wavefronts, shifting the signal from PHz down to \sim kHz. Thus, measuring the amplitude and phase of an oscillating signal are critical for the work described here, which have been accomplished with a technique called digital demodulation.

Consider a generic oscillating signal with amplitude A , frequency f , and arbitrary phase ϕ : $A(t) = A \cos(2\pi ft + \phi)$, which, in our case, can represent an oscillating voltage from the output of a photodetector and transimpedance amplification circuit. Now, assume that this signal is phase-locked with the sample clock of an ADC, which can easily be accomplished by sharing timebases between the signal generator(s) and the ADC, with built-in hardware. If the ADC is sampling at a frequency f_{ADC} such that $f = 4(f \% f_{\text{ADC}})$ where $\%$ is the modulus operator, then successive ADC samples A_i are spaced by $(2\pi n + \pi/2)$, with n an integer, and thus,

$$A_i = A \begin{cases} \sin(\phi) & i = 1, 5, \dots \\ -\cos(\phi) & i = 2, 6, \dots \\ -\sin(\phi) & i = 3, 7, \dots \\ \cos(\phi) & i = 4, 8, \dots \end{cases}, \quad (\text{D.1})$$

which is depicted schematically in Figure D.1.

The amplitude and phase, i.e. the variables encoding our desired measurement, can then

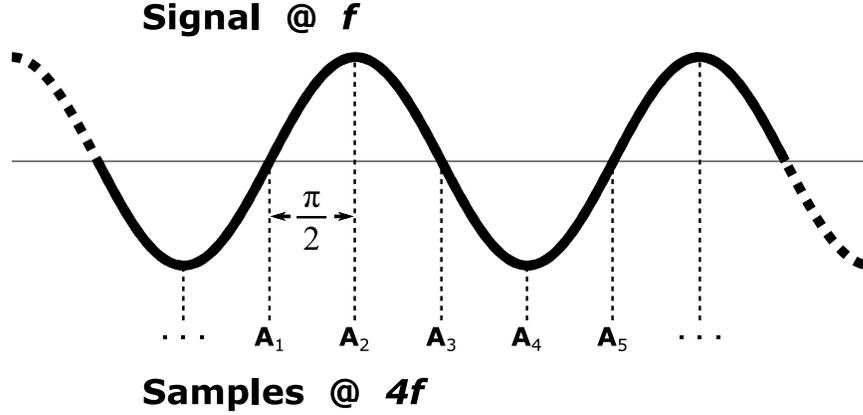


Figure D.1: A cartoon depiction of digital demodulation. The arbitrary phase ϕ from Equation. (D.1) was set to 0 for simplicity.

be recovered from the signal samples as,

$$A = \sqrt{A_i^2 + A_{i+1}^2} \quad \phi = \begin{cases} \arctan\left(\frac{-A_i}{-A_{i+1}}\right) & i = 1, 5, \dots \\ \arctan\left(\frac{-A_{i+1}}{-A_i}\right) & i = 2, 6, \dots \\ \arctan\left(\frac{-A_i}{A_{i+1}}\right) & i = 3, 7, \dots \\ \arctan\left(\frac{A_{i+1}}{A_i}\right) & i = 4, 8, \dots \end{cases}, \quad (\text{D.2})$$

making use of the identities $\sin^2 + \cos^2 = 1$ and $\tan = \sin / \cos$. If the amplitude, A , and phase, ϕ , are time-dependent, but their changes are slow compared to the carrier frequency f , then these variables provide robust estimates of the amplitude and phase of the signal of interest, in real-time. These estimates of A and ϕ are computed within the FPGA that is integrated with the ADCs sampling the signal. The details of this implementation are discussed in the next section.

This procedure naturally depends on the spectral purity of the signal to be demodulated. Significant oscillating components at different frequencies will still “make it through” the digital demodulation algorithm, introducing offsets and/or oscillating estimates of the amplitude and phase. The amplifier circuits for the QPD and the retroreflected photodiode include significant active filtering to ensure the output is nominally a single tone, although the bandwidth is not particularly narrow.

D.4.3 FPGA Signal Processing

There are five oscillating signals that are demodulated: four from the quadrant photodiode, and a single one from the retroreflected photodiode. An identical demodulation procedure is used for each of the five signals, so the following discussion will pertain to one such signal. Firstly, each oscillating signal is digitized with a 16-bit differential ADC, resulting in a signed 16-bit integer for each sample. Although reconfigurable, the ADC is usually fixed to sample at 500 kHz, for a phase-locked interference signal at either 125 kHz or 625 kHz. There is some inherent and common-mode phase shift between the sample clock and the interference signals, likely from the synchronization of timebases, but this is unimportant for the horizontal degrees of freedom, while the “origin” of the vertical degree of freedom is determined empirically as discussed in Section 2.6.

The sampling rate of the ADC drives all of the successive operations within the FPGA, so that they all operate at 500 kHz or integer divisions of this sampling frequency, as calculations based on these signals need only occur each time there is a new sample. Generally speaking, within the FPGA there are a number of looping operations that run in parallel and share data with one another. For example, the simultaneous sampling happens within a single loop, with samples being sent to a demodulation loop, the output of which is sent to a downsampling loop, which in turn is sent to the feedback and data transfer loop, all of which are operating synchronously, triggered by data transfer from a previous loop in the pipelined analysis chain. Generally, we followed guidelines and suggestions, i.e. data pipelining and parallel execution, provided by National Instruments in order to make the best use of the FPGA [224].

Since estimates of the amplitude and phase, Equation (D.2), require neighboring samples, the FPGA retains two samples at all times, from the current timestep i , A_i , as well as the previous timestep, A_{i-1} . The amplitude is estimated directly by computing the quantity $A = \sqrt{A_i^2 + A_{i+1}^2}$, where the 16-bit integers are first converted to a datatype with a larger bit-width to accommodate the possible increase in dynamic range and/or requisite numerical precision⁵. To estimate the phase, a running modulo-4 counter is maintained to keep track of which sample is in the numerator/denominator of the argument of the arctangent from

⁵National Instruments and LabVIEW makes use of a fixed-point datatype, wherein numbers are represented by a fixed number of bits, with some being dedicated to representing the integer portion of the number, and the other being dedicated to representing the decimal portion of the number. The fixed bit-width makes mathematical operations more tractable, increasing data throughput significantly. Although hugely important, we leave the details of understanding the fixed-point datatype to the interested reader.

Equation (D.2), and use the Coordinate Rotation Digital Computer (CORDIC) algorithm to rapidly compute the arctangent with the appropriately signed and ordered ratio of the successive samples.

The above procedure results in 10 distinct quantities, the amplitudes, $\{A_0, A_I, A_{II}, A_{III}, A_{IV}\}$, and phases $\{\phi_0, \phi_I, \phi_{II}, \phi_{III}, \phi_{IV}\}$ of each of the interference signals, with ‘0’ for the retroreflected photodiode signal and as before for each quadrant of the QPD: ‘I’ for $x, y > 0$, ‘II’ for $x < 0 \wedge y > 0$, ‘III’ for $x, y < 0$, and ‘IV’ for $x > 0 \wedge y < 0$. These quantities are computed at the same rate the signals are digitized, i.e. 500 kHz, although shifted in time by $3 \times dt$ where $dt = 1/(500 \text{ kHz})$ is the basic unit of time for the programmed signal processing in the FPGA, due to the use of data-pipelining as discussed in Reference [224].

The amplitudes and phases are used to estimate x , y , and z as

$$x = \frac{A_I + A_{IV} - A_{II} - A_{III}}{\sum_i A_i} \quad \text{and} \quad y = \frac{A_I + A_{II} - A_{III} - A_{IV}}{\sum_i A_i} \quad \text{and} \quad z = \phi_0. \quad (\text{D.3})$$

These estimates are further processed in a down-sampling loop, which computes the sum of 100 estimates and then divides the result by $2^7 = 128$ to approximate an average, as fixed-point arithmetic with powers of 2 is exceedingly fast, involving only a shift of the binary decimal point within the representation of the number itself. Once 100 estimates for each coordinate are collected and an approximation of their average computed, the down-sampled signals are sent to the feedback loop, which necessarily operates $\times 100$ slower, at a rate of 5 kHz. Thus, the previous delay of $3 \times dt$ represents a negligible phase shift on the down-sampled signal, which is used for feedback and offline analysis.

D.4.4 PID Feedback

A general implementation of PID feedback was discussed in Section 2.6, together with an approximate method for initializing feedback on a newly trapped microsphere. Here, the FPGA implementation of the various feedback terms will be discussed. It should first be noted that although the exact datatype (most importantly, the bit-width, which is representative of the dynamic range) has not been explicitly mentioned throughout, great care was taken so that numerical precision is preserved throughout the signal processing. Basically this amounts to using a datatype with a larger bit-width than necessary, applying any necessary digital scaling to the numerical value, and casting the final result to whatever precision is required. This practice certainly makes non-optimal use of the FPGA resources,

but protects against the loss of information and precision for poorly handled digital typing.

In this apparatus, the initial signals are digitized with 16-bit resolution, which in practice usually amounts to an actual precision somewhere between 14-bit to 15-bit. The amplitude estimates can have higher precision than the initially digitized values given the square-root operation, while the CORDIC algorithm has a user-defined precision, as it's implemented with a look-up table that is built when the code is compiled for the FPGA. LabVIEW automatically determines the necessary output datatype for any mathematical operation given the 16-bit input signals, but the automatic typing can be somewhat opaque and difficult to adapt. Additionally, further arithmetic operations may require higher dynamic range, so for these intermediate steps, the amplitude and phase estimates are cast to a 32-bit fixed-point datatype, which in turn is reinterpreted as a 32-bit integer⁶. It is these 32-bit integers that are down-sampled and subsequently processed for feedback.

As in Section 2.6, the calculated estimates of $x_i = \{x, y, z\}$ are compared to user-defined setpoints, $x_{0,i}$ in order to generate error signals $e_i(t_n) = x_i(t_n) - x_{0,i}$, where the time argument has been explicitly discretized from $t \rightarrow t_n$ to indicate the finite sampling rate. With these error signals, generating proportional feedback amounts to simple arithmetic operations, where either input (the error signal or the proportional gain $K_{P,i}$) can be scaled arbitrarily to prevent both overflow and underflow.

The integral feedback is implemented with an infinite-time accumulator, wherein successive values of the error signal $e_i(t_n)$ are added to a running sum (stored as a 64-bit integer) which serves as the integral in Equation (2.6). The running sum can be set to 0 arbitrarily to 'reset the memory' of the integral feedback, but this usually need only be done during the initialization of a trapped microsphere when the z signal is changing significantly while the microsphere is lowered into position near the focus. The running sum is then multiplied by the integral gain $K_{I,i}$ to generate integral feedback, with the same flexibility in numerical scaling to preserve information content.

Finally, the derivative of the signal is approximated by a finite difference between successive samples: $(dx_i/dt)_{t_n} \sim dx_i = x_i(t_n) - x_i(t_{n-1})$, where the timestep is absorbed into any arbitrary scaling applied. Feedback generated by this approximation, and pure derivative feedback in general, is susceptible to ringing when phase shifts become non-negligible near

⁶It should be noted that the position of the decimal point within the fixed-point datatype is somewhat irrelevant since the calibration from observed displacements to applied forces is determined empirically, so the amplitudes and phases are usually 'reinterpreted' as 32-bit integers, wherein their bit patterns remain the same, but the numerical value is scaled arbitrarily.

the feedback loop's highest operating frequency. To slightly mitigate this risk, two values of the finite difference are averaged together, so that the derivative feedback term looks like $K_{D,i}(x_i(t_n) - x_i(t_{n-2}))$. A more astute solution would be the approximation of a derivative generated from the simple linear regression to a limited number of data points (say, 3-4 points in total), since linear regression amounts to just a matrix operation which can be performed quite rapidly on an FPGA.

All together, the PID signal for a single coordinate at time t_n

$$PID_i = K_{P,i} e_i(t_n) + K_{D,i} [x_i(t_n) - x_i(t_{n-2})] + K_{I,i} \sum_{k=0}^{\infty} e_i(t_{n-k}). \quad (D.4)$$

Once this signal has been computed, with all of the appropriate scaling to prevent over- and underflow, the result is cast to a 16-bit signed integer in order to conform to the DAC limitations. The output feedback signals for x and y are sent to an active level control circuit, which converts signals of ± 10 V to signals over the range 0 V to 5 V, which is compatible with the controller for the piezoelectrically driven deflection mirror. The output feedback signal for z is sent through a passive voltage divider to reduce the range from ± 10 V to ± 1 V, compatible with the amplitude modulation input on the SG382, which drives the trapping beam AOM and thus controls the trapping beam power. Both of these level-control circuits are intended to make the most effective use of the dynamic range of the DACs, using every bit to span the full range of control.

D.4.5 Data Acquisition for Offline Analysis

Once the estimates of x , y , and z have been computed, reinterpreted as 32-bit integers, and finally down-sampled, they are written to DMA FIFO buffers for access by the host computer, together with the 16-bit values generated by the PID feedback described previously. The amplitudes and phases of the interference signals, as well as digitized signal from the power photodiode, are also down-sampled by the same factor and written to DMA FIFOs. Due to the intricacies of the actual implementation, LabVIEW only allows for the construction of FIFO buffers with a length equal to a power of 2, although any number of samples may be drawn from the FIFO during a read operation, and successive read operations can occur with 0 sample delay in order to collect long, continuous digitizations. Additionally, the FPGA module within this apparatus has a limited number of DMA FIFOs (3 transfer channels, compared to ~ 20 unique signals) that can be used to traffic data to the host

PC. Thus, the samples of x , y , z , the feedback signals, the amplitudes and phases of the interference signals, and the power are actually interleaved into three different FIFOs, not to be de-interleaved until offline analysis following transfer to the host PC. The data are also interleaved with a timestamp from the FPGA's 64-bit time counter, so that proper de-interleaving can be verified by finding the regularly spaced, monotonically increasing "timestamp signal".

When a user requests some data to be acquired through the LabVIEW interface, the host PC first configures the dedicated digitizers (setting up the sample clock and specifying the number of desired samples). Once these modules are ready and waiting for a trigger, the host PC starts up the DMA FIFOs and begins extracting data from the FPGA. Immediately after the FIFO read begins, the PXI chassis triggers the dedicated digitizer modules to begin collecting data, as well as sending the synchronization signal discussed previously to the FPGA. Once the primary digitization task is complete, the FPGA continues storing values briefly to guarantee overlap, before stopping the FIFOs and resuming normal operation. During this acquisition period, data is stored within the RAM of the host PC, only to be written to disk when the acquisition period has finished. This particular architecture naturally sets a limit on the longest digitization period achievable, as the RAM will eventually fill at which point LabVIEW encounters an error and ceases execution.

For both the dedicated digitizer modules and the FPGA, data is written to disk in a binary format, stored within two HDF5 files: one for the dedicated digitizer modules and environmental monitors (managed by the same LabVIEW VI) and one for the FPGA-acquired data. A few text strings are stored within an XML-encoded "attribute" file, specifying a PC-derived timestamp for general sorting, as well as the requested sampling rate, number of samples, and other defining features. Thus for each requested integration, there are three associated files with identical base-names. The excess integration time required by the FPGA, as well as the requisite time for all modules to save their data to disk, ends up requiring approximately 2 s to 3 s between each continuous integration, although more adept LabVIEW programmers can likely reduce this with proper handshaking as opposed to hard-coded time delays.

D.4.6 Secondary FPGA for Rotational Degrees of Freedom

In addition to the primary FPGA module within the PXI chassis, there is a secondary FPGA module with a PCIe form factor that sits directly inside the host PC. This secondary

FPGA is dedicated to generating the DDS electrode drive signals discussed previously, as well as digitizing and demodulating the single output from the cross-polarized light monitor (see Section 2.5). The architecture of the digitization and demodulation is identical to that which is implemented within the primary FPGA, just limited to a single channel instead of five, as the secondary FPGA module is slightly older and contains hardware with significant limitations as compared to the primary.

D.5 Digitization

The primary translational motion and power signals (x , y , z , and \mathcal{P}) are necessarily digitized by the FPGA as part of the active feedback algorithm, and transferred to the host PC via a DMA system. The remaining signals of interest are the cross-polarized light monitor, electrode and stage monitor signals, various environmental variables, as well as the occasional extra photodiode for debugging or alignment purposes. The electrode and stage monitors provide information about the known (or unknown, but of interest to investigations of new physics) forces being applied to a trapped microsphere and are thus critical for force-sensing applications, while the environmental variables are used to better understand background forces that may be influenced by conditions around the apparatus.

Generally speaking, the techniques and experiments discussed in this thesis make use of continuous excitation and measurement. Given that infinite length data files are not tractable, the digitization is segmented in a scheme where signals are acquired continuously for a short length of time, usually 10 s, and then written to a binary file. The various drive signals and digitization tasks are then reset and re-executed as many times as is necessary to perform the measurement. Both types of tasks are triggered simultaneously by dedicated trigger lines within the backplane of the PXI chassis itself.

D.5.1 Fast Monitor Signals

All of the “fast” signals are digitized by a single PXI module, the PXI-6259, which also includes analog and digital output channels as a sort of all-in-one solution. The acquisition portion of the module consists of a pair of 16-bit ADCs, each capable of operating at 1 MHz. Multi-channel acquisition is then accomplished by multiplexing the desired inputs, an operation that happens entirely within the module itself. For example, if 4 channels are required, then each can be digitized with a maximum frequency of 500 kHz given the bandwidth

limitation of the underlying ADCs. Most often, the system is operated with digitization frequencies of 5 kHz so this rarely presents an issue. For the electrically driven rotation discussed in Chapter 5, we are interested in much faster signals of order ~ 100 kHz, which limits the number of channels that can be simultaneously digitized. This is not an issue usually, as information about the attractor stage position is unimportant to the electrically driven rotation measurements, as a specific example.

Regardless, all signals of interest are input to the digitizer module, and when the user configures an acquisition task within LabVIEW, a subset of the signals can be chosen for digitization, leaving the others terminated. For force-sensing measurements, the attractor stage monitors and electrode drive monitors are critical, so these are all digitized, usually at the slower rate of 5 kHz used to investigate the translational motion of a trapped microsphere. If the microsphere is also rotating, the cross-polarized light monitor may contain useful information, although the rotation frequencies, both driven and observed in the absence of any drive, are often well above this bandwidth. For rotation measurements, usually the cross-polarized light monitor and a single electrode drive signal are digitized, where the voltage on the other three electrodes used to generate the spinning field are inferred from the main drive signal, given that they are generated by the same DDS module, with common-mode scale and symmetric amplification. If the rotation frequency is sufficiently low, ≤ 25 kHz, then all electrode drive signals are digitized.

The acquired signals are then stored in a single binary file separate from the FPGA data, then synchronized to the FPGA data during offline analysis following the procedure outlined in the previous section.

D.5.2 Slow Environmental Variables

Two environmental variables may have a non-negligible effect on measurements made with this apparatus: temperature and atmospheric pressure. The temperature of the lab is nominally controlled by an HVAC system to within 0.1 °C, although this control results in cyclical changes of the temperature,⁷ which have been found to be correlated with certain measurement backgrounds. The air temperature affects air density (and thus index of refraction) which changes the optical path length of the light in the free-space optics. Temperature

⁷When the lab warms above the upper threshold of the thermostat's setting, the HVAC turns on and cools the lab to lower threshold, at which point it turns off and the lab begins to heat again. This repeats indefinitely.

fluctuations of the mechanical components themselves can also induce thermal expansion, changing the physical distance between various components of the apparatus. Thankfully, the thermal conductivity of air-matter interfaces is such that the cyclical temperature fluctuations of the air only couple very weakly to the temperature of the mechanical components. Atmospheric pressure fluctuations have a similar effect to air temperature fluctuations, as they change the refractive index within the free-space optics and change the optical path length.

Historically with this apparatus, the temperatures of the air and of the optical table itself were measured with thermocouples and a dedicated thermocouple readout instrument. Temperature values were then extracted from the readout instrument via GPIB. Atmospheric pressure was not actively monitored, although superstitious graduate students would often check local weather conditions in attempt to transfer blame of a non-functional apparatus to uncontrollable environmental conditions.

With the development of the newest iteration of the apparatus, the environmental variable readout was updated significantly. Both the air temperature as well as the chamber temperature are now measured with thermistors in a Wheatstone bridge configuration (see Appendix E). The outputs of these thermistor preamplifiers are sent to a cRIO module controlled by a dedicated PC. Similarly, an amplified pressure-transducer signal is sent to the same module. Measurements of temperature and pressure are then acquired at regular intervals, timestamped, and stored on this dedicated PC. Offline analysis can then make use of this data, synchronizing it based on the universal timestamping used for all acquired data.

D.6 Comment on Details Not Included

A large set of details was not included in this Appendix most notably any information regarding the various interconnects involved in operating the wide range of hardware. This includes both the connection of electrical signals from high-density connectors to individual coaxial cables (usually referred as “break-out” boxes), as well as the vacuum feedthroughs to bring electrical signals inside the main experimental chamber. The solutions implemented in this apparatus are varied, often custom in nature, and underwent many changes as system capabilities were expanded. Thus it doesn’t serve to explain them here.

Appendix E

Custom Analog Circuits

Wherever possible in this apparatus, consumer/off-the-shelf electronics were used, as these tend to be reliable and well-tested solutions to known problems. Occasionally, however, a specific amplification or coupling behavior is required that may not be commonplace and thus has no consumer solution. In these few instances, we devised and constructed our own circuits.

For a few of the earlier circuits, a very basic schematic and printed circuit board (PCB) layout program suite was used: ExpressSCH and ExpressPCB, the latter also being the name of the board manufacturer. These programs have a few features that make them attractive as an entry-level design suite, including a built-in online ordering system, as well as layout tools that only allow the user to draw features that are actually manufacturable. However, earlier versions of the software did not include automatic PCB and schematic linking and synchronization, and custom board sizing beyond an automatic standard was not feasible.

As a result of these (relatively minor) limitations, circuit design was transitioned to EAGLE, part of the Autodesk suite. EAGLE is incredibly sophisticated by comparison and has capabilities well beyond anything required. Design files for the circuits documented here can be requested from the author. This appendix will only include surface level documentation, given that representation of full circuit schematics or PCB layouts in a thesis formatted for the standard US-letter sized paper is highly suboptimal.

E.1 Direct QPD Readout

The transmitted light interferometry discussed in Chapter 2 makes use of a quadrant photodiode (QPD) in order to image the radial motions of a trapped microsphere. The expected signal is an interfering optical power, which results in an oscillating photocurrent on each of the QPD's segments, expected at 125 kHz or 625 kHz. In the scheme developed for this thesis work, each quadrant is individually AC-coupled and amplified, with the resulting interference signals sent directly to demodulating analog-to-digital converters (ADCs). This is accomplished with a single circuit board to which the QPD is directly mounted.

The QPD itself is in a common-cathode configuration, so that all quadrants can share a reverse bias while producing four distinct photocurrent signals. Each quadrant's photocurrent is first sent through a transimpedance amplifier, which includes a resistive and inductive shunt in a T-configuration with a AC-coupling capacitor at the non-inverting input of the op-amp in order to separate the DC component of the interference from the desired oscillating portion of the interference. This first stage is shown schematically in Figure E.1, together with a plot of the simulated frequency response. In the passband, the transimpedance gain is dominated by the feedback resistor, yielding $G_t = 10^4 \text{ V/A} (= 10 \text{ k}\Omega)$.

Following the initial photocurrent-to-voltage conversion, the oscillating voltage is further amplified by an AC-coupled, non-inverting amplifier, which includes a precision variable resistor to tune the gain of the circuit, from $\leq 1 \text{ V/V}$ up to 50 V/V , and a compensation capacitor in the feedback path so that this final stage has low-pass filter behavior for anti-aliasing. The output of the second stage is sent to an board-mounted SMA connector, which is in turn connected to the analog input of the FPGA. A screenshot of a portion of the full EAGLE schematic is shown in Figure E.2, detailing this second amplification stage.

The two-stage amplification topology is duplicated a number of times so that the final board has four independent readouts. The board also includes an excitation input and jumpers, so that a known current can be injected into the transimpedance amplifier input in order to tune and calibrate the electronic gain. These gains have been set so that with a typical microsphere levitated near the focus, an oscillating signal with an amplitude of a few volts is output from the readout board. This takes advantage of most of the dynamic range of the $\pm 10 \text{ V}$ ADC inputs, while also avoiding any saturation effects from operating closer to the voltage limits.

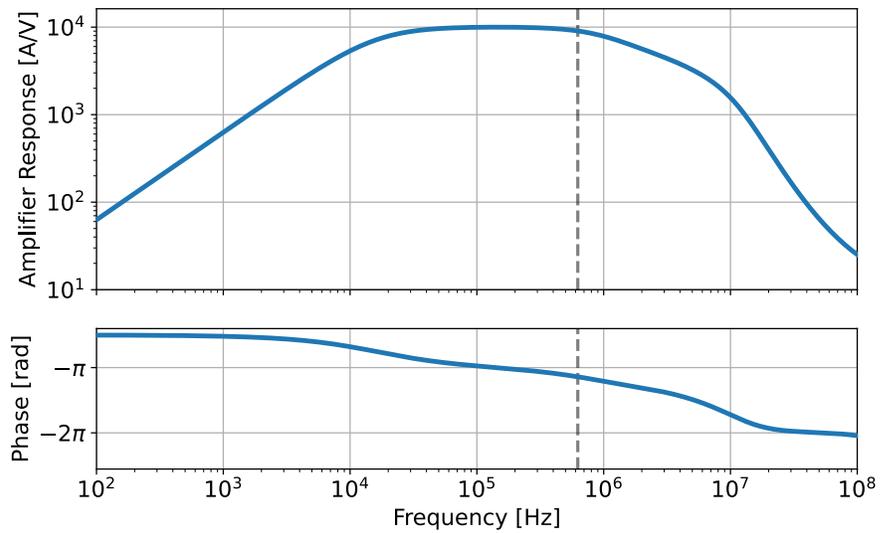
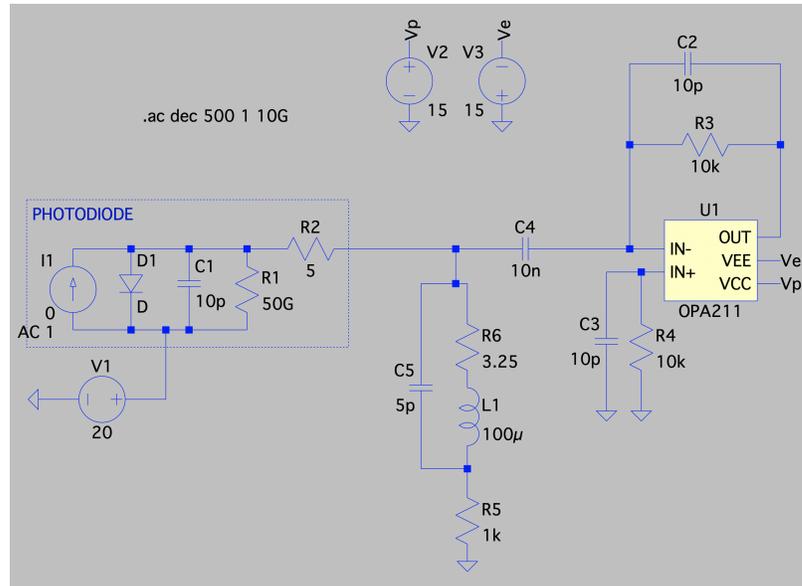


Figure E.1: (upper) An LTSpice implementation of a single QPD quadrant and the first stage of amplification, including a realistic model of both the photodiode itself, as well as the inductor used to shunt the DC component of the photodiode signal. (lower) Bode plot of the amplitude and phase of the amplifier response as a function of frequency. A dashed vertical line indicates the chosen interference frequency of 625 kHz.

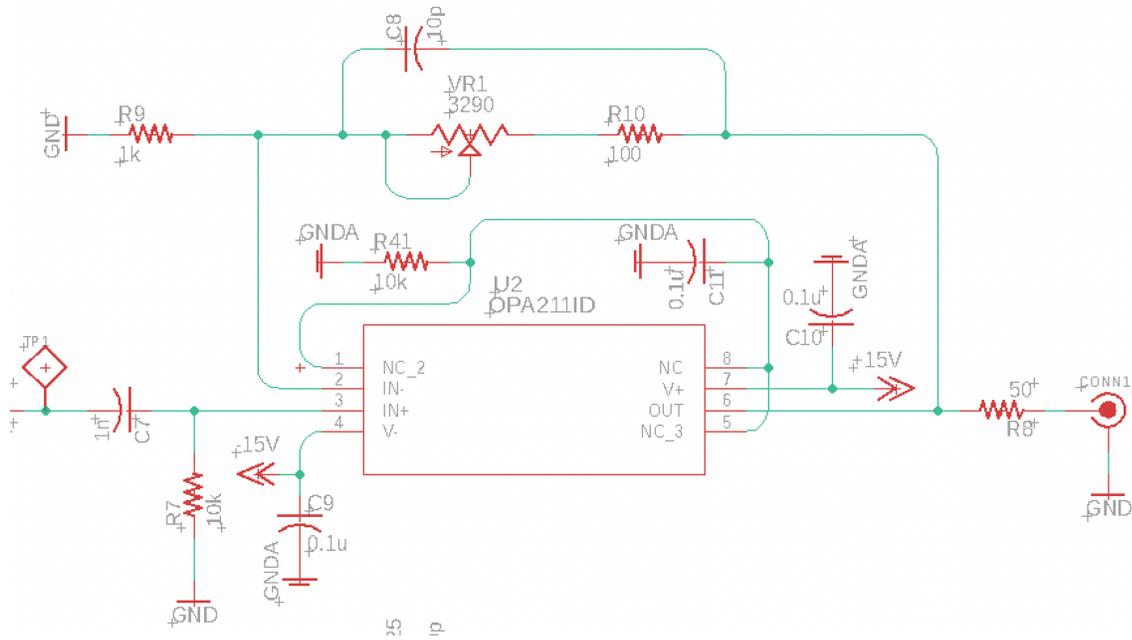


Figure E.2: The second stage of amplification in the direct QPD readout board. An AC-coupled, non-inverting amplifier, with low-pass behavior for anti-aliasing. “No Connection” (NC) terminals of the op-amp have been terminated with a resistor and capacitor in parallel, following guidelines in the manufacturer datasheet.

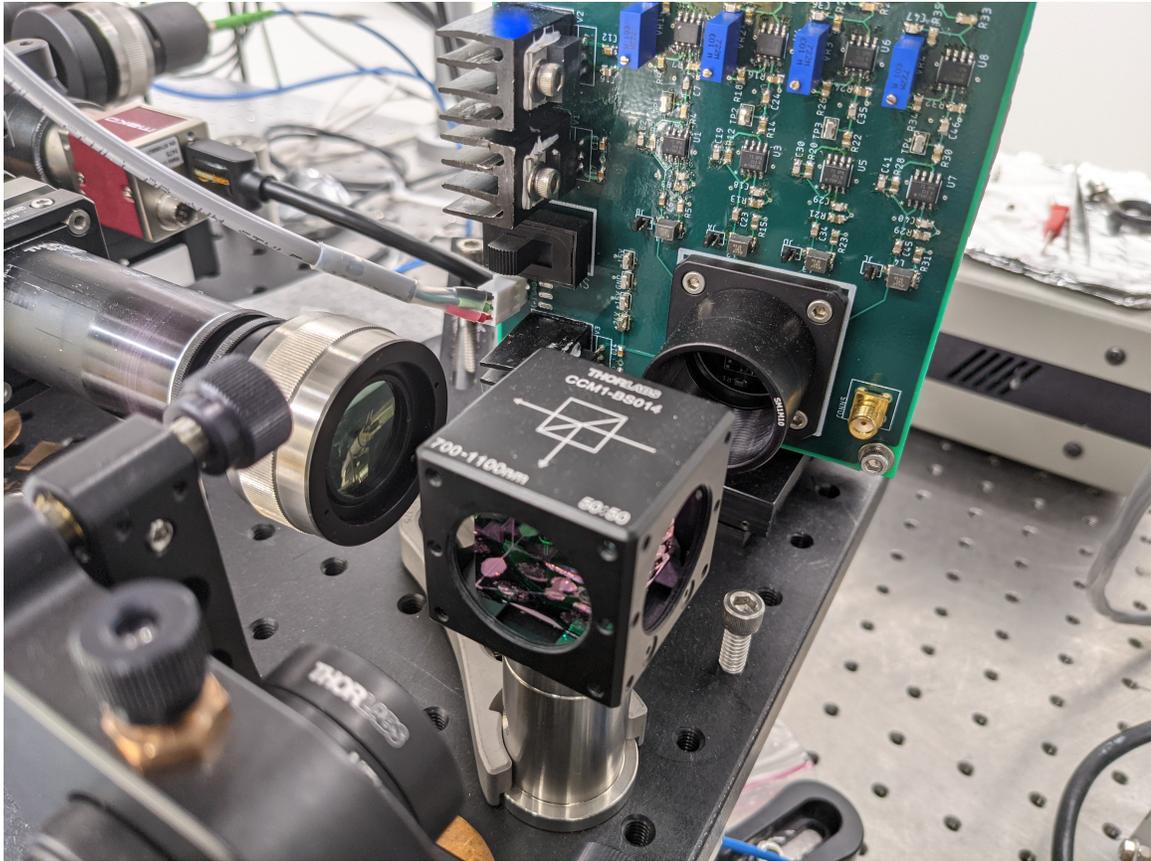


Figure E.3: A photograph of the QPD circuit, with the output beam reduction telescope visible on the left, as well as the reference beam collimator visible in the lower left. A superposition of the transmitted light and reference beam light is generated with the beamsplitter cube, and subsequently projected onto the QPD, which has a lens tube serving as shroud in this image.

E.2 Retroreflected Interferometry Readout

Naïvely, one might be inclined to use an identical amplification scheme, now with just a single channel, for the retroreflected photodiode signal. However, this is not the approach of the currently implemented readout, as the expected power of the retroreflected light is much smaller than that of the transmitted light.¹ It turns out that electronic pickup (as of yet untraceable) or anomalous pure tones in either the reference beam AOM's synthesizer or the associated RF amplifier become non-negligible components of the interference signal when the reference beam has significantly larger optical power compared to the retroreflected beam. Unfortunately, these frequencies are somewhat close to the chosen difference frequency (for example, one tone has $f_{\text{pickup}} \approx 88 \text{ kHz}$ while $\Delta f = 625 \text{ kHz}$), so a reasonably broad passband, such as that used for the QPD readout, also amplifies these anomalous signals, contaminating the measurement.

The first stage is identical to that within the QPD readout as seen in Figure E.1, as this AC-coupled first stage is still necessary to shunt the large DC component of the photodiode signal (now much larger given the mismatch between reference beam and retroreflected beam powers). The oscillating voltage from this first stage is then further amplified by an LC resonant circuit. The quality factor of this circuit is intentionally lower than one might expect, so that phase modulations of the interference signal, i.e. the portion of the signal encoding the z -position of a trapped microsphere, are not significantly altered or attenuated. This is accomplished with a non-inverting amplifier configuration, where an inductor and capacitor are placed in parallel with the feedback resistor. An LTSpice schematic and simulated frequency response are shown in Figure E.4. The output of this circuit is sent directly to a demodulating ADC within the FPGA, in order to recover the phase of the interference signal and thus the z -position of a trapped microsphere.

¹This is a result of both the low reflection coefficient, as well as a limited collection efficiency due to apertures inherent to the isolator.

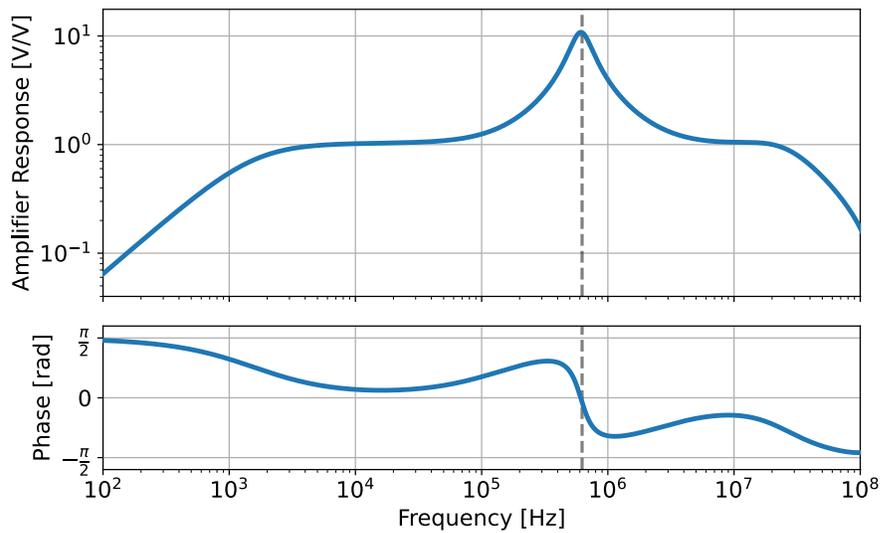
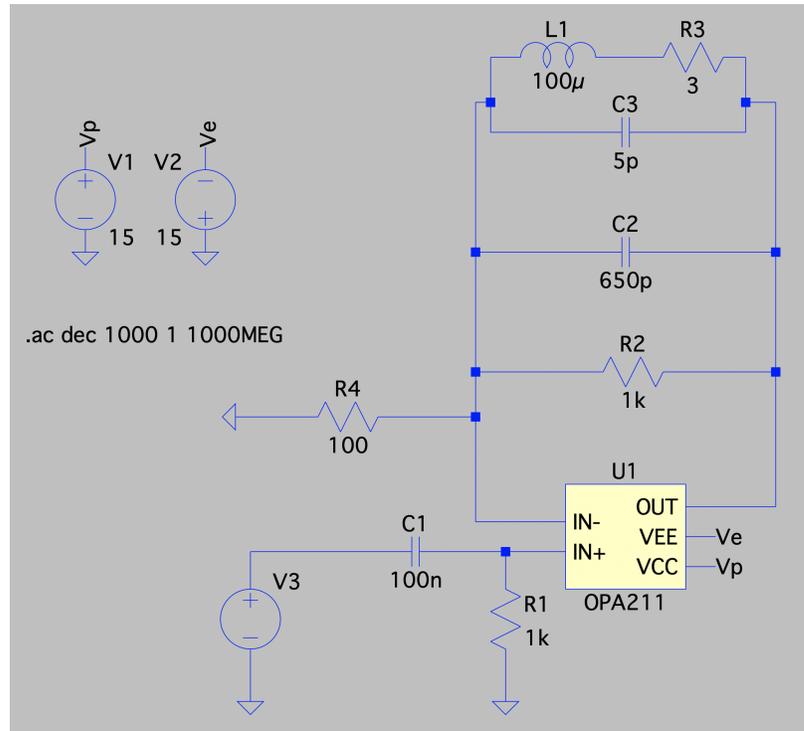


Figure E.4: (upper) An LTSpice implementation of the resonant non-inverting amplifier used as a second stage of the retroreflected interferometry readout (lower) Bode plot of the amplitude and phase of the amplifier response as a function of frequency. A dashed vertical line indicates the chosen interference frequency of 625 kHz.

E.3 Feedback Dynamic Range Control

To take full advantage of the dynamic range of the DACs operating within the active feedback loop, a level matching circuit is necessary. The level matching circuit is designed to accept the 16-bit, $\pm 10\text{ V}$ output from the DACs, scaling it down by some factor and adding an offset, the latter feature being necessary for applications requiring positive-definite voltages. This is accomplished with relatively simple and well-understood circuit topologies. The FPGA's output is first buffered with an instrumentation amplifier operating in a unity-gain configuration. The instrumentation amplifier drives a passive voltage divider which includes a precision variable resistor to provide control over the scaling factor. The output of the voltage divider is further buffered by a non-inverting op-amp with unity-gain, where the op-amp chosen (Texas Instruments, OPA211) is designed to be explicitly stable in this configuration. The scaled down and buffered output is sent to a simple summing amplifier composed of an op-amp in an inverting configuration (suggesting that we must be careful with our factors of -1).

The summing amplifier is used to add a DC bias to the scaled output from the FPGA's DACs. A tunable offset in the range -5 V to 5 V is first generated with a pair of precision 5 V voltage references. This is accomplished by tying one reference between ground and the positive voltage supply, and tying another between ground and the negative voltage supply. Constant current is drawn through each of the references in order to maintain stable and continuous operation, while a variable resistor positioned between the referenced outputs allows selection of an offset voltage anywhere within the range $\pm 5\text{ V}$. A schematic of one such level-matching channel is shown in Figure E.5.

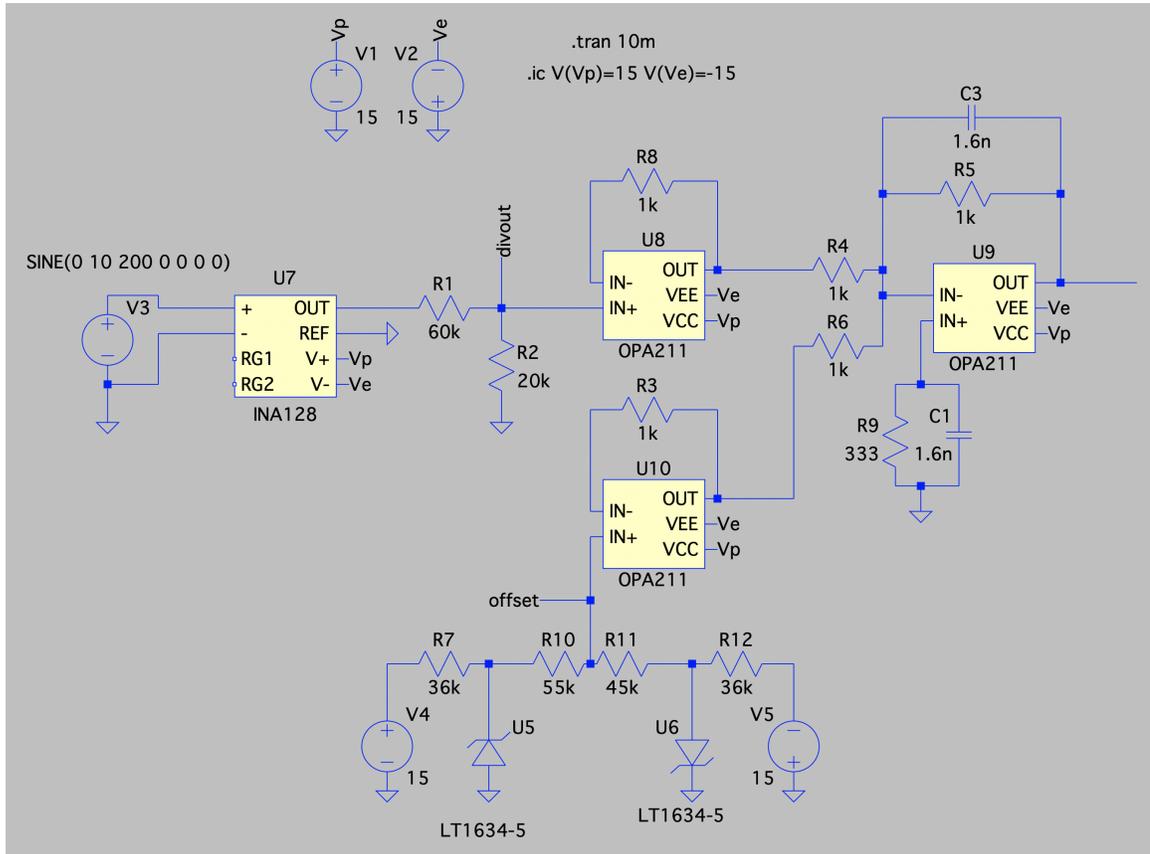


Figure E.5: A schematic of the level matching circuit implemented to take full advantage of the dynamic range of our analog drivers. The output voltage from a DAC within the FPGA is first buffered by an instrumentation amplifier, whose output drives a resistive divider in order to scale down the amplitude of the voltage. The output of the resistive divider is buffered by a unity-gain op-amp and sent to one of two inputs of a summing amplifier. A tunable DC bias is generated with a pair of precision 5 V voltage references, whose outputs are connected across a variable resistor such that the wiper terminal of the variable resistor can be anywhere in the range ± 5 V.

E.4 Thermistor Readout

Environmental temperature monitoring, both of the air and of the optical tables, is accomplished with thermistors, passive electronic components whose resistances change deterministically with temperature. There are a variety of methods to measure these changes in resistance, and we make use of a Wheatstone bridge resistor configuration, capable of sensitive measurements via the differential readout. In a Wheatstone bridge, a constant voltage source drives two parallel paths of resistors, where each path consists of two resistors in series. Assuming all of the resistances are identical to begin with, subsequent changes in resistance result in a differential voltage between the midpoints of these two paths. Both resistances along one path are assumed to be known, while the second path consists of a variable resistor and the thermistor itself. The known resistances and the midpoint value of the variable resistor are all set to $10\text{ k}\Omega$, while the thermistor (Omega, SA1-TH-44006-40-T) also has a nominal resistance of $10\text{ k}\Omega$ at room temperature.

As changes in temperature raise and lower the thermistor's resistance, the differential voltage across the bridge changes, where a 0 V signal is expected at room temperature. This differential voltage is then measured by an instrumentation amplifier and then output to the dedicated environmental monitoring computer. The variable resistor can be adjusted to tune an offset voltage, which is necessary given that the cRIO module for the environmental monitoring computer (see Appendix D) has ADCs with a range of 0 V to 5 V . The PCB for the thermistor readout includes 5 channels, and Figure E.6 details a schematic of one such channel.

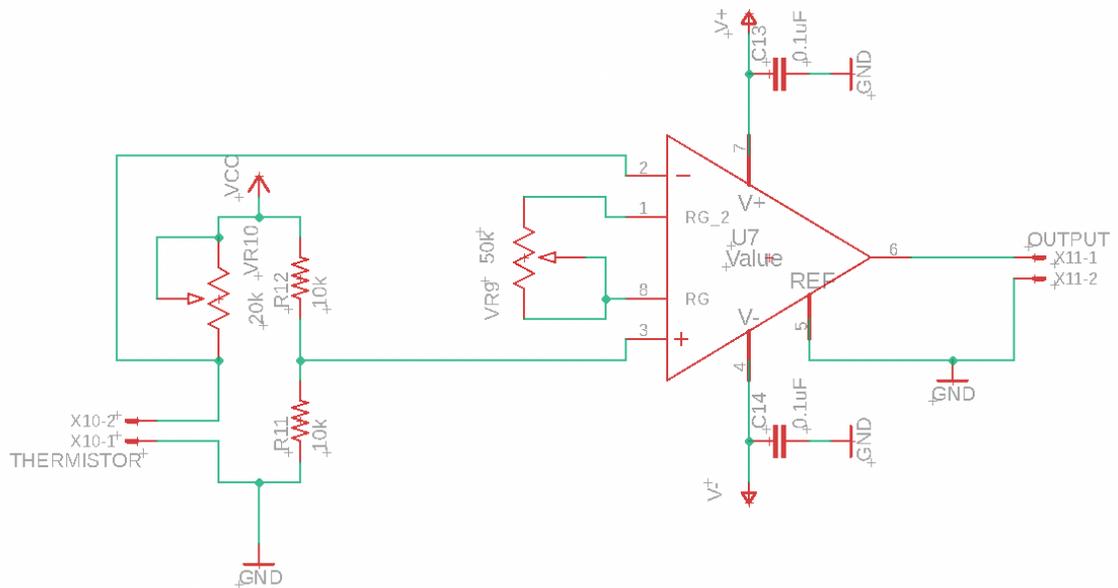


Figure E.6: A schematic of the Wheatstone bridge readout for a single thermistor. The differential voltage is measured and amplified by an instrumentation amplifier, and is expected to be 0 V at room temperature and with the variable resistor set to 10 k Ω . The bridge can be biased to produce a positive-definite voltage by adjusting this variable resistor.

Appendix F

The Cosmological Constant

Although there are a variety of theories that can result in modifications to the inverse square law of Newtonian gravity (see Chapter 6), some of the most compelling are those that can answer open questions in physics, of which there are many. Of particular interest to the author are those theories that can bridge cosmology and with small-scale physics, even if only in some small way.

Cosmology can be thought of as the study of both the origin and fate of the observable universe, in the broadest sense possible. The basic picture of cosmology is built upon solutions to Einstein’s equations of General Relativity, and posit that the universe is a continuous fluid with just a few distinct constituents. The different components of the fluid (to be formally introduced below: matter, dark matter, dark energy) interact only through their effects on the nature of our spacetime.

This description has proven to very closely match observations of the universe made from Earth, but necessarily invokes the existence of dark matter and dark energy, both which remain elusive and have not been “detected” aside from their presumed effect on the nature of our universe at the largest scales.

F.1 Historical Context

Shortly after his original formulation of General Relativity [139] in 1915, Albert Einstein came to the conclusion that his description of a dynamic universe was incorrect, and thus introduced into his equations the famous, and somewhat contrived at the time, cosmological constant Λ , with a very specific value intended to counteract the force of gravity and produce

a static and perpetual universe [225]. This was, however, at odds with the work of Vesto Slipher [226, 227] who, by measuring (and discovering!) the Doppler shift of stellar spectra, observed that nearly all nearby galaxies were in fact receding from the Milky Way. Although the implications weren't immediately obvious to Slipher and colleagues, this was some of the first evidence of an expanding universe.

Further theoretical work, performed independently by Alexander Friedman [228] and Georges Lemaître [229], demonstrated that a dynamical and expanding universe can be consistently described by the equations of General Relativity, assuming that on very large scales of order 100 Mpc that the universe is homogeneous and isotropic. In fact, the extrapolation of an expanding universe backward in time led Lemaître to hypothesize the “primeval atom”, which marked the beginning of the well-known Big Bang Theory to describe the origins of the observable universe.

It wasn't until 1929, however, that Edwin Hubble was able to combine measurements of galactic distances, using the period-luminosity relationship of Cepheid variable stars discovered by Henrietta Leavitt [230, 231], with Slipher's redshift measurements in order to unequivocally establish that the universe was indeed expanding. This was made clear by an observed linear relationship between distance and radial velocity of nearby galaxies, the latter quantity being derived from redshift measurements [232]. From this linear relationship, Hubble extracted the Hubble constant $H_0 \sim 500 \text{ km/s Mpc}$, although more accurate measurements of galactic distance brought this number closer to 100 km/s Mpc . All in all, this result would seem to imply that Λ itself is a useful parameterization as an extension to General Relativity, but that the value proposed by Einstein was possibly incorrect.

Since that initial discovery, successively more advanced cosmological observations have indicated that the universe is not only expanding with $H_0 \approx 70 \text{ km/s Mpc}$, but that the rate of expansion is accelerating [147, 165, 233, 234]. Nothing in either the Standard Model or General Relativity can readily explain this observation. Instead, the currently accepted, macroscopic description of the universe assumes that the accelerating expansion is driven by “dark energy”, whose exact nature remains unknown. Essentially, dark energy is conceptualized as the ubiquitous vacuum energy of an undetected field (or fields), whose effect at cosmological scales can be parameterized with the cosmological constant Λ , appearing as a source term in the equations of General Relativity.

A number of potential solutions have been proposed that would give rise to the observed expansion, and “explain” the cosmological constant. Interestingly, many of these solutions

also have observable consequences in the short-distance regime, where techniques using optically levitated microspheres as force sensors may become useful.

F.2 A Friedman Universe

In order to quantify some of these effects, a bit of math must be introduced. The interested reader can find all of the following information worked in much more detail in References [185, 235], and this should serve only as a primer. It is assumed that the reader has a rudimentary knowledge of General Relativity, such as familiarity with concepts like the Einstein field equations, the metric tensor, causality, and the like.

To begin, make the assumption that at very large length scales, the universe is homogeneous and isotropic, allowing a conceptualization wherein the universe is a fluid of matter and energy. This allows one to write a spacetime metric of the form,

$$ds^2 = a(t)^2 ds_3^2 - c^2 dt^2, \quad (\text{F.1})$$

where $a(t)$ is a scale factor that describes how the proper distance between *any* two points in spacetime evolves as the universe expands or contracts, and $a(t = \text{now}) = 1$ by construction. This is the famous Friedman-Lemaître-Robertson-Walker metric assumed in cosmology, and the scale factor is directly related to the redshift, $z = 1 - \lambda_{\text{observed}}/\lambda_{\text{emitted}}$, commonly measured by astronomers (such as Slipher) via the expression $a = (1 + z)^{-1}$. If one inserts this metric into Einstein's field equations (including the cosmological constant Λ), using the stress-energy tensor of a fluid, it is possible to relate the time-evolution of $a(t)$ to density and pressure of the universe, formalized in the Friedman equations:

$$\begin{aligned} H^2 &= \left(\frac{\dot{a}}{a}\right)^2 = \frac{8\pi G}{3}\rho + \frac{\Lambda c^2}{3} - \frac{kc^2}{a^2} \\ \dot{H} + H^2 &= \frac{\ddot{a}}{a} = -\frac{4\pi G}{3}\left(\rho + \frac{3p}{c^2}\right) + \frac{\Lambda c^2}{3}, \end{aligned} \quad (\text{F.2})$$

where H is the Hubble parameter, ρ and p are the volumetric mass density and pressure of the universe, respectively, G is Newton's gravitational constant, c is the speed of light, k is constant describing the large-scale curvature (in a differential geometry sense) of the universe, and the dots represent time derivatives. The first equation can be understood as a measure of the expansion rate since it describes \dot{a} , while the second equation is related to

the acceleration of the expansion \ddot{a} .

Sometimes, the cosmological constant is omitted from the full equations and instead incorporated into the total density and pressure of the universe,

$$\rho \rightarrow \rho - \frac{\Lambda c^2}{8\pi G} \quad \text{and} \quad p \rightarrow p + \frac{\Lambda c^4}{8\pi G}, \quad (\text{F.3})$$

such that the cosmological constant can be thought of as peculiar form of energy, “dark energy”, wherein $p_\Lambda = -\rho_\Lambda c^2$. The total density and pressure of the universe are then sums of the constituent densities and pressures, respectively. Two things should be noted here. Firstly, note that a sufficiently large negative pressure in the Friedman equations, $p_{\text{tot}} > -\rho_{\text{tot}}c^2/3$, will drive accelerated expansion since $(\ddot{a}/a) > 0$, and dark energy has a strictly negative pressure under our construction. Secondly, this is only one particular description of dark energy, and a necessarily simple one. More complex models often invoke a dark energy equation of state: $p_\Lambda = w(z)\rho_\Lambda c^2$, where $w(z)$ is a parameter that may depend on the redshift z , and thus the age of the universe.

Maintaining a constant value of $w(z) = -1$, which is generally in line with the concept of inflation in the early universe, if it is further assumed that the universe is flat, $k = 0$, a reasonably well-supported assumption in light of recent measurements [146, 147, 165], then the first Friedman equation can be used to derive a critical total density (including Λ as in Equation (F.3)) consistent with the observed expansion,

$$\begin{aligned} \rho_{\text{crit}} &= \frac{3H_0^2}{8\pi G} \approx 9.2 \times 10^{-27} \text{ kg/m}^3 \\ &\approx 5-6 p^+/\text{m}^3, \end{aligned} \quad (\text{F.4})$$

where if $\rho_{\text{tot}} > \rho_{\text{crit}}$, the universe is said to be “closed” and will eventually collapse back into itself by gravitational interaction, while if $\rho_{\text{tot}} < \rho_{\text{crit}}$, the universe is said to be “open” and will continue expanding forever (where $H \rightarrow 0$ as $t \rightarrow \infty$ is included in this latter category). The first Friedman equation can be rewritten by defining a density parameter $\Omega = \rho/\rho_{\text{crit}}$, and separating the disparate terms from the constituents. On cosmological scales, the universe can be divided into three categories of mass-energy based on the nature of their contributions to the stress-energy tensor:¹

¹In complete descriptions, there is also a density parameter related to the curvature of the universe, k , but given that current measurements suggest $k = 0$ [146, 147, 165], we will work from this assumption.

1. Massive particles/matter - both baryonic and dark matter
2. Relativistic particles/radiation - photons and neutrinos
3. Dark energy - Λ

Which leads us to a density evolution equation,

$$\left(\frac{H}{H_0}\right)^2 = \frac{\Omega_R}{a^4} + \frac{\Omega_M}{a^3} + \Omega_\Lambda, \quad (\text{F.5})$$

where Ω_R , $\Omega_M = \Omega_b + \Omega_c$, and Ω_Λ are the density parameters for radiation, matter (baryonic and cold dark matter), and dark energy, respectively. The massive term Ω_M includes a single massive neutrino eigenstate with $\Sigma m_\nu = 0.06$ eV, although the full Λ CDM model includes 3 relativistic degrees of freedom corresponding to the 3 species of neutrino [147, 235].

The density evolution equation is derived assuming that matter density scales as $\rho_M \propto a^{-3}$ since a describes the expansion of spatial coordinates, and also assuming that radiation scales as $\rho_R \propto a^{-4}$ with a^{-3} again coming from the expansion of spatial coordinates and another a^{-1} from the change in wavelength since photon energy is given by $E = hc/\lambda$. At late times, $a \geq 1$ and the dark energy/cosmological constant term clearly dominates.

Although this is by no means a complete representation, these equations are the basis for Λ CDM cosmology, the reigning description of our universe at the largest scales.

F.3 Observations in Our Friedman Universe

These equations are all well and good, but one might ask how an experimentalist can get a handle on the quantities involved and test this description of the universe. The usual procedure involves first developing a model for some collection of observables where a number of free parameters inform the structure of this model, then the model is fit to the measured data in order to extract best-fit values of the parameters, potentially subject to constraints from first principles or other related measurements. While more modern observations following the form of Slipher and Hubble allow precise determinations of the local value of H_0 [165, 233], robust extraction of the density parameters requires a different technique, governed by a different model. In particular, the cosmic microwave background (CMB) and baryon acoustic oscillations (BAO) may help ferret out the details.

Briefly, in the wake of the Big Bang and after neutrino de-coupling, the universe was a hot dense plasma of photons, baryons, and dark matter. The plasma is thought to have exhibited density fluctuations that trace their origin to quantum fluctuations during the inflationary period immediately after the Big Bang. Outward photon pressure in the over-dense plasma regions, combined with the gravitational attraction toward the centers of the over-dense regions, resulted in acoustic oscillations of the plasma, propagating radially outward from the centers of the over-dense regions.

As the universe continued expanding, eventually the average energy density fell sufficiently to allow electrons and protons to bind and form neutral hydrogen. The hydrogen atoms were typically in excited electronic states, and their relaxation to the ground state involved the emission of photons. Because this recombination (a historical misnomer given the electrons and protons had never been previously combined) necessarily reduced the number of free electrons, photons, especially those emitted from the de-excitation, streamed freely without significant matter interactions and thus preserved information about the state of the universe at the epoch of recombination.

An immediate consequence of photons de-coupling from the primordial plasma is that the acoustic oscillations of the plasma no longer had a driving source of outward pressure, and thus the pattern of baryon density induced by the oscillations, BAO, was preserved as the universe continued to expand. Simultaneously, free-streaming photons emitted during recombination were subsequently redshifted by the expansion, all the way into the microwave spectrum where they're observed today as the CMB. This is a relatively simple picture that serves our purposes well, but the interested reader can find a much more detailed treatment in some of the first manuscripts on recombination [236, 237], as well as some more contemporary works which include a focus on calculation for the inevitable comparison to measurement [238–241].

Importantly, if one assumes a Friedman universe described by the equations from the previous section, then various properties of the observed CMB, as well as the large-scale structure of matter density encoded by the BAO, are directly related to the values of H_0 , $\Omega_M = \Omega_b + \Omega_c$, Ω_Λ , and other cosmological parameters. This was the impetus for a number of ground-based galaxy surveys and satellite-based microwave telescopes, most recently the Baryon Oscillation Spectroscopic Survey (BOSS, part of the Sloan Digital Sky Survey) [146]

and the Planck observatory [147]. Their most recent data releases gave the following values,

$$\begin{aligned} H_0 &= (67.66 \pm 0.42) \text{ km/s Mpc} \\ \Omega_M &= 0.321 \pm 0.013 \quad \Omega_\Lambda = 0.6889 \pm 0.0056. \end{aligned} \tag{F.6}$$

The energy density of the radiative components, Ω_R , is degenerate with other parameters that inform the structure of the CMB, and thus is not directly extracted from the Planck results. However, it can be derived briefly by considering the two most significant contributors: the CMB photons themselves, as well as the cosmic neutrino background, sometimes called relic neutrinos. The latter population arises from a similar mechanism as that which generates the CMB, where the energy density of the expanding universe lowers sufficiently such that weak-scale interactions between neutrinos and protons/neutrons that were in thermal equilibrium eventually “freeze-out”, and the neutrinos propagate freely [242–244].

If a thermal population of CMB photons is assumed, with temperature $T_\gamma = (2.7255 \pm 0.0006) \text{ K}$ [245], the Stefan-Boltzmann Law for black-body radiation, integrated over the surface of the black-body, can be used to compute the associated energy density as $U_\gamma = 4\sigma T_\gamma^4/c$. Because neutrinos were relativistic prior to the epoch of recombination, their energy density is related to the photon energy density by the expression $U_\nu = N_{\text{eff}}(7/8)(4/11)^{4/3}U_\gamma$ [242, 243] where $N_{\text{eff}} = 3.046$ [147] is the effective number of relativistic particle species (aside from photons). The number is slightly larger than 3 because the neutrino decoupling is more complicated than the simple picture painted here [238, 242]. The results can then be compared to the critical density to determine $\Omega_R = \Omega_\gamma + \Omega_\nu$,

$$\begin{aligned} \Omega_\gamma &= \frac{\rho_\gamma}{\rho_{\text{crit}}} = \frac{U_\gamma}{c^2} \cdot \frac{8\pi G}{3H_0^2} \\ \implies \Omega_\gamma &\approx 5.4 \times 10^{-5} \quad \Omega_\nu \approx 3.7 \times 10^{-5}. \end{aligned} \tag{F.7}$$

Thus, the radiative terms are found to be smaller than the measurement uncertainty on Ω_M , and are often ignored when considering late-time behavior of the universe. It is dark energy and the cosmological constant Λ that play a significant role in the fate of our universe, especially at the current epoch with $a = 1$ (see Equation (F.5)). Most investigations of Λ to date have necessarily been carried out with astronomical observations, but there are intrinsically compelling reasons to investigate Λ at other length-scales.

F.4 Relevant Length Scales

With the Planck and BOSS results in hand, we can begin to consider how laboratory experiments might probe Λ . Since observations currently support that the universe is flat, the mass-energy density of dark energy can be calculated from the value of the critical density as $U = \Omega \rho_{\text{crit}} c^2$, and thus,

$$U_{\Lambda} = \rho_{\Lambda} c^2 = \Omega_{\Lambda} \left(\frac{3H_0^2}{8\pi G} \right) c^2 \approx 5.32 \times 10^{-10} \text{ J/m}^3. \quad (\text{F.8})$$

Here, one must take a small leap of faith. Despite the fact that dark energy does not appear to behave like any other forms of mass-energy (in particular $p_{\Lambda} = -\rho_{\Lambda} c^2$), let's suppose it's another quantum field, call it ϕ , with the usual concepts of quantization and mass. A dark energy quanta with rest energy E_{Λ} would thus have a Compton wavelength $l_{\Lambda} = \hbar c / E_{\Lambda}$.² If it is further assumed that any quanta of this field are homogeneously distributed, a relevant energy scale can be computed from the mass-energy density given by the Planck result:

$$\begin{aligned} E_{\Lambda} &= U_{\Lambda} l_{\Lambda}^3 = U_{\Lambda} \left(\frac{\hbar c}{E_{\Lambda}} \right)^3 \\ \implies E_{\Lambda} &= [U_{\Lambda} (\hbar c)^3]^{1/4} \approx 3.6 \times 10^{-22} \text{ J} \approx 2.2 \text{ meV}. \end{aligned} \quad (\text{F.9})$$

This is sometimes referred to generically as the “dark energy scale”, which has an associated length scale l_{Λ} ,

$$\begin{aligned} l_{\Lambda} &= \frac{\hbar c}{E_{\Lambda}} = \frac{\hbar c}{U_{\Lambda} l_{\Lambda}^3} \\ \implies l_{\Lambda} &= \left(\frac{\hbar c}{U_{\Lambda}} \right)^{1/4} \approx 88 \text{ } \mu\text{m}. \end{aligned} \quad (\text{F.10})$$

Following the arguments presented in References [140, 141], at distances above this length-scale l_{Λ} , the primary effect of the field ϕ is on the vacuum energy density and pressure, i.e. it produces the requisite energy density to explain measurements of Ω_{Λ} from cosmology. By contrast, at distances below l_{Λ} the dynamics of ϕ become important, and the field itself may couple to matter, where the nature of this coupling can vary significantly depending

²Despite the Compton wavelength usually being defined with the non-reduced Planck constant (i.e. h not \hbar), the convention in cosmology appears to make these calculations with \hbar . We follow this convention for the sake of consistency.

on the construction of ϕ . It should be noted that some string theories and modified gravity theories do not necessarily require an entirely new field, as will be discussed shortly, but there are a wealth of proposals that introduce a new field with a Compton wavelength of order l_Λ in order to explain the cosmological constant, and this construction leads to a very natural parameterization.

Since the cosmological constant is fundamentally related to General Relativity and thus gravity, it is instructive to compare the length scale l_Λ to another fundamental scales associated with gravity: the Planck length. Consider a particle of mass m , with Compton wavelength $\lambda_c = \hbar/mc$. The Schwarzschild radius r_s associated with a spherical mass m , i.e. the event horizon of a non-rotating black hole of mass m , is given by $r_s = 2Gm/c^2$. If the two values are assumed to be equal, $\lambda_c = r_s \equiv l_P$, a scale at which quantum effects of gravity should become readily apparent, and use the two expressions to eliminate the mass,

$$m = \frac{l_P c^2}{2G} = \frac{\hbar}{l_P c} \quad \implies \quad l_P \approx \sqrt{\frac{G\hbar}{c^3}} \approx 10^{-35} \text{ m}, \quad (\text{F.11})$$

where we've dropped the factor of 2 for simplicity, as l_P is then just a product of fundamental constants. Comparing this quantity to the dark energy scale: $l_P/l_\Lambda \sim 10^{-31}$, a monstrous difference. The cosmological constant is often interpreted very simply as a vacuum energy of some gravitational field, while the Planck length should represents a characteristic scale for vacuum fluctuations of a quantum gravitational field. Naïvely, one would expect these numbers to be the same, but it is clear that this is far from the case. The mismatch in scales is often called the ‘‘Cosmological Constant Problem’’, which still has no satisfactory solution.

Appendix G

Attractor and Shield Fabrication

A number of nanofabricated devices have been designed and fabricated specifically for use in this collection of work, including the source/attractor masses used to search for both screened interactions associated to dark energy (Chapter 7) and non-Newtonian gravity (Chapter 8), as well as the free-standing electrostatic shield implemented in the latter of the two searches. The fabrication of the the attractor masses was developed initially by the author’s predecessor Alexander Rider, together with Qidong Wang, a collaborator based at The Institute of Microelectronics of the Chinese Academy of Sciences, while significant improvements were driven by the inseparable contributions from the collaboration between the author and A. Rider. Fabrication of the free-standing electrostatic shield makes use of a number of processes that had been developed for the attractors, although the synthesis of the processes into a coherent recipe, as well as the use of a few new techniques, was completed primarily by the author, with the invaluable assistance of Alexander Fieguth.

Nearly every fabrication process detailed here was performed in the nano@Stanford laboratories, while device characterization was performed in the Stanford Nano Shared Facilities (SNSF), both of which are supported by the National Science Foundation as part of the National Nanotechnology Coordinated Infrastructure under Award No. ECCS-1542152. A select few processes were performed in-house, rather than in the dedicated facilities.

Generically speaking, the knowledge and tools made accessible through the nano@Stanford labs are exceptional resources for research and development of niche fabrication tasks. There are a wide variety of instruments, some that make use of “tried and true” techniques, as well as instruments closer to the forefront of nanofabrication as a field. The performance and limitations of the instruments are generally well-known, allowing users (such as the author

and collaborators) to step-in with minimal process development and/or refinement.

As a result, the author's knowledge of the finer details of the fabrication (such as the underlying chemistry, how the process scales with certain parameters, and even some of the underlying physics), is limited. Often times, only an effective understanding is necessary, especially if the desired fabrication doesn't push the boundaries of the field. As an example, when depositing photoresist on a silicon wafer, knowing details of the viscosity of a standard photoresist and its adhesion to various surface types is unimportant, as the staff members of the nano@Stanford labs (as well the manufacturers of the tools themselves) have tuned the instruments such that a user can request a "1.6 μm layer of SPR3612", and the tool will do it. Of course, knowing more about the fundamentals of the process can help when adjustments are needed to better achieve one's goals, but these fundamentals are not usually investigated until they become explicitly necessary in order to surmount a difficulty and make progress.

G.1 Basic Processes

As written and used, the fabrication recipes can have close to 100 individual steps. However, many of these steps are repeated in chunks, which allows for a natural structuring of the recipe into more basic processes. For clarity, number of these processes will be defined and assigned specific names, so that longer recipes can be constructed with some degree of brevity. Individual processes will be given a name, such as "`process()`", their individual steps will be described, and then they will be strung together into complete recipes.

In general, the subtleties of using ANY process/tool involved in nanofabrication are great, and success only comes with experience and good training. The same can be said for ensuring chemical compatibility between containers/vessels/baths, their intended contents, and any tools that may come into contact with the contents, which is left as an exercise for the readers. As such, the contents of this Appendix should not be taken as a complete resource for anyone trying to duplicate and/or adapt these processes, but rather as a rough outline, with some important feature or precautions mentioned. Any proper facility (assuming the nano@Stanford labs are a reasonably representative set of examples) will not only have significant documentation on best practices, but will also require some degree of training before a process can be completed by a new user. Performing processes on cheap dummy wafers is also a standard practice, so that things like etch rates can be characterized based on current conditions of the tool. Thus, it doesn't make sense to include every possible

detail here. The Stanford Nanofabrication Facility (SNF) website [246], is a reasonably comprehensive resource and, where appropriate, includes links to other external sites with pertinent information.

G.1.1 Cleaning

We begin by defining a number of cleaning processes that are used throughout and enumerated in Table G.1. These include a gentle de-ionized (DI) water rinse, an aggressive acid cleaning, and a more general multi-step cleaning, primarily used before the initial processing. The usefulness and applicability of the DI rinse should be self-evident. The aggressive acid clean is primarily for removing photoresist and polymer residue after an etching or deposition step, but will assuredly attack exposed metals and should thus be used only in the appropriate context. The final multi-step process, often called the RCA clean [247], was developed in 1970 as a standard method to prepare wafers. The first step serves to remove organic contaminants and particulates; the second step removes oxides allowing access to the bare silicon underneath; the third and final step removes metallic contaminants and passivates the surface of the wafer. The RCA clean can also be used as an intermediate cleaning step, provided that any deposited metals can withstand minor etching.

The majority of the cleaning procedures are performed in a dedicated wet-bench, specifically for use with corrosive chemicals. Proper personal protective equipment is critical to perform these tasks safely, as well as extreme care and caution, given the significant hazards present. Standard procedures must be strictly adhered to, in order to protect oneself as well as others. If in doubt, don't do it.

Table G.1: Cleaning processes.

Process Tool name	Individual steps involved
<i>continued on next page</i>	

<i>continued from previous page</i>	
spin_rinse_dry() SRD	Basic water cleaning for wafers, usually initial and/or final step. High throughput automated tool <ol style="list-style-type: none"> 1. Continuously rinses wafers with DI water, low-velocity spinning to ensure full coverage 2. Medium-velocity spinning to remove most water 3. High-velocity, long-duration spin for complete drying
piranha_clean() wbflexcorr	Aggressive and comprehensive acid cleaning for wafers. Resist and post-etch/post-deposition cleaning <ol style="list-style-type: none"> 1. Fill quartz bath 9:1 sulfuric acid:H₂O₂ 2. Raise temperature to 120 °C, submerge wafers for 20 min 3. Dump rinse,^a 5 times 4. spin_rinse_dry()
rca_clean() wbflexcorr	RCA Corporation standard cleaning procedure for new wafers <ol style="list-style-type: none"> 1. SC1: quartz bath 5:1:1 H₂O:H₂O₂:NH₄OH 2. Raise temperature to 50 °C, submerge wafers for 20 min 3. Dump rinse, 5 times 4. Fill Teflon bath 50:1 H₂O:HF 5. Submerge wafers for 20 min 6. Dump rinse, 5 times 7. SC2: quartz bath 5:1:1 H₂O:H₂O₂:HCl 8. Raise temperature to 50 °C, submerge wafers for 20 min 9. Dump rinse, 5 times 10. spin_rinse_dry()

^aSpecialized bath that fills completely with a large volume of DI water and then drains the entire bath. This process is usually repeated a number of times.

G.1.2 Lithography

One of the most critical processes in nanofabrication is the use of lithography, which allows a specific two-dimensional pattern to be worked into the surface of a silicon wafer. Essentially the process proceeds as follows: photosensitive (or electron-beam sensitive) material called photoresist (or more generally, resist) is deposited onto the surface of a wafer. With appropriate masking or localization, specific parts of the resist surface can be “exposed”, usually either to ultra-violet (UV) photons or an electron beam. The surface of the now-exposed resist is then “developed”, wherein the portions that have been exposed are removed (using

a positive resist), or the portions that haven't been exposed are removed (using a negative resist).

Once the resist has been developed, the surface directly beneath has been partially exposed in a very specific pattern. Layers of the wafer can then be etched following this pattern of exposed and developed resist, or further layers can be deposited in the pattern set forth by the lithography. Repeated lithography/etching/deposition is the core of nanofabrication, which is really made possible by the lithography itself.

In the fabrication procedures described here, we will make use of photolithography specifically, employing photoresist that is sensitive to UV light. The processes used are detailed in Table G.2. Historically, photolithography has required the use of “masks”, where some desired pattern is first drawn onto a transparent substrate, using opaque materials to define the pattern. A photoresist-coated wafer is then exposed by shining a UV lamp at the wafer, with the mask in between such that it casts a shadow. An example masking material is chrome or iron oxide deposited on a quartz substrate. Maskless photolithography is a relatively new field, made accessible by the development of UV diode lasers which can be focused to small spot sizes in order to achieve direct-write capability.

Table G.2: Processes related to photolithography. All such processes should be undertaken in areas dedicated to photolithography, in order to avoid accidental exposure of the photoresist.

Process Tool name	Individual steps involved
<code>vapor_prime()</code> YES oven	Automatic dehydration and wafer priming for optimal photoresist adhesion <ol style="list-style-type: none"> 1. <code>spin_rinse_dry()</code> 2. Dehydration in N₂ at 150 °C 3. Alternate N₂ and vacuum a few times 4. Expose heated wafers to HMDS^a vapor 5. Remove HMDS vapor, alternate N₂ and vacuum
<code>resist_coat(<i>resist_type</i>, <i>thickness</i>)</code> svgcoat	Photoresist spinning, high throughput automated tool <ol style="list-style-type: none"> 1. <code>vapor_prime()</code> 2. Place sender (with wafers to be coated) and receiver cassettes in designated locations 3. Select preset recipe, ensuring that it includes EBR^b 4. Initiate processing and carefully monitor throughout^c 5. Includes medium- and high-velocity spinning steps for uniform resist distribution
<i>continued on next page</i>	

^aHexamethyldisilazane - a adhesion promoter.

^bEdge bead removal - cleans photoresist from a ring around the edge of the wafer so that downstream tools can handle/move the wafer about without touching the “sticky” resist.

^cFor recipes using viscous photoresist for thick layers (such as SPR220-7), extra care must be taken so that near the end of the spinning cycle, the photoresist nozzle doesn't drip on the coated wafer.

<i>continued from previous page</i>	
<pre>resist_expose(mask_name, wavelength, intensity, duration, cycles, delay) karlsuss</pre>	<p>Contact mask alignment^a for photoresist exposure. Uses prefabricated chrome-on-quartz hard masks</p> <ol style="list-style-type: none"> 1. Clean masks, first with solvent rinses (acetone, then isopropanol, then methanol). 2. Dry masks with N₂ by “pushing” methanol droplets to the bottom edge (to be absorbed by cleanroom wipe) 3. Finish with dedicated mask cleaner, using DI water and high-velocity rotation 4. Program tool with desired exposure parameters: wavelength/intensity limited control; duration/cycles/delay complete control 5. Install mask in tool, ensuring operation of vacuum chuck 6. Locate alignment marks on mask 7. Install resist-coated wafer with appropriate chuck 8. Co-align wafer marks with mask marks 9. Bring mask/wafer into contact, expose unmasked region 10. Post-exposure bake if necessary
<i>continued on next page</i>	

^aDirect-write, maskless exposure technologies exist and would likely be optimal for the scale of features we are fabricating and the expected wafer throughput (<10), but the author was not trained on these tools, and the capabilities of the “old-school” contact aligner were more than sufficient.

<i>continued from previous page</i>	
resist_develop(<i>resist_type</i>, <i>thickness</i>) svgdev	Photoresist development, high throughput automated tool <ol style="list-style-type: none"> 1. Place sender (with exposed wafers to be developed) and receiver cassettes in designated locations 2. Select preset recipe, dependent on resist type/thickness 3. Initiate processing and carefully monitor throughout 4. Wafers are lightly baked (recipe dependent), covered with developer solution for a set duration, then rinsed to complete the process 5. Assess development with optical microscope 6. spin_rinse_dry()

G.1.3 Deposition

Lithography is usually a precursor step to either an etch or a deposition, as the process imprints a very specific pattern onto the resist, allowing selective and directed modifications to the wafer surface beneath the resist. Once a pattern has been drawn, material can be deposited in the form of that pattern in order to build a device, such as a metal deposition forming a surface electrode or electrically connecting features to one another on the wafer. Although the list of materials that one can deposit is expansive, very common depositions include oxides (such as SiO_2 , or Al_2O_3 and other metal oxides), metals, and nitrides. The fabrication processes detailed here are restricted to SiO_2 and pure metals, with two distinct deposition techniques for each type of material.

For SiO_2 , thermal oxides and plasma-enhanced chemical vapor deposition (PECVD) are employed. In the former process, a wafer is subjected to temperatures of $\sim 1000^\circ\text{C}$ in the presence of molecular oxygen or water vapor for an extended period of time, which oxidizes the outermost layer of the wafer, on all sides. For the latter process, a mixture of gases is ignited into a plasma, and the constituent charged particles are electrostatically accelerated toward the surface of a wafer where they adsorb and react to deposit a material. PECVD is a versatile technique that can also grow nitrides and carbides with the appropriate reactant gases.

Metal depositions are historically performed with either evaporation or sputtering, the former being very directional and the latter intentionally less so. An evaporation chamber usually contains some sort of ablation mechanism wherein, for example, a reservoir of the

desired metal is evaporated by an electron beam. If the wafers are positioned above the metal target, with their device side facing down, the evaporated metal atoms will impact the surface and adhere, resulting in a uniform layer, with directional deposition a result of the usually long (~ 0.5 m) distance between the electron-beam source, target, and wafers to be plated. Sputtering involves generating a localized plasma above the surface of a desired deposition material, where accelerated particles within the plasma kick off some of the target material, which is then “swirled” around a vacuum chamber with a magnetic field. By rotating the wafer(s) to be coated in the presence of this swirling plasma, nominally uniform coverage is achieved on all surfaces, unless there are some high aspect ratio ($\gtrsim 5$) features.

Table G.3: Processes related to deposition.

Process Tool name	Individual steps involved
thermal_oxide(<i>thickness</i>) thermco4	Oxidize the outermost layer of a silicon wafer to produce a uniform thermal oxide layer. <ol style="list-style-type: none"> 1. piranha_clean()^a 2. rca_clean() 3. Program furnace recipe for desired oxide thickness^b 4. Once furnace boat is fully extended, load wafers with quartz vacuum wand 5. Initiate thermal processing 6. When completed, assess oxide thickness
<i>continued on next page</i>	

^aThese cleaning steps are recommended generically to prepare surfaces for a subsequent deposition. However, if the wafer has photoresist or other sensitive layers, these cleaning steps may not be appropriate. Process compatibility should be ensured before executing initial cleaning steps.

^bRecipe depends on wafer crystal orientation, wet vs. dry environment, furnace temperature, thickness of pre-existing oxide layer [248].

<i>continued from previous page</i>	
<p>pecvd_oxide(<i>thickness</i>)</p> <p>ccp-dep</p>	<p>Plasma-enhanced chemical vapor deposition of SiO_x, with parameters such that $x \rightarrow 2$</p> <ol style="list-style-type: none"> 1. piranha_clean() 2. rca_clean() 3. Bake at 110 °C for 10 min 4. Load chamber without any wafers 5. Run SF₆ and N₂O plasma chamber clean 10 min/μm of previous recipe 6. Load wafers (max of 4) with deposition side up, including dummies to fill platen 7. Run recipe "SIO350-0" 8. Remove wafers and rerun cleaning recipe
<p>metal_evap(<i>species,</i> <i>thickness</i>)</p> <p>innotech</p>	<p>Electron-beam evaporation of metals or metal-oxides, directional deposition</p> <ol style="list-style-type: none"> 1. piranha_clean() 2. rca_clean() 3. Remove blanks in wafer platen, insert wafers with deposition side <i>down</i> 4. With chamber evacuated, initiate platen rotation 5. Turn on electron beam and steer to target crucible location (with metal) 6. Wait until thickness readout reaches desired level 7. Turn off electron-beam and remove coated wafers
<i>continued on next page</i>	

<i>continued from previous page</i>	
<p><code>metal_sputter(<i>species</i>, <i>thickness</i>)</code></p> <p>lesker-sputter</p>	<p>Magnetron sputtering tool for metal or metal-oxide non-directional deposition</p> <ol style="list-style-type: none"> 1. <code>piranha_clean()</code> 2. <code>rca_clean()</code> 3. Load wafer onto standard chuck (or load custom device mount) in load-lock, bring to process chamber when vacuum pressure conditions are met 4. Configure recipe parameters for desired thickness, i.e. RF power and plasma pressure 5. Start recipe, using viewport to confirm plasma activity and platen rotation. 6. Remove wafers once automated recipe is complete
<p><code>metal_plating(<i>species</i>, <i>thickness</i>)</code></p> <p>Custom tool (not SNF)</p>	<p>Electroplating to build structures from evaporated seed layers</p> <ol style="list-style-type: none"> 1. <code>spin_rinse_dry()</code> 2. Package appropriately for removing from cleanroom environment 3. Install wafer into plating apparatus, fill with filtered plating solution 4. Use new gold leads (no buildup) to make electrical contact with wafer 5. Initiate stir bar and pulsed plating supply^a to drive electrochemical cell 6. Assess growth rate with dummy wafer first. Approximately 150 nm/min with these parameters

^aThis process was not optimized significantly, but reasonable parameters are 45 mA forward current, 1 kHz and 50 % duty cycle square wave drive.

G.1.4 Etching

Etching is the yin to deposition's yang. Exactly the same principle holds wherein lithography defines a specific pattern on a device wafer, which can then be used to etch that pattern into the surface. Sometimes, a feature is etched into the surface of a wafer, and subsequently followed by a deposition to fill that etched feature with a desired material. Etching can either be directed, or isotropic, where directionality is achieved by specifying the direction of the etchant (i.e. with an electric field), or by taking advantage of the orientation of a

crystal structure, such as in KOH based etching of silicon wafers.

Nearly all of the etching performed as a part of these fabrication processes is a form of reactive ion etching (RIE). With this technique, a plasma is formed in a specific combination of gases, and the constituent ions within the plasma are directed toward the surface to be etched by an electric field, yielding marginal directionality in the over etch profile. The ions adsorb into the surface and subsequently react to form a volatile product, which then leaves surface exposing more of the material to be etched. The process of accelerating ions toward a surface results in some non-reactive ablation, such that all materials are etched slightly within such a plasma. Different combinations of gases etch distinct materials differentially, both via reaction and ablation, allowing selectivity to be tuned to some extent.

One may notice when perusing Table G.4 that there are two oxide RIE etches, corresponding to two distinct tools with different gas mixtures and plasma configurations. Incidentally, both systems yield approximately the same etch rate of SiO_2 , 200 nm/min, as well as the same selectivity for SiO_2 relative to bare silicon or photoresist, $\gtrsim 2$. Different electric field and plasma configurations have benefits and drawbacks which are beyond the scope of this work. In practice, both tools have been used for the same oxide etching steps, where the choice of tool is dictated by whichever is less busy, given that the tools are shared among a large user base.

A distinct and all-important etching process here is anisotropic deep reactive ion etching (DRIE), specifically deep silicon etching. This process is time-multiplexed with two etching modes that result in vertical structures, often with high aspect ratio. The first step is a standard RIE etch, much like those above, wherein ions from a plasma are directed toward a surface. A second step deposits a passivation layer, commonly a fluorocarbon like Teflon. Repeating the first step with the directional etching inherent to RIE preferentially attacks the bottom of the etched feature and not the passivated sidewalls, first removing the passivation on the bottom, and subsequently etching the material underneath. Successive repetitions etch the patterned feature straight down.

Two isotropic etching techniques are used, a hydrofluoric acid vapor etch to remove sacrificial oxide layers, and a hydrofluoric acid wet etch remove metals, specifically Ti. Both are relatively straightforward: the wafer to be etched is surrounded by the acid vapor or submersed in an aqueous bath for a set period of time, then removed and rinsed/cleaned.

Table G.4: Processes related to etching. Many of the tools used have automated wafer handling and pre-programmed recipes. *Always* start from a pre-programmed recipe known to work for other users, and then tweak as necessary based on dummy wafer performance.

Process Tool name	Individual steps involved
<p>descum()</p> <p>drytek2, oxford-rie, technics</p>	<p>Basic O₂ plasma etch to remove any contaminants that may have found their way to the wafer surface.</p> <ol style="list-style-type: none"> 1. RF power and O₂ pressure depend on the specific chamber, which should be known and/or characterized prior to use 2. For drytek2/oxford-rie: 500 W, 150 mtorr 3. For technics: 50 W, 300 mtorr 4. Clean chamber with O₂ plasma etch, but no wafers, 5 to 10 min 5. Load wafer(s) in requisite position within plasma chamber 6. Initiate short duration etch, 1 to 2 min
<i>continued on next page</i>	

<i>continued from previous page</i>	
<p>oxide_rie_ccp(<i>depth</i>)</p> <p>oxford-rie</p>	<p>Plasma etch with CHF₃ and CF₄ in an Argon carrier gas, performed in a capacitively-coupled plasma chamber. Selective to SiO₂ and other oxides/nitrides.</p> <ol style="list-style-type: none"> 1. descum() 2. Clean chamber with O₂ plasma etch, blank silicon wafer, 10 to 30 min with standard cleaning recipe (similar to descum) 3. Etch dummy wafer to condition chamber before etching device wafer 4. Gas flow rates for standard oxide etch: 30 sccm Ar, 45 sccm CHF₃, and 15 sccm CF₄^a 5. Overall plasma parameters: 500 W, 100 mtorr 6. Backside He cooling at 10 torr with 20 °C setpoint 7. Max etch time ~20 min before another chamber cleaning is required
<i>continued on next page</i>	

^asccm - standard cubic centimeter per minute

<i>continued from previous page</i>	
<p><code>oxide_rie_icp(<i>depth</i>)</code></p> <p>pt-ox</p>	<p>Plasma etch with CHF₃ and O₂, performed in an inductively-couple plasma chamber. Selective to SiO₂ and other oxides/nitrides</p> <ol style="list-style-type: none"> 1. <code>descum()</code> 2. Clean chamber with O₂ plasma etch, blank silicon wafer, 10 to 30 min with standard cleaning recipe (similar to <code>descum</code>) 3. Gas flow rates for standard oxide etch: 2 sccm O₂ and 45 sccm CHF₃ 4. Overall plasma parameters: 600 W for main plasma, 50 W bias point for wafer, 5 mtorr 5. Backside He cooling at 5 torr with 10 °C setpoint 6. Max etch time ~10 min before another chamber cleaning is required
<p><code>oxide_hf_vapor(<i>endpoint</i>)</code></p> <p>uetch</p>	<p>Isotropic vapor etch using hydrofluoric acid and ethanol. Extremely high selectivity for oxide</p> <ol style="list-style-type: none"> 1. <code>piranha_clean()</code> 2. <code>rca_clean()</code> 3. Bake wafer at 250 °C or <code>descum()</code> 4. Select one pre-programmed recipe R1 through R5, increasing etch rate from 25 nm/min to 150 nm/min at the cost of uniformity 5. Tune etch time and cycle number, nothing else 6. Run chamber through 10, 60 s cycles with no wafer 7. Insert wafer and run through desired cycles 8. Assess completion and repeat if necessary
<i>continued on next page</i>	

<i>continued from previous page</i>	
<p><code>silicon_drie(<i>depth</i>)</code></p> <p>pt-dse</p>	<p>Time-multiplexed, anisotropic, deep silicon etching with the Bosch process [92]</p> <ol style="list-style-type: none"> 1. <code>descum()</code> 2. Run DSE-Clean, O₂ plasma recipe for 30 min 3. Load either aggressive recipe for ultra-deep backside etches (DSE-FAT, ~1 μm/cycle), or sensitive recipe for fine topside etch (DSE-Nano, ~100 nm/cycle) 4. Condition chamber with dummy wafer and 20 cycles of main etch 5. Etch device wafer, no more than 150 cycles 6. Clean chamber for 10 min 7. Repeat etch step and cleaning until completed
<p><code>metal_wet_etch(<i>depth</i>)</code></p> <p>wbflexcorr</p>	<p>Hydrofluoric acid wet etch to remove thin metal layers</p> <ol style="list-style-type: none"> 1. <code>spin_rinse_dry()</code> 2. In Teflon containers, prepare 20:1:1 mixture of H₂O:H₂O₂:HF(49%) with an etch rate of 1000 nm/min for Ti, together with two stages of H₂O rinsing 3. Immerse wafer in etch solution and gently stir 4. Rinse in first station, rinse in second station 5. <code>spin_rinse_dry()</code>

G.1.5 Other processes

A few other processes don't fit nicely into the above categories and are thus included here.

Table G.5: Miscellaneous processes.

Process Tool name	Individual steps involved
cmp(<i>endpoint</i>) cmp	Chemical mechanical polishing for surface planarization <ol style="list-style-type: none"> 1. Soak polishing pad with DI water for 30 minutes, scrub clean 2. For Si polishing, use Ultra-Sol S-10 slurry and prepare with aggressive mechanical agitation 3. Flush slurry and water lines, active flow rate of ~few drops per second 4. Pad conditioning run with slurry active and recipe "Auto 2" 5. Dummy wafer conditioning run with recipe "Auto 3" 6. Load device wafer into chuck and initiate polishing 7. Release wafer from chuck at regular intervals and assess polishing with optical microscope 8. Typical settings: platen, 100 rpm; polishing head, 60 rpm; polishing pressure, 250 g/cm²; retainer pressure, 350 g/cm²; polishing time, 100 s 9. DO NOT LET SLURRY DRY ON ANYTHING. LITERALLY IMPOSSIBLE TO REMOVE.

G.2 Density-Staggered Attractor Recipe

For the scientific program investigating modifications to gravity, we choose to excite an oscillating gravitational response of a trapped microsphere using a source mass, deemed the attractor, which has a regularly patterned structure of materials with significantly different densities. The density modulation is microscopic in scale with periodicity $50\ \mu\text{m}$, making it ideal for short-distance modifications to Newtonian gravity. Given the availability of tools and resources to fabricate mechanical devices of this size, nanofabrication techniques were used to define the low-density portion of the attractor in silicon ($\rho_{\text{Si}} = 2.33\ \text{g}/\text{cm}^3$), filling the voids within the silicon outline with gold ($\rho_{\text{Au}} = 19.3\ \text{g}/\text{cm}^3$) to produce the high-density portion of the attractor.

Different metals or substrate materials could also be used, with the expected signal scaling roughly as $\sim\Delta\rho$. However, given the range of densities of materials known to humankind, one can only gain factors of order unity by optimizing over material parameter space. Fabrication feasibility, as well as possible backgrounds introduced by materials with undesirable properties, often plays a more significant role in the design of an attractor mass.

Below is a fabrication recipe using the previously defined steps. This particular recipe has been adapted slightly from the original detailed in Reference [69], although the end result should be nearly identical. Step numbering matches that shown in Figure G.1.

(-) Initial cleaning

- piranha_clean()
- rca_clean()

(1) Photolithography for trenches

- resist_coat(SCR3612, $1.6\ \mu\text{m}$)
- resist_expose(trenches, 365 nm, $15\ \text{mW}/\text{cm}^2$, 1.5 s, 1, 0)
- resist_develop(SCR3612, $1.6\ \mu\text{m}$)

(2) Trench etch

- silicon_drie($10\ \mu\text{m}$), fully through device layer
- piranha_clean()

- (3) `thermal_oxide(250 nm)`
- (4) Deposit electroplating seed layer
 - `metal_evap(Ti, 20 nm)`, adhesion layer
 - `metal_evap(Au, 50 nm)`
- (5) `cmp(remove Au/Ti on topside)`
- (6) Grow gold in trenches
 - `metal_plating(Au, 10 μm)`, filling trenches
 - `cmp(planarize grown gold with wafer surface)`
- (7) Photolithography for device outline
 - `resist_coat(SCR3612, 1.6 μm)`
 - `resist_expose(device outline, 365 nm, 15 mW/cm2, 1.5 s, 1, 0)`
 - `resist_develop(SCR3612, 1.6 μm)`
- (8) Device outline etch
 - `oxide_rie_ccp(250 nm)`, through oxide polishing mask
 - `silicon_drie(10 μm)`, again fully through device layer
 - `piranha_clean()`
- (9) PECVD oxide hardmask for backside release etch
 - `descum()`
 - `pecvd_oxide(4 μm)`
 - `piranha_clean()`
- (10) Photolithography for device release
 - `resist_coat(SCR220-7, 7 μm)`
 - `resist_expose(device outline, 365 nm, 15 mW/cm2, 2.5 s, 4, 120 s)`
 - `resist_develop(SCR220-7, 7 μm)`

(11) Release etch

- `oxide_rie_ccp(4 μm)`, through pecvd oxide hard mask
- `silicon_drie(~500 μm)`, fully through handle layer
- `piranha_clean()`

(12) `oxide_hf_vapor(remove buried/sacrificial oxide)`

(-) `metal_sputter(Ti, 50 nm)`

(-) `metal_sputter(Au, 150 nm)`

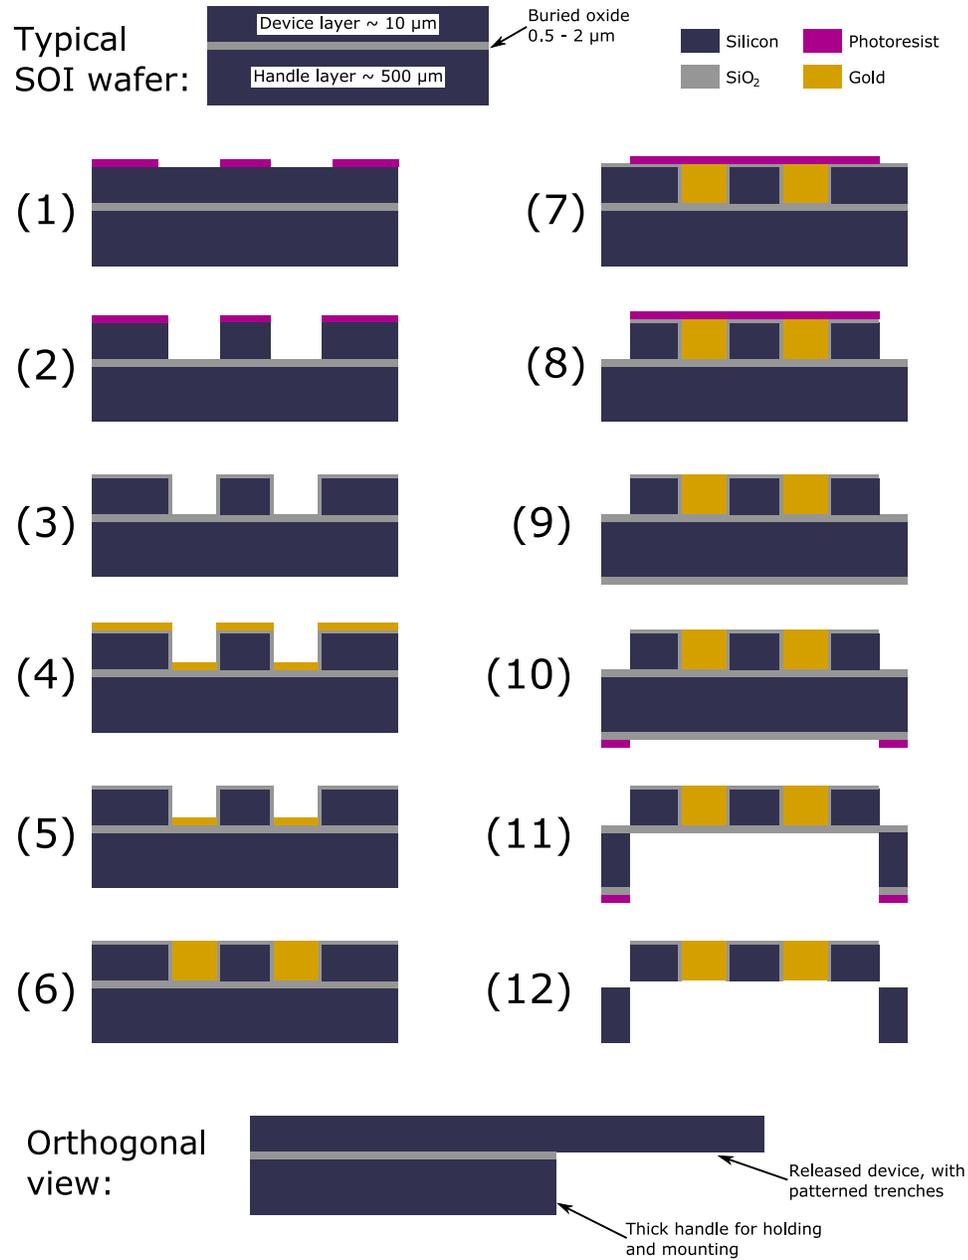


Figure G.1: Selected steps from the attractor fabrication process, with numbers corresponding to the enumeration in the text. (1) trench lithography, (2) trench etch, (3) grow thermal oxide, (4) evaporate gold, (5) CMP topside, (6) electroform in trench with CMP after, (7) outline lithography, (8) outline etch, (9) backside PECVD oxide, (10) release lithography, (11) deep release etch, (12) HF vapor etch. Devices are then snapped free of the carrier wafer and sputter-coated. *Significantly* not-to-scale.

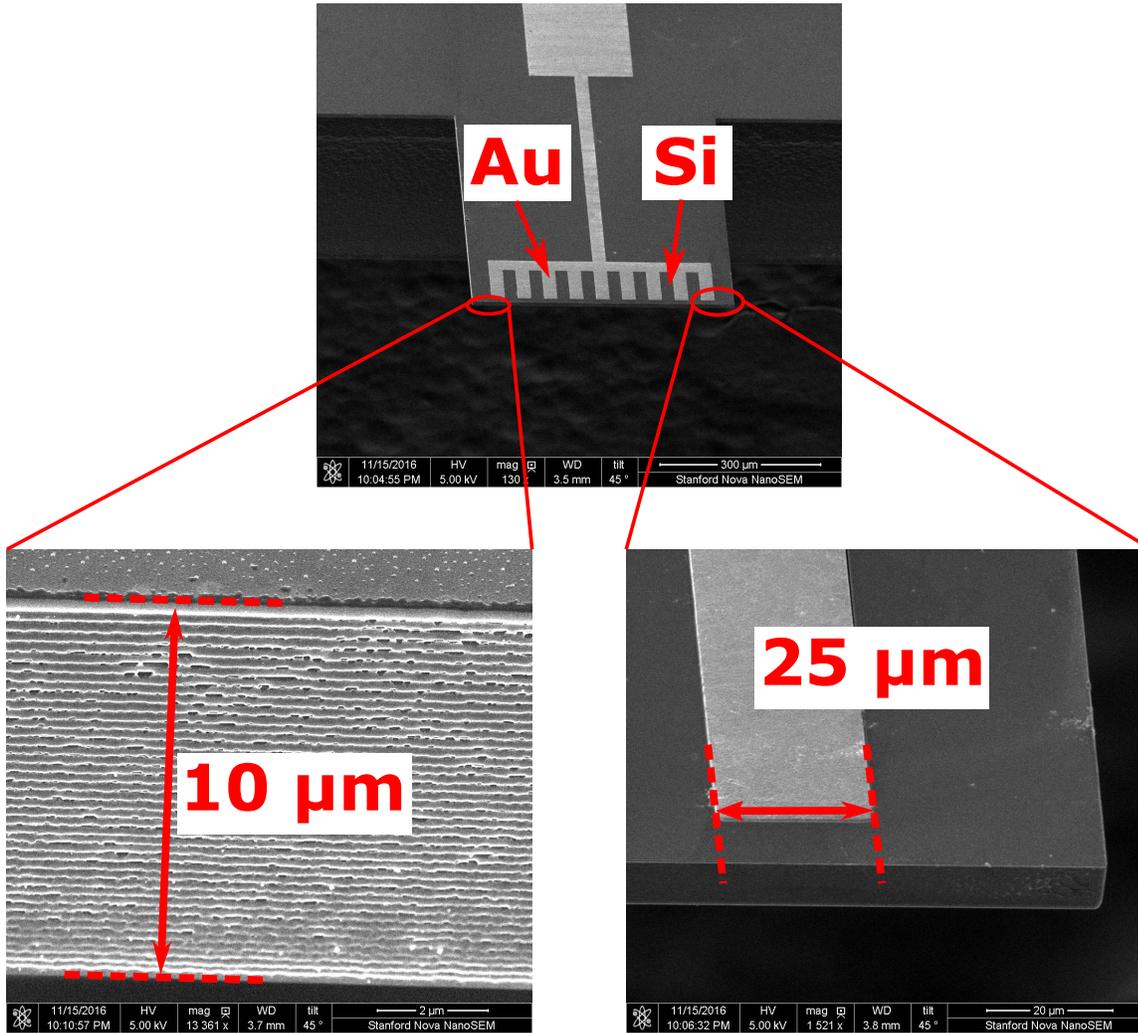


Figure G.2: Annotated SEM micrographs of a completed attractor device, prior to deposition of the final metallic coating. Two details show the front face of the attractor with the characteristic scalloping from the deep silicon etching with the Bosch process [92], as well as one of the electroformed and polished gold portions of the density modulation. A few length scales are indicated.

G.3 Free-standing Electrostatic Shield Recipe

Any application of force-sensing must contend with the possibility of background forces that may contaminate the measurement. The strong and weak forces are limited to sub-atomic distance scales and thus do not factor into consideration here, but electromagnetism has an infinite range and thus will always present some level of background that must be mitigated. This is especially difficult given that gravity, and gravity-like forces, are weak compared to electromagnetism.

Electrically neutralizing the test and source masses removes the monopole component of the associated charge distributions, which is a possible source of significant background forces. However, permanent higher order moments, dipole/quadrupole/etc., still remain in silica microspheres and appear to be highly material and manufacturer dependent. Polarizability also complicates this description further, and cannot be assumed to be isotropic, nor uniform. Clearly a naïve solution is to electrostatically shield the test and source masses from one another. This has been attempted with a free-standing conductive device, positioned stationary between test and source masses. The shield is designed with an L-shaped cross section to simultaneously block retroreflections from the attractor itself, which would otherwise affect the z position measurement.

Below is a fabrication recipe using the previously defined steps. This particular recipe is unpublished was developed primarily by the author. Step numbering matches that shown in Figure G.3.

- (-) Initial cleaning
 - `piranha_clean()`
 - `rca_clean()`
- (1) `thermal_oxide(250 nm)`
- (2) Photolithography for “pocket feature”
 - `resist_coat(SCR3612, 1.6 μm)`
 - `resist_expose(trenches, 365 nm, 15 mW/cm2, 1.5 s, 1, 0)`
 - `resist_develop(SCR3612, 1.6 μm)`
- (3) Etch pocket into oxide hard mask only (to register it’s position)

- oxide_rie_ccp(250 nm), stop on silicon
- (4) Strip resist then photolithography for channel and device outline
- piranha_clean()
 - resist_coat(SCR3612, 1.6 μm)
 - resist_expose(trenches, 365 nm, 15 mW/cm², 1.5 s, 1, 0)
 - resist_develop(SCR3612, 1.6 μm)
- (5) Channel and device outline *partial* etch
- oxide_rie_ccp(250 nm), through oxide hard mask
 - silicon_drie(3 μm), *partially* through device layer¹
 - piranha_clean()
- (6) Strip resist completely: piranha_clean()
- (7) Simultaneously finish device channel/outline etch and pocket etch
- silicon_drie(22 μm), *now* fully through device layer
 - piranha_clean()
- (8) PECVD oxide hardmask for backside release etch
- descum()
 - pecvd_oxide(4 μm)
 - piranha_clean()
- (9) Photolithography for device release
- resist_coat(SCR220-7, 7 μm)
 - resist_expose(device outline, 365 nm, 15 mW/cm², 2.5 s, 4, 120 s)
 - resist_develop(SCR220-7, 7 μm)
- (10) Release etch
- oxide_rie_ccp(4 μm), through pecvd oxide hard mask

¹This etch depth roughly defines how thick the bottom portion of the L-shaped cross-section is.

- silicon_drie(~500 μm), fully through handle layer
- piranha_clean()

(11) oxide_hf_vapor(remove buried/sacrificial oxide)

(-) metal_sputter(Ti, 50 nm)

(-) metal_sputter(Au, 150 nm)

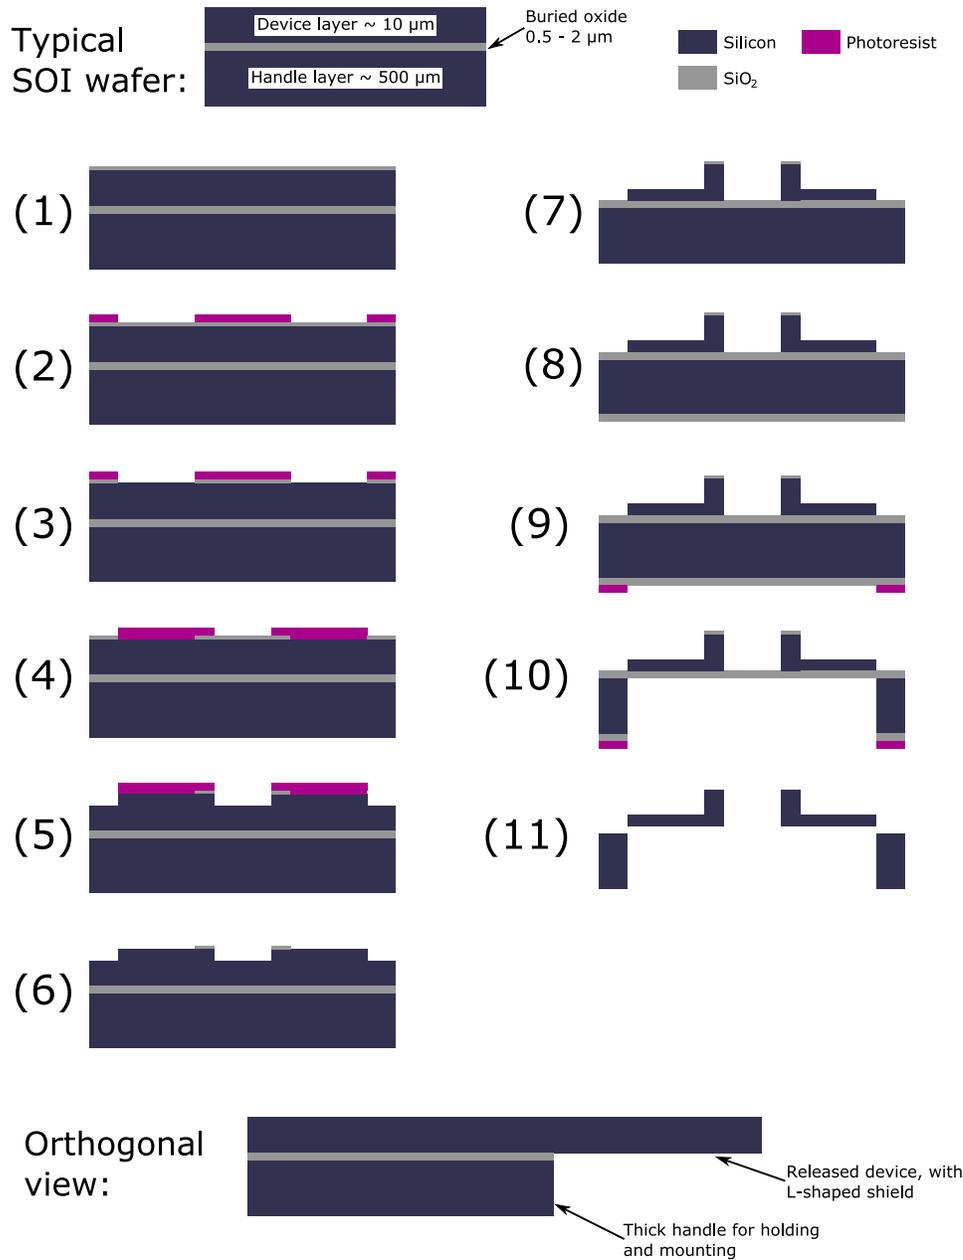


Figure G.3: Selected steps from the free-standing electrostatic shield fabrication process, with numbers to match the text. (1) Grow thermal oxide, (2) “pocket” lithography, (3) etch pocket in oxide mask, (4) strip resist then outline lithography, (5) etch oxide hard mask and partial device layer, (6) strip resist, (7) etch pocket and outline simultaneously to buried oxide stop, (8) backside PECVD oxide, (9) release lithography, (10) deep release etch, (11) HF vapor etch. Devices are snapped free of the carrier wafer and sputter-coated. *Significantly* not-to-scale.

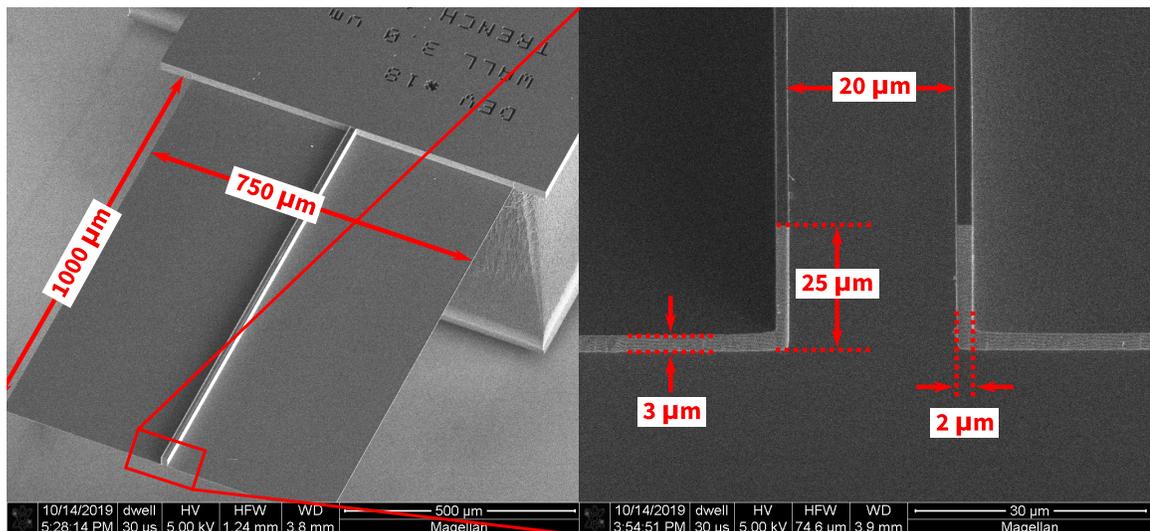
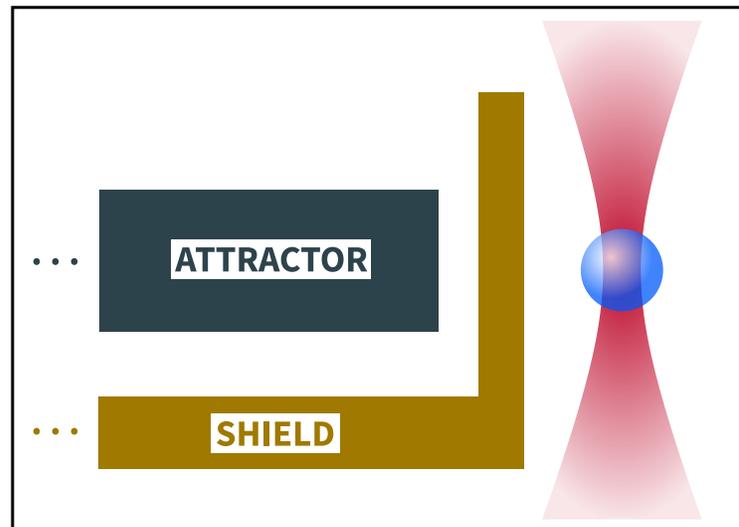


Figure G.4: (upper) A schematic depiction of the attractor and shield positioning near a trapped microsphere, emphasizing the L-shaped cross-section. (lower) Annotated SEM micrographs of a completed shield device, prior to deposition of the final metallic coating, with both images using a 40° inclination in viewing angle to make the three-dimensional structure clearly visible. A detail shows the end of the shield and the symmetric pockets for this particular device. A few length scales are indicated.

Bibliography

- [1] A. Ashkin, “Acceleration and trapping of particles by radiation pressure”, [Physical Review Letters](#) **24**, 156–159 (1970).
- [2] J. C. Maxwell, “Viii. a dynamical theory of the electromagnetic field”, [Philosophical Transactions of The Royal Society](#) **155**, 459–512 (1865).
- [3] O. Heaviside, *Electromagnetic Theory*, 3 vols. (The Electrician Printing and Publishing Co, London, 1893).
- [4] J. D. Jackson, *Classical Electrodynamics* (John Wiley & Sons, New York, 1962).
- [5] E. Hecht and A. Zajac, *Optics* (Addison Wesley, Boston, 1974).
- [6] A. Ashkin, J. M. Dziedzic, J. E. Bjorkholm, and S. Chu, “Observation of a single-beam gradient force optical trap for dielectric particles”, [Optics Letters](#) **11**, 288–290 (1986).
- [7] A. Ashkin and J. M. Dziedzic, “Optical levitation by radiation pressure”, [Applied Physics Letters](#) **19**, 283–285 (1971).
- [8] A. Ashkin and J. M. Dziedzic, “Optical levitation in high vacuum”, [Applied Physics Letters](#) **28**, 333–335 (1976).
- [9] A. Ashkin and J. M. Dziedzic, “Feedback stabilization of optically levitated particles”, [Applied Physics Letters](#) **30**, 202, 202 (1977).
- [10] G. Mie, “Beiträge zur optik trüber medien, speziell kolloidaler metallösungen”, [Annalen der Physik](#) **330**, 377–445 (1908).
- [11] P. C. Waterman, “Symmetry, unitarity, and geometry in electromagnetic scattering”, [Physical Review D](#) **3**, 825–839 (1971).

- [12] B. Maheu, G. Gouesbet, and G. Grehan, “A concise presentation of the generalized lorenz-mie theory for arbitrary location of the scatterer in an arbitrary incident profile”, *Journal of Optics* **19**, 59–67 (1988).
- [13] T. A. Nieminen, V. L. Y. Loke, A. B. Stilgoe, G. Knöner, A. M. Brańczyk, N. R. Heckenberg, and H. Rubinsztein-Dunlop, “Optical tweezers computational toolbox”, *Journal of Optics A: Pure and Applied Optics* **9**, S196 (2007).
- [14] T. Nieminen, H. Rubinsztein-Dunlop, and N. Heckenberg, “Calculation of the t-matrix: general considerations and application of the point-matching method”, *Journal of Quantitative Spectroscopy and Radiative Transfer* **79-80**, *Electromagnetic and Light Scattering by Non-Spherical Particles*, 1019–1029 (2003).
- [15] T. Nieminen, H. Rubinsztein-Dunlop, and N. Heckenberg, “Multipole expansion of strongly focussed laser beams”, *Journal of Quantitative Spectroscopy and Radiative Transfer* **79-80**, *Electromagnetic and Light Scattering by Non-Spherical Particles*, 1005–1017 (2003).
- [16] W. Stöber, A. Fink, and E. Bohn, “Controlled growth of monodisperse silica spheres in the micron size range”, *Journal of Colloid and Interface Science* **26**, 62–69 (1968).
- [17] I. Ab Rahman and P. Vejayakumaran, “Synthesis of silica nanoparticles by sol-gel: size-dependent properties, surface modification, and applications in silica-polymer nanocomposites—a review”, *Journal of Nanomaterials* **2012**, 132424 (2012).
- [18] S. R. Parnell, A. L. Washington, A. J. Parnell, A. Walsh, R. M. Dalgliesh, F. Li, W. A. Hamilton, S. Prevost, J. P. A. Fairclough, and R. Pynn, “Porosity of silica stöber particles determined by spin-echo small angle neutron scattering”, *Soft Matter* **12**, 4709–4714 (2016).
- [19] C. P. Blakemore, A. D. Rider, S. Roy, A. Fieguth, A. Kawasaki, N. Priel, and G. Gratta, “Precision mass and density measurement of individual optically levitated microspheres”, *Physical Review Applied* **12**, 024037 (2019).
- [20] A. E. Siegman, *Lasers* (University Science Books, 1986).
- [21] T. Li, S. Kheifets, and M. G. Raizen, “Millikelvin cooling of an optically trapped microsphere in vacuum”, *Nature Physics* **7**, 527–530 (2011).
- [22] L. D. Hinkle and B. R. F. Kendall, “Brownian motion of a particle levitated in vacuum”, *Journal of Vacuum Science & Technology A* **10**, 243–247 (1992).

- [23] K. Berg-Sørensen and H. Flyvbjerg, “Power spectrum analysis for optical tweezers”, [Review of Scientific Instruments](#) **75**, 594–612 (2004).
- [24] D. C. Moore, A. D. Rider, and G. Gratta, “Search for millicharged particles using optically levitated microspheres”, [Physical Review Letters](#) **113**, 251801 (2014).
- [25] A. D. Rider, D. C. Moore, C. P. Blakemore, M. Louis, M. Lu, and G. Gratta, “Search for screened interactions associated with dark energy below the 100 μm length scale”, [Physical Review Letters](#) **117**, 101101 (2016).
- [26] A. D. Rider, C. P. Blakemore, G. Gratta, and D. C. Moore, “Single-beam dielectric-microsphere trapping with optical heterodyne detection”, [Physical Review A](#) **97**, 013842 (2018).
- [27] C. P. Blakemore, A. D. Rider, S. Roy, Q. Wang, A. Kawasaki, and G. Gratta, “Three-dimensional force-field microscopy with optically levitated microspheres”, [Physical Review A](#) **99**, 023816 (2019).
- [28] C. P. Blakemore, D. Martin, A. Fieguth, A. Kawasaki, N. Priel, A. D. Rider, and G. Gratta, “Absolute pressure and gas species identification with an optically levitated rotor”, [Journal of Vacuum Science & Technology B](#) **38**, 024201 (2020).
- [29] C. P. Blakemore, A. Fieguth, A. Kawasaki, N. Priel, D. Martin, A. D. Rider, Q. Wang, and G. Gratta, “Search for non-Newtonian interactions at micrometer scale with a levitated test mass”, (2021).
- [30] Y. Roichman, B. Sun, A. Stolarski, and D. G. Grier, “Influence of nonconservative optical forces on the dynamics of optically trapped colloidal spheres: the fountain of probability”, [Physical Review Letters](#) **101**, 128301 (2008).
- [31] T. Li, Ph.D. Thesis, UT Austin (2013).
- [32] J. Gieseler, B. Deutsch, R. Quidant, and L. Novotny, “Subkelvin parametric feedback cooling of a laser-trapped nanoparticle”, [Physical Review Letters](#) **109**, 103603 (2012).
- [33] J. Vovrosh, M. Rashid, D. Hempston, Bateman, M. Paternostro, and H. Ulbricht, “Parametric feedback cooling of levitated optomechanics in a parabolic mirror trap”, [Journal of the Optical Society of America B](#) **34**, 1421 (2017).
- [34] F. Monteiro, W. Li, G. Afek, C. Li, M. Mossman, and D. C. Moore, “Force and acceleration sensing with optically levitated nanogram masses at microkelvin temperatures”, [Physical Review A](#) **101**, 053835 (2020).

- [35] F. Tebbenjohanns, M. L. Mattana, M. Rossi, M. Frimmer, and L. Novotny, “Quantum control of a nanoparticle optically levitated in cryogenic free space”, *Nature* **595**, 378–382 (2021).
- [36] U. Delić, M. Reisenbauer, K. Dare, D. Grass, V. Vuletić, N. Kiesel, and M. Aspelmeyer, “Cooling of a levitated nanoparticle to the motional quantum ground state”, *Science* **367**, 892–895 (2020).
- [37] A. Kawasaki, A. Fieguth, N. Priel, C. P. Blakemore, D. Martin, and G. Gratta, “High sensitivity, levitated microsphere apparatus for short-distance force measurements”, *Review of Scientific Instruments* **91**, 083201 (2020).
- [38] A. D. Rider, C. P. Blakemore, A. Kawasaki, N. Priel, S. Roy, and G. Gratta, “Electrically driven, optically levitated microscopic rotors”, *Physical Review A* **99**, 041802 (2019).
- [39] Kimball Physics, [<https://www.kimballphysics.com/>].
- [40] MKS Instruments, [<https://www.mksinst.com/>].
- [41] E. R. I. Abraham and E. A. Cornell, “Teflon feedthrough for coupling optical fibers into ultrahigh vacuum systems”, *Applied Optics* **37**, 1762–1763 (1998).
- [42] Pfeiffer Vacuum, [<https://www.pfeiffer-vacuum.com/>].
- [43] Kurt J. Lesker Company, [<https://www.lesker.com/>].
- [44] A. E. Siegman, “How to (maybe) measure laser beam quality”, in *DPSS (Diode Pumped Solid State) Lasers: Applications and Issues* (1998), MQ1.
- [45] C. Xu, Y. Zi, A. C. Wang, H. Zou, Y. Dai, X. He, P. Wang, Y. Wang, P. Feng, D. Li, and Z. L. Wang, “On the electron-transfer mechanism in the contact-electrification effect”, *Advanced Materials* **30**, 1706790 (2018).
- [46] A. A. Geraci, S. B. Papp, and J. Kitching, “Short-range force detection using optically cooled levitated microspheres”, *Physical Review Letters* **105**, 101101 (2010).
- [47] M. Frimmer, K. Luszcz, S. Ferreira, V. Jain, E. Hebestreit, and L. Novotny, “Controlling the net charge on a nanoparticle optically levitated in vacuum”, *Physical Review A* **95**, 061801 (2017).

- [48] F. Ricci, M. T. Cuairan, G. P. Conangla, A. W. Schell, and R. Quidant, “Accurate mass measurement of a levitated nanomechanical resonator for precision force sensing”, *Nano Letters* **19**, 6711 (2019).
- [49] G. P. Conangla, F. Ricci, M. T. Cuairan, A. W. Schell, N. Meyer, and R. Quidant, “Optimal feedback cooling of a charged levitated nanoparticle with adaptive control”, *Physical Review Letters* **122**, 223602 (2019).
- [50] C. P. Blakemore, [https://github.com/charlesblakemore/opt_lev_analysis/].
- [51] G. Binnig, C. F. Quate, and C. Gerber, “Atomic force microscope”, *Physical Review Letters* **56**, 930–933 (1986).
- [52] A. I. Bishop, T. A. Nieminen, N. R. Heckenberg, and H. Rubinsztein-Dunlop, “Optical microrheology using rotating laser-trapped particles”, *Physical Review Letters* **92**, 198104 (2004).
- [53] D. E. Chang, C. A. Regal, S. B. Papp, D. J. Wilson, J. Ye, O. Painter, H. J. Kimble, and P. Zoller, “Cavity opto-mechanics using an optically levitated nanosphere”, *Proceedings of the National Academy of Sciences of the United States of America* **107**, 1005–1010 (2010).
- [54] Z. Q. Yin, A. A. Geraci, and T. Li, “Optomechanics of levitated dielectric particles”, *International Journal of Modern Physics A* **B27**, 1330018 (2013).
- [55] P. Asenbaum, S. Kuhn, S. Nimmrichter, U. Sezer, and M. Arndt, “Cavity cooling of free silicon nanoparticles in high vacuum”, *Nature Communications* **4**, 2743, 2743 (2013).
- [56] G. Ranjit, D. P. Atherton, J. H. Stutz, M. Cunningham, and A. A. Geraci, “Attonewton force detection using microspheres in a dual-beam optical trap in high vacuum”, *Physical Review A* **91**, 051805 (2015).
- [57] J. Millen, P. Z. G. Fonseca, T. Mavrogordatos, T. S. Monteiro, and P. F. Barker, “Cavity cooling a single charged levitated nanosphere”, *Physical Review Letters* **114**, 123602 (2015).
- [58] V. Jain, J. Gieseler, C. Moritz, C. Dellago, R. Quidant, and L. Novotny, “Direct measurement of photon recoil from a levitated nanoparticle”, *Physical Review Letters* **116**, 243601 (2016).

- [59] T. M. Hoang, Y. Ma, J. Ahn, J. Bang, F. Robicieux, Z. Q. Yin, and T. Li, “Torsional optomechanics of a levitated nonspherical nanoparticle”, *Physical Review Letters* **117**, 123604 (2016).
- [60] G. Ranjit, M. Cunningham, K. Casey, and A. A. Geraci, “Zeptonewton force sensing with nanospheres in an optical lattice”, *Physical Review A* **93**, 053801 (2016).
- [61] M. Mazilu, Y. Arita, T. Vettenburg, J. M. Auñón, E. M. Wright, and K. Dholakia, “Orbital-angular-momentum transfer to optically levitated microparticles in vacuum”, *Physical Review A* **94**, 053821 (2016).
- [62] P. Z. G. Fonseca, E. B. Aranas, J. Millen, T. S. Monteiro, and P. F. Barker, “Nonlinear dynamics and strong cavity cooling of levitated nanoparticles”, *Physical Review Letters* **117**, 173602 (2016).
- [63] F. Monteiro, S. Ghosh, A. G. Fine, and D. C. Moore, “Optical levitation of 10-ng spheres with nano- g acceleration sensitivity”, *Physical Review A* **96**, 063841 (2017).
- [64] D. Hempston, J. Vovrosh, M. Toroš, G. Winstone, M. Rashid, and H. Ulbricht, “Force sensing with an optically levitated charged nanoparticle”, *Applied Physics Letters* **111**, 133111 (2017).
- [65] F. Monteiro, S. Ghosh, E. C. van Assendelft, and D. C. Moore, “Optical rotation of levitated spheres in high vacuum”, *Physical Review A* **97**, 051802(R) (2018).
- [66] R. Diehl, E. Hebestreit, R. Reimann, F. Tebbenjohanns, M. Frimmer, and L. Novotny, “Optical levitation and feedback cooling of a nanoparticle at subwavelength distances from a membrane”, *Physical Review A* **98**, 013851 (2018).
- [67] D. S. Ether., L. B. Pires, S. Umrath, D. Martinez, Y. Ayala, B. Pontes, G. R. de S. Araújo, S. Frases, G.-L. Ingold, F. S. S. Rosa, N. B. Viana, H. M. Nussenzveig, and P. A. M. Neto, “Probing the casimir force with optical tweezers”, *Europhysics Letters* **112**, 44001 (2015).
- [68] G. Winstone, R. Bennett, M. Rademacher, M. Rashid, S. Buhmann, and H. Ulbricht, “Direct measurement of the electrostatic image force of a levitated charged nanoparticle close to a surface”, *Physical Review A* **98**, 053831 (2018).
- [69] Q. Wang, A. D. Rider, D. C. Moore, C. P. Blakemore, L. Cao, and G. Gratta, “A density staggered cantilever for micron length gravity probing”, *IEEE 67th Electronic Components and Technology Conference (ECTC)* **67**, 1773–1778 (2017).

- [70] Q. A. Turchette, D. Kielpinski, B. E. King, D. Leibfried, D. M. Meekhof, C. J. Myatt, M. A. Rowe, C. A. Sackett, C. S. Wood, W. M. Itano, C. Monroe, and D. J. Wineland, “Heating of trapped ions from the quantum ground state”, *Physical Review A* **61**, 063418 (2000).
- [71] G. H. Low, P. F. Herskind, and I. L. Chuang, “Finite-geometry models of electric field noise from patch potentials in ion traps”, *Physical Review A* **84**, 053425 (2011).
- [72] J. Chiaverini, S. J. Smullin, A. A. Geraci, D. M. Weld, and A. Kapitulnik, “New experimental constraints on non-newtonian forces below 100 μm ”, *Physical Review Letters* **90**, 151101 (2003).
- [73] A. B. Churnside and T. T. Perkins, “Ultrastable atomic force microscopy: improved force and positional stability”, *FEBS Letters* **588**, 3621–3630 (2014).
- [74] F. Rossi and G. I. Opat, “Observations of the effects of adsorbates on patch potentials”, *Journal of Physics D: Applied Physics* **25**, 1349 (1992).
- [75] N. A. Robertson, J. R. Blackwood, S. Buchman, R. L. Byer, J. Camp, D. Gill, J. Hanson, S. Williams, and P. Zhou, “Kelvin probe measurements: investigations of the patch effect with applications to st-7 and lisa”, *Classical and Quantum Gravity* **23**, 2665 (2006).
- [76] Bangs Laboratories, [<http://www.bangslabs.com/>].
- [77] C. C. Speake and C. Trenkel, “Forces between conducting surfaces due to spatial variations of surface potential”, *Physical Review Letters* **90**, 160403 (2003).
- [78] J. L. Garrett, D. Somers, and J. N. Munday, “The effect of patch potentials in casimir force measurements determined by heterodyne kelvin probe force microscopy”, *Journal of Physics: Condensed Matter* **27**, 214012 (2015).
- [79] A. Arvanitaki and A. A. Geraci, “Detecting high-frequency gravitational waves with optically levitated sensors”, *Physical Review Letters* **110**, 071105 (2013).
- [80] O. Romero-Isart, M. L. Juan, R. Quidant, and J. I. Cirac, “Toward quantum superposition of living organisms”, *New Journal of Physics* **12**, 033015, 033015 (2010).
- [81] O. Romero-Isart, A. C. Pflanzer, F. Blaser, R. Kaltenbaek, N. Kiesel, M. Aspelmeyer, and J. I. Cirac, “Large quantum superpositions and interference of massive nanometer-sized objects”, *Physical Review Letters* **107**, 020405 (2011).

- [82] J. Chaste, A. Eichler, J. Moser, G. Ceballos, R. Rurali, and A. Bachtold, “A nanomechanical mass sensor with yoctogram resolution”, *Nature Nanotechnology* **7**, 301 (2012).
- [83] G. Göbel, T. Wriedt, and K. Bauckhage, “Periodic drag force and particle size measurement in a double ring electrodynamic trap”, *Review of Scientific Instruments* **68**, 3046–3052 (1997).
- [84] C. Heinisch, J. B. Wills, J. P. Reid, T. Tschudi, and C. Tropea, “Temperature measurement of single evaporating water droplets in a nitrogen flow using spontaneous raman scattering”, *Physical Chemistry Chemical Physics* **11**, 9720–9728 (2009).
- [85] G. Hargreaves, N.-O. A. Kwamena, Y. H. Zhang, J. R. Butler, S. Rushworth, S. L. Clegg, and J. P. Reid, “Measurements of the equilibrium size of supersaturated aqueous sodium chloride droplets at low relative humidity using aerosol optical tweezers and an electrodynamic balance”, *Journal of Physical Chemistry A* **114**, 1806–1815 (2010).
- [86] U. K. Krieger, C. Marcolli, and J. P. Reid, “Exploring the complexity of aerosol particle properties and processes using single particle techniques”, *Chemical Society Reviews* **41**, 6631–6662 (2012).
- [87] M. Singh, Y. S. Mayya, J. Gaware, and R. M. Thaokar, “Levitation dynamics of a collection of charged droplets in an electrodynamic balance”, *Journal of Applied Physics* **121**, 054503 (2017).
- [88] microParticles GmbH, [<https://microparticles.de/en>].
- [89] Planotec and Pelcotec, [https://www.tedpella.com/calibration_html/SEM_Magnification_Calibration_Standards.htm].
- [90] N. K. Pavlis, S. A. Holmes, S. C. Kenyon, and J. K. Factor, “An earth gravitational model to degree 2160: EGM2008”, in (Apr. 2008).
- [91] Thorlabs, Inc., [<https://www.thorlabs.com/>].
- [92] F. Laermer and A. Schilp, “Method of anisotropic etching of silicon”, *US Patent 6 531 068* (2003).
- [93] Corning Incorporated, [<https://www.corning.com/>].
- [94] S. Li, Q. Wan, Z. Qin, Y. Fu, and Y. Gu, “Understanding stöber silica’s pore characteristics measured by gas adsorption”, *Langmuir* **31**, 824–832 (2015).

- [95] D. A. Kurdyukov, D. A. Eurov, D. A. Kirilenko, V. V. Sokolov, and V. G. Golubev, “Tailoring the size and microporosity of stöber silica particles”, [Microporous and Mesoporous Materials](#) **258**, 205–210 (2018).
- [96] microParticles GmbH, personal communication, Mar. 19, 2020.
- [97] I. H. Malitson, “Interspecimen comparison of the refractive index of fused silica”, [Journal of the Optical Society of America](#) **55**, 1205–1209 (1965).
- [98] D. Goldwater, B. A. Stickler, L. Martinetz, T. E. Northup, K. Hornberger, and J. Millen, “Levitated electromechanics: all-electrical cooling of charged nano- and micro-particles”, [Quantum Science and Technology](#) **4**, 024003 (2019).
- [99] O. Romero-Isart, “Quantum superposition of massive objects and collapse models”, [Physical Review A](#) **84**, 052121 (2011).
- [100] B. R. Slezak, C. W. Lewandowski, J. Hsu, and B. D’Urso, “Cooling the motion of a silica microsphere in a magneto-gravitational trap in ultra-high vacuum”, [New Journal of Physics](#) **20**, 063028 (2018).
- [101] A. T. M. Anishur Rahman and P. F. Barker, “Laser refrigeration, alignment and rotation of levitated yb^{3+} nanocrystals”, [Nature Photonics](#) **11**, 634–638 (2017).
- [102] J. Ahn, Z. Xu, J. Bang, Y. Deng, T. M. Hoang, Q. Han, R. Ma, and T. Li, “Optically levitated nanodumbbell torsion balance and ghz nanomechanical rotor”, [Physical Review Letters](#) **121**, 033603 (2018).
- [103] R. Reimann, M. Doderer, E. Hebestreit, R. Diehl, M. Frimmer, D. Windey, F. Tebbenjohanns, and L. Novotny, “Ghz rotation of an optically trapped nanoparticle in vacuum”, [Physical Review Letters](#) **121**, 033602 (2018).
- [104] S. Kuhn, B. A. Stickler, A. Kosloff, F. Patolsky, K. Hornberger, M. Arndt, and J. Millen, “Optically driven ultra-stable nanomechanical rotor”, [Nature Communications](#) **8**, 1670 (2017).
- [105] Y. Arita, A. W. McKinley, M. Mazilu, H. Rubinsztein-Dunlop, and K. Dholakia, “Picoliter rheology of gaseous media using a rotating optically trapped birefringent microparticle”, [Analytical Chemistry](#) **83**, 8855–8858 (2011).
- [106] Y. Arita, E. M. Wright, and K. Dholakia, “Optical binding of two cooled micro-gyroscopes levitated in vacuum”, [Optica](#) **5**, 910–917 (2018).

- [107] Y. Arita, M. Mazilu, and K. Dholakia, “Laser-induced rotation and cooling of a trapped microgyroscope in vacuum”, *Nature Communications* **4**, 2374 (2013).
- [108] H. He, M. E. J. Friese, N. R. Heckenberg, and H. Rubinsztein-Dunlop, “Direct observation of transfer of angular momentum to absorptive particles from a laser beam with a phase singularity”, *Physical Review Letters* **75**, 826–829 (1995).
- [109] J. S. Bennett, L. J. Gibson, R. M. Kelly, E. Brousse, B. Baudisch, D. Preece, T. A. Nieminen, T. Nicholson, N. R. Heckenberg, and H. Rubinsztein-Dunlop, “Spatially-resolved rotational microrheology with an optically-trapped sphere”, *Scientific Reports* **3**, 1759 (2013).
- [110] W. Li, N. Li, Y. Shen, Z. Fu, H. Su, and H. Hu, “Dynamic analysis and rotation experiment of an optical-trapped microsphere in air”, *Applied Optics* **57**, 823–828 (2018).
- [111] M. Rashid, M. Toro š, A. Setter, and H. Ulbricht, “Precession motion in levitated optomechanics”, *Physical Review Letters* **121**, 253601 (2018).
- [112] P. Nagornykh, J. E. Coppock, J. P. J. Murphy, and B. E. Kane, “Optical and magnetic measurements of gyroscopically stabilized graphene nanoplatelets levitated in an ion trap”, *Physical Review B* **96**, 035402 (2017).
- [113] J. E. Coppock, P. Nagornykh, J. P. J. Murphy, and B. E. Kane, “Phase locking of the rotation of a graphene nanoplatelet to an rf electric field in a quadrupole ion trap”, in *Optical trapping and optical micromanipulation XIII*, Vol. 9922 (2016), pp. 56–65.
- [114] J. F. O’Hanlon, *A user’s guide to vacuum technology* (John Wiley & Sons, Hoboken, NJ, 2003).
- [115] Leybold GmbH, [<https://www.leybold.com/us/>].
- [116] E. O. Lawrence, “The ionization of atoms by electron impact”, *Physical Review* **28**, 947–961 (1926).
- [117] I. Langmuir and H. A. Jones, “Collisions between electrons and gas molecules”, *Physical Review* **31**, 357–404 (1928).
- [118] R. A. Gulyaev, “Ionization of neutral atoms by electron impact”, *Soviet Astronomy* **10**, 761 (1967).
- [119] J. C. Maxwell, “Xiii. the bakerian lecture.-on the viscosity or internal friction of air and other gases”, *Philosophical Transactions of the Royal Society* **156**, 249 (1866).

- [120] J. W. Beams, D. M. Spitzer, and J. P. Wade, “Spinning rotor pressure gauge”, [Review of Scientific Instruments](#) **33**, 151 (1962).
- [121] G. Comsa, J. K. Fremerey, B. E. Lindenau, G. Messer, and P. Röhl, “Calibration of a spinning rotor gas friction gauge against a fundamental vacuum pressure standard”, [Journal of Vacuum Science & Technology](#) **17**, 642 (1980).
- [122] K. E. McCulloh, “Calibration of molecular drag vacuum gages”, [Journal of Vacuum Science & Technology A](#) **1**, 168 (1980).
- [123] A. Cavalleri, G. Ciani, R. Dolesi, M. Hueller, D. Nicolodi, D. Tombolato, S. Vitale, P. Wass, and W. Weber, “Gas damping force noise on a macroscopic test body in an infinite gas reservoir”, [Physics Letters A](#) **374**, 3365–3369 (2010).
- [124] L. Martinetz, K. Hornberger, and B. A. Stickler, “Gas-induced friction and diffusion of rigid rotors”, [Physical Review E](#) **97**, 052112 (2018).
- [125] J. K. Fremerey, “The spinning rotor gauge”, [Journal of Vacuum Science & Technology A](#) **3**, 1715 (1985).
- [126] R. F. Chang and P. J. Abbott, “Factors affecting the reproducibility of the accommodation coefficient of the spinning rotor gauge”, [Journal of Vacuum Science & Technology A](#) **25**, 1567 (2007).
- [127] M. E. J. Friese, T. A. Nieminen, N. R. Heckenberg, and H. Rubinsztein-Dunlop, “Optical alignment and spinning of laser-trapped microscopic particles”, [Nature](#) **394**, 348 (1998).
- [128] R. C. Jones, “A new calculus for the treatment of optical systems. description and discussion of the calculus”, [Journal of the Optical Society of America](#) **31**, 488–493 (1941).
- [129] A. La Porta and M. D. Wang, “Optical torque wrench: angular trapping, rotation, and torque detection of quartz microparticles”, [Physical Review Letters](#) **92**, 190801 (2004).
- [130] R. F. Berg and J. A. Fedchak, “NIST calibration services for spinning rotor gauge calibrations”, [NIST Special Publications](#) **10**, 11505–11510 (2015).
- [131] P. A. Sturrock, “Stochastic acceleration”, [Physical Review](#) **141**, 186–191 (1966).

- [132] J. Millen, T. Deesuwan, P. Barker, and J. Anders, “Nanoscale temperature measurements using non-equilibrium brownian dynamics of a levitated nanosphere”, *Nature Nanotechnology* **9**, 425–429 (2014).
- [133] Y. Arita, J. M. Richards, M. Mazilu, G. C. Spalding, S. E. Skelton Spesyvtseva, D. Craig, and K. Dholakia, “Rotational dynamics and heating of trapped nanovaterite particles”, *ACS Nano* **10**, 11505–11510 (2016).
- [134] E. Hebestreit, R. Reimann, M. Frimmer, and L. Novotny, “Measuring the internal temperature of a levitated nanoparticle in high vacuum”, *Physical Review A* **97**, 043803 (2018).
- [135] S. Dittmann, B. E. Lindenau, and C. R. Tilford, “The molecular drag gauge as a calibration standard”, *Journal of Vacuum Science & Technology A* **7**, 3356 (1989).
- [136] Z. Lu and H. Owens, “A method to improve the quality of silica nanoparticles (snps) over increasing storage durations”, *Journal of Nanoparticle Research* **20**, 213–213 (2018).
- [137] X. Dai, K. D. Rasamani, G. Hall, R. Makrypodis, and Y. Sun, “Geometric symmetry of dielectric antenna influencing light absorption in quantum-sized metal nanocrystals: a comparative study”, *Frontiers in Chemistry* **6**, 494 (2018).
- [138] I. Newton, *Philosophiae naturalis principia mathematica* (Royal Society, London, 1687).
- [139] A. Einstein, “Die grundlage der allgemeinen relativitätstheorie”, *Annalen der Physik* **354**, 769–822 (1916).
- [140] S. R. Beane, “On the importance of testing gravity at distances less than 1cm”, *General Relativity and Gravitation* **29**, 945–951 (1997).
- [141] E. G. Adelberger, B. R. Heckel, and A. E. Nelson, “Tests of the gravitational inverse-square law”, *Annual Review of Nuclear and Particle Science* **53**, 77–121 (2003).
- [142] G. Arnison et al. (UA1 Collaboration), “Experimental observation of isolated large transverse energy electrons with associated missing energy at s=540 gev”, *Physics Letters B* **122**, 103–116 (1983).
- [143] G. Arnison et al. (UA1 Collaboration), “Experimental observation of lepton pairs of invariant mass around 95 gev/c² at the cern sps collider”, *Physics Letters B* **126**, 398–410 (1983).

- [144] M. Banner et al. (UA2 Collaboration), “Observation of single isolated electrons of high transverse momentum in events with missing transverse energy at the cern pp collider”, *Physics Letters B* **122**, 476–485 (1983).
- [145] P. Bagnaia et al. (UA2 Collaboration), “Evidence for $z_0 \rightarrow e^+e^-$ at the cern pp collider”, *Physics Letters B* **129**, 130–140 (1983).
- [146] S. Alam et al., “The clustering of galaxies in the completed sdss-iii baryon oscillation spectroscopic survey: cosmological analysis of the dr12 galaxy sample”, *Monthly Notices of the Royal Astronomical Society* **470**, 2617–2652 (2017).
- [147] N. Aghanim et al. (Planck Collaboration), “Planck 2018 results - vi. cosmological parameters”, *Astronomy and Astrophysics* **641**, A6 (2020).
- [148] A. S. Goldhaber and M. M. Nieto, “Terrestrial and extraterrestrial limits on the photon mass”, *Reviews of Modern Physics* **43**, 277–296 (1971).
- [149] R. Lakes, “Experimental limits on the photon mass and cosmic magnetic vector potential”, *Physical Review Letters* **80**, 1826–1829 (1998).
- [150] L.-C. Tu and J. Luo, “Experimental tests of coulomb’s law and the photon rest mass”, *Metrologia* **41**, S136–S146 (2004).
- [151] A. E. Nelson and J. Scholtz, “Dark light, dark matter, and the misalignment mechanism”, *Physical Review D* **84**, 103501 (2011).
- [152] P. W. Graham, J. Mardon, and S. Rajendran, “Vector dark matter from inflationary fluctuations”, *Physical Review D* **93**, 103520 (2016).
- [153] K. Aoki and S. Mukohyama, “Massive gravitons as dark matter and gravitational waves”, *Physical Review D* **94**, 024001 (2016).
- [154] R. D. Peccei, “The strong cp problem and axions”, in *Axions: theory, cosmology, and experimental searches*, edited by M. Kuster, G. Raffelt, and B. Beltrán (Springer Berlin Heidelberg, Berlin, Heidelberg, 2008), pp. 3–17.
- [155] N. Arkani-Hamed, S. Dimopoulos, and G. Dvali, “The hierarchy problem and new dimensions at a millimeter”, *Physics Letters B* **429**, 263–272 (1998).
- [156] I. Antoniadis, N. Arkani-Hamed, S. Dimopoulos, and G. Dvali, “New dimensions at a millimeter to a fermi and superstrings at a tev”, *Physics Letters B* **436**, 257–263 (1998).

- [157] L. Randall and R. Sundrum, “Large mass hierarchy from a small extra dimension”, *Physical Review Letters* **83**, 3370–3373 (1999).
- [158] L. Randall and R. Sundrum, “An alternative to compactification”, *Physical Review Letters* **83**, 4690–4693 (1999).
- [159] R. Sundrum, “Fat gravitons, the cosmological constant and submillimeter tests”, *Physical Review D* **69**, 044014 (2004).
- [160] E. G. Adelberger, J. H. Gundlach, B. R. Heckel, S. Hoedl, and S. Schlamminger, “Torsion balance experiments: a low-energy frontier of particle physics”, *Progress in Particle and Nuclear Physics* **62**, 102–134 (2009).
- [161] J. E. Moody and F. Wilczek, “New macroscopic forces?”, *Physical Review D* **30**, 130–138 (1984).
- [162] G. Aad et al. (ATLAS Collaboration), “Search for new phenomena in events with an energetic jet and missing transverse momentum in pp collisions at $\sqrt{s} = 13$ TeV with the atlas detector”, *Physical Review D* **103**, 112006 (2021).
- [163] A. M. Sirunyan et al. (CMS Collaboration), “Music: a model-unspecific search for new physics in proton–proton collisions at $\sqrt{s} = 13$ TeV”, *The European Physical Journal C* **81**, 629 (2021).
- [164] E. Verlinde and H. Verlinde, “RG-flow, gravity and the cosmological constant”, *Journal of High Energy Physics* **2000**, 034–034 (2000).
- [165] A. G. Riess et al., “A 2.4% determination of the local value of the hubble constant”, *The Astrophysical Journal* **826**, 56 (2016).
- [166] B. Ratra and P. J. E. Peebles, “Cosmological consequences of a rolling homogeneous scalar field”, *Physical Review D* **37**, 3406–3427 (1988).
- [167] P. W. Graham, D. E. Kaplan, and S. Rajendran, “Cosmological relaxation of the electroweak scale”, *Physical Review Letters* **115**, 221801 (2015).
- [168] A. Joyce, B. Jain, J. Khoury, and M. Trodden, “Beyond the cosmological standard model”, *Physics Reports* **568**, 1–98 (2015).
- [169] J. Khoury and A. Weltman, “Chameleon fields: awaiting surprises for tests of gravity in space”, *Physical Review Letters* **93**, 171104 (2004).

- [170] J. Khoury and A. Weltman, “Chameleon cosmology”, *Physical Review D* **69**, 044026 (2004).
- [171] P. Brax, C. van de Bruck, A. C. Davis, J. Khoury, and A. Weltman, “Chameleon dark energy”, *AIP Conference Proceedings* **736**, edited by C. J. A. P. Martins, P. P. Avelino, M. S. Costa, K. Mack, M. F. Mota, and M. Parry, 105–110 (2004).
- [172] B. Elder, J. Khoury, P. Haslinger, M. Jaffe, H. Müller, and P. Hamilton, “Chameleon dark energy and atom interferometry”, *Physical Review D* **94**, 044051 (2016).
- [173] C. Burrage, E. J. Copeland, and E. A. Hinds, “Probing dark energy with atom interferometry”, *Journal of Cosmology and Astroparticle Physics* **1503**, 042 (2015).
- [174] A. Upadhye, S. S. Gubser, and J. Khoury, “Unveiling chameleon fields in tests of the gravitational inverse-square law”, *Physical Review D* **74**, 104024 (2006).
- [175] A. Upadhye, “Dark energy fifth forces in torsion pendulum experiments”, *Physical Review D* **86**, 102003 (2012).
- [176] C. Wetterich, “Cosmology and the fate of dilatation symmetry”, *Nuclear Physics B* **302**, 668–696 (1988).
- [177] D. B. Kaplan and M. B. Wise, “Couplings of a light dilaton and violations of the equivalence principle”, *Journal of High Energy Physics* **2000**, 037–037 (2000).
- [178] T. C. Scott, X. Zhang, R. B. Mann, and G. J. Fee, “Canonical reduction for dilatonic gravity in $3 + 1$ dimensions”, *Physical Review D* **93**, 084017 (2016).
- [179] S. Dimopoulos and G. F. Giudice, “Macroscopic forces from supersymmetry”, *Physics Letters B* **379**, 105–114 (1996).
- [180] G. Dvali, G. Gabadadze, and G. Senjanović, “Constraints on extra time dimensions”, in *The many faces of the superworld* (2000), pp. 525–532.
- [181] T. Kaluza, “On the unification problem in physics”, *International Journal of Modern Physics D* **27**, German original in: *Sitzungsberichte der Preussischen Akademie der Wissenschaften, Berlin (Physikalisch-Mathematische Klasse)* 1921, 966–972, 1870001 (2018).
- [182] O. Klein, “The atomicity of electricity as a quantum theory law”, *Nature* **118**, 516–516 (1926).

- [183] C. Brans and R. H. Dicke, “Mach’s principle and a relativistic theory of gravitation”, [Physical Review](#) **124**, 925–935 (1961).
- [184] G. Dvali, G. Gabadadze, and M. Shifman, “Diluting the cosmological constant in infinite volume extra dimensions”, [Physical Review D](#) **67**, 044020 (2003).
- [185] C. W. Misner, K. S. Thorne, and J. A. Wheeler, *Gravitation* (W. H. Freeman, San Francisco, 1973).
- [186] K. Hinterbichler, “Theoretical aspects of massive gravity”, [Reviews of Modern Physics](#) **84**, 671–710 (2012).
- [187] C. de Rham, “Massive gravity”, [Living Reviews in Relativity](#) **17**, 7 (2014).
- [188] M. Fierz, W. E. Pauli, and P. A. M. Dirac, “On relativistic wave equations for particles of arbitrary spin in an electromagnetic field”, [Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences](#) **173**, 211–232 (1939).
- [189] C. de Rham and G. Gabadadze, “Generalization of the fierz-pauli action”, [Physical Review D](#) **82**, 044020 (2010).
- [190] S. F. Hassan and R. A. Rosen, “Bimetric gravity from ghost-free massive gravity”, [Journal of High Energy Physics](#) **2012**, 126 (2012).
- [191] P. J. Shelus, “Mlrs: a lunar/artificial satellite laser ranging facility at the mcdonald observatory”, [IEEE Transactions on Geoscience and Remote Sensing](#) **GE-23**, 385–390 (1985).
- [192] E. Samain, J. F. Mangin, C. Veillet, J. M. Torre, P. Fridelance, J. E. Chabaudie, D. Féraudy, M. Glentzlin, J. Pham Van, M. Furia, A. Journet, and G. Vigouroux, “Millimetric lunar laser ranging at oca (observatoire de la côte d’azur)”, [Astronomy and Astrophysics Supplement Series](#) **130**, 235–244 (1998).
- [193] C. D. Hoyle, D. J. Kapner, B. R. Heckel, E. G. Adelberger, J. H. Gundlach, U. Schmidt, and H. E. Swanson, “Submillimeter tests of the gravitational inverse-square law”, [Physical Review D](#) **70**, 042004 (2004).
- [194] A. A. Geraci, S. J. Smullin, D. M. Weld, J. Chiaverini, and A. Kapitulnik, “Improved constraints on non-newtonian forces at 10 microns”, [Physical Review D](#) **78**, 022002 (2008).

- [195] A. O. Sushkov, W. J. Kim, D. A. R. Dalvit, and S. K. Lamoreaux, “New experimental limits on non-newtonian forces in the micrometer range”, [Physical Review Letters](#) **107**, 171101 (2011).
- [196] Y.-J. Chen, W. K. Tham, D. E. Krause, D. López, E. Fischbach, and R. S. Decca, “Stronger limits on hypothetical yukawa interactions in the 30-8000 nm range”, [Physical Review Letters](#) **116**, 221102 (2016).
- [197] S.-Q. Yang, B.-F. Zhan, Q.-L. Wang, C.-G. Shao, L.-C. Tu, W.-H. Tan, and J. Luo, “Test of the gravitational inverse square law at millimeter ranges”, [Physical Review Letters](#) **108**, 081101 (2012).
- [198] J. G. Lee, E. G. Adelberger, T. S. Cook, S. M. Fleischer, and B. R. Heckel, “New test of the gravitational $1/r^2$ law at separations down to $52 \mu\text{m}$ ”, [Physical Review Letters](#) **124**, 101101 (2020).
- [199] W.-H. Tan, A.-B. Du, W.-C. Dong, S.-Q. Yang, C.-G. Shao, S.-G. Guan, Q.-L. Wang, B.-F. Zhan, P.-S. Luo, L.-C. Tu, and J. Luo, “Improvement for testing the gravitational inverse-square law at the submillimeter range”, [Physical Review Letters](#) **124**, 051301 (2020).
- [200] H. Cavendish, “Experiments to determine the density of the earth”, *Philosophical Transactions of the Royal Society of London*, (part II), 469 (1798).
- [201] R. Spero, J. K. Hoskins, R. Newman, J. Pellam, and J. Schultz, “Test of the gravitational inverse-square law at laboratory distances”, [Physical Review Letters](#) **44**, 1645–1648 (1980).
- [202] J. K. Hoskins, R. D. Newman, R. Spero, and J. Schultz, “Experimental tests of the gravitational inverse-square law for mass separations from 2 to 105 cm”, [Physical Review D](#) **32**, 3084–3095 (1985).
- [203] M. V. Moody and H. J. Paik, “Gauss’s law test of gravity at short range”, [Physical Review Letters](#) **70**, 1195–1198 (1993).
- [204] 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”, [Physical Review Letters](#) **98**, 021101 (2007).

- [205] D. M. Weld, J. Xia, B. Cabrera, and A. Kapitulnik, “New apparatus for detecting micron-scale deviations from newtonian gravity”, *Physical Review D* **77**, 062006 (2008).
- [206] E. Tiesinga, P. J. Mohr, D. B. Newell, and B. N. Taylor, “CODATA recommended values of the fundamental physical constants: 2018”, *Reviews of Modern Physics* **93**, 025010 (2021).
- [207] B. Jain and J. Khoury, “Cosmological tests of gravity”, *Annals of Physics* **325**, 1479–1516 (2010).
- [208] D. F. Mota and D. J. Shaw, “Strongly coupled chameleon fields: new horizons in scalar field theory”, *Physical Review Letters* **97**, 151102 (2006).
- [209] W.-H. Tan, S.-Q. Yang, C.-G. Shao, J. Li, A.-B. Du, B.-F. Zhan, Q.-L. Wang, P.-S. Luo, L.-C. Tu, and J. Luo, “New test of the gravitational inverse-square law at the submillimeter range with dual modulation and compensation”, *Physical Review Letters* **116**, 131101 (2016).
- [210] P. Hamilton, M. Jaffe, P. Haslinger, Q. Simmons, H. Müller, and J. Khoury, “Atom-interferometry constraints on dark energy”, *Science* **349**, 849–851 (2015).
- [211] K. Li, M. Arif, D. G. Cory, R. Haun, B. Heacock, M. G. Huber, J. Nsofini, D. A. Pushin, P. Saggi, D. Sarenac, C. B. Shahi, V. Skavysh, W. M. Snow, and A. R. Young (The INDEX Collaboration), “Neutron limit on the strongly-coupled chameleon field”, *Physical Review D* **93**, 062001 (2016).
- [212] H. Lemmel, P. Brax, A. N. Ivanov, T. Jenke, G. Pignol, M. Pitschmann, T. Potocar, M. Wellenzohn, M. Zawisky, and H. Abele, “Neutron interferometry constrains dark energy chameleon fields”, *Physics Letters B* **743**, 310–314 (2015).
- [213] T. Jenke, G. Cronenberg, J. Burgdörfer, L. A. Chizhova, P. Geltenbort, A. N. Ivanov, T. Lauer, T. Lins, S. Rotter, H. Saul, U. Schmidt, and H. Abele, “Gravity resonance spectroscopy constrains dark energy and dark matter scenarios”, *Physical Review Letters* **112**, 151105 (2014).
- [214] N. Kiesel, F. Blaser, U. Delic, D. Grass, R. Kaltenbaek, and M. Aspelmeyer, “Cavity cooling of an optically levitated submicron particle”, *Proceedings of the National Academy of Sciences of the United States of America* **110**, 14180–14185 (2013).

- [215] A. A. Geraci, S. B. Papp, and J. Kitching, “Short-range force detection using optically cooled levitated microspheres”, *Physical Review Letters* **105**, 101101 (2010).
- [216] J. Gieseler, L. Novotny, and R. Quidant, “Thermal nonlinearities in a nanomechanical oscillator”, *Nature Physics* **9**, 806–810 (2013).
- [217] E. G. Adelberger, B. R. Heckel, S. A. Hoedl, C. D. Hoyle, D. J. Kapner, and A. Upadhye, “Particle physics implications of a recent test of the gravitational inverse square law”, *Physical Review Letters* **98**, 131104 (2007).
- [218] S. S. Wilks, “The large-sample distribution of the likelihood ratio for testing composite hypotheses”, *The Annals of Mathematical Statistics* **9**, 60–62 (1938).
- [219] G. Cowan, *Statistical data analysis*, Oxford science publications (Clarendon Press, 1998).
- [220] G. J. Simpson, C. F. Wilson, K.-H. Gericke, and R. N. Zare, “Coupled electrorotation: two proximate microspheres spin in registry with an ac electric field”, *ChemPhysChem* **3**, 416–423 (2002).
- [221] G. Cowan, K. Cranmer, E. Gross, and O. Vitells, “Asymptotic formulae for likelihood-based tests of new physics”, *European Physical Journal C* **71**, [Erratum: *Eur.Phys.J.C* **73**, 2501 (2013)], 1554 (2011).
- [222] R. Hooke, *Lectures de potentia restitutiva, or of spring explaining the power of springing bodies*, Cutlerian lecture (John Martyn, 1678).
- [223] C. P. Blakemore, A. D. Rider, and D. C. Moore, [https://github.com/stanfordbeads/opt_lev_controls].
- [224] National Instruments, [<https://www.ni.com/en-us/support/documentation/supplemental/18/labview-fpga-code-optimization.html>], Accessed: 2021-06-14.
- [225] S. M. Carroll, W. H. Press, and E. L. Turner, “The cosmological constant”, *Annual Review of Astronomy and Astrophysics* **30**, 499–542 (1992).
- [226] V. M. Slipher, “The radial velocity of the andromeda nebula”, *Lowell Observatory Bulletin* **1**, 56–57 (1913).
- [227] V. M. Slipher, “Nebulae”, *Proceedings of the American Philosophical Society* **56**, 403–409 (1917).

- [228] A. Friedman, “Über die Krümmung des Raumes”, *Zeitschrift für Physik* **10**, 377–386 (1922).
- [229] G. Lemaître, “Un univers homogène de masse constante et de rayon croissant rendant compte de la vitesse radiale des nébuleuses extra-galactiques”, *Annales de la Société Scientifique de Bruxelles* **47**, 49–59 (1927).
- [230] H. S. Leavitt, “1777 variables in the Magellanic clouds”, *Annals of Harvard College Observatory* **60**, 87–108.3 (1908).
- [231] H. S. Leavitt and E. C. Pickering, “Periods of 25 variable stars in the small Magellanic cloud”, *Harvard College Observatory Circular* **173**, 1–3 (1912).
- [232] E. Hubble, “A relation between distance and radial velocity among extra-galactic nebulae”, *Proceedings of the National Academy of Sciences* **15**, 168–173 (1929).
- [233] A. G. Riess et al. (Supernova Search Team), “Observational evidence from supernovae for an accelerating universe and a cosmological constant”, *The Astronomical Journal* **116**, 1009–1038 (1998).
- [234] S. Perlmutter et al. (Supernova Cosmology Project), “Measurements of Ω and Λ from 42 high redshift supernovae”, *The Astrophysical Journal* **517**, 565–586 (1999).
- [235] J. J. Condon and A. M. Matthews, “ Λ CDM cosmology for astronomers”, *Publications of the Astronomical Society of the Pacific* **130**, 073001 (2018).
- [236] P. J. E. Peebles, “Recombination of the primeval plasma”, *The Astrophysical Journal* **153**, 1 (1968).
- [237] Y. B. Zeldovich, V. G. Kurt, and R. A. Syunyaev, “Recombination of hydrogen in the hot model of the universe”, *Zhurnal Eksperimentalnoi i Teoreticheskoi Fiziki* **55**, 278–286 (1968).
- [238] P. J. E. Peebles, *Principles of physical cosmology* (Princeton University Press, Princeton, 1993).
- [239] W. Hu, D. Scott, N. Sugiyama, and M. White, “Effect of physical assumptions on the calculation of microwave background anisotropies”, *Physical Review D* **52**, 5498–5515 (1995).

- [240] U. Seljak, N. Sugiyama, M. White, and M. Zaldarriaga, “Comparison of cosmological boltzmann codes: are we ready for high precision cosmology?”, *Physical Review D* **68**, 083507 (2003).
- [241] W. A. Fendt, J. Chluba, J. A. Rubiño-Martín, and B. D. Wandelt, “RICO: A new approach for fast and accurate representation of the cosmological recombination history”, *The Astrophysical Journal Supplement Series* **181**, 627–638 (2009).
- [242] N. Y. Gnedin and O. Y. Gnedin, “Cosmological neutrino background revisited”, *The Astrophysical Journal* **509**, 11–15 (1998).
- [243] A. Faessler, R. Hodak, S. Kovalenko, and F. Simkovic, “Search for the cosmic neutrino background”, *Journal of Physics: Conference Series* **580**, 012040 (2015).
- [244] D. Baumann, F. Beutler, R. Flauger, D. Green, A. Slosar, M. Vargas-Magaña, B. Wallisch, and C. Yèche, “First constraint on the neutrino-induced phase shift in the spectrum of baryon acoustic oscillations”, *Nature Physics* **15**, 465–469 (2019).
- [245] D. J. Fixsen, “The temperature of the cosmic microwave background”, *The Astrophysical Journal* **707**, 916–920 (2009).
- [246] Stanford University, [<https://snfexfab.stanford.edu/>], Accessed: 2021-07-15.
- [247] W. Kern and D. A. Puotinen, “Cleaning solutions based on hydrogen peroxide for use in silicon semiconductor technology”, *RCA Review* **31**, 187 (1970).
- [248] BYU Integrated Microfabrication Lab, [<https://cleanroom.byu.edu/oxidetimecalc>], Accessed: 2021-07-15.